



National Library
of Canada

Bibliothèque nationale
du Canada

Canadian Theses Service

Services des thèses canadiennes

Ottawa, Canada
K1A 0N4

CANADIAN THESES

THÈSES CANADIENNES

NOTICE

The quality of this microfiche is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Previously copyrighted materials (journal articles, published tests, etc.) are not filmed.

Reproduction in full or in part of this film is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30.

**THIS DISSERTATION
HAS BEEN MICROFILMED
EXACTLY AS RECEIVED**

AVIS

La qualité de cette microfiche dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

Les documents qui font déjà l'objet d'un droit d'auteur (articles de revue, examens publiés, etc.) ne sont pas microfilmés.

La reproduction, même partielle, de ce microfilm est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30.

**LA THÈSE A ÉTÉ
MICROFILMÉE TELLE QUE
NOUS L'AVONS REÇUE**

PERCEPTION AND JUDGEMENT

by

Rockney Jacobsen

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES AND RESEARCH

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHILOSOPHY

EDMONTON, ALBERTA

SPRING, 1987

Permission has been granted to the National Library of Canada to microfilm this thesis and to lend or sell copies of the film.

The author (copyright owner) has reserved other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without his/her written permission.

L'autorisation a été accordée à la Bibliothèque nationale du Canada de microfilmer cette thèse et de prêter ou de vendre des exemplaires du film.

L'auteur (titulaire du droit d'auteur) se réserve les autres droits de publication; ni la thèse ni de longs extraits de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation écrite.

ISBN 0-315-37624-4

THE UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR Rockney Jacobsen

TITLE OF THESIS Perception and Judgement

DEGREE FOR WHICH THESIS WAS PRESENTED Ph.D.

YEAR THIS DEGREE GRANTED 1987

Permission is hereby granted to THE UNIVERSITY OF ALBERTA LIBRARY to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

Rocky A Jack
Rockney Jacobsen

PERMANENT ADDRESS: #105 266 Kingswood Dr
Kitchener, Ont
N2E 2K2

DATED March 30 1987

THE UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research, for acceptance, a thesis entitled PERCEPTION AND JUDGEMENT submitted by Rockney Jacobsen in partial fulfilment of the requirements for the degree of Doctor of Philosophy.

Richard Bosley

Supervisor

Fred Ditch

Lucia Hunte

Allen Carlson

Melba Mattie

Beaudet Mrs

Date

March 26 '87

to my parents,
Kermit and Shirley Jacobsen

ABSTRACT

Beginning with a criticism of the adequacy of a causal framework for discussing human action and with a minimal explication of a teleological framework, an account of the acts subordinated to language learning and use is developed and defended. This account emphasizes the nature of the semantic conventions of a language and the role of a community in establishing and maintaining those conventions.

The teleological framework is brought to bear to explain the intensional character of indirect discourse. Well known isomorphisms between indirect discourse and belief reports then make it possible to extend the account of acts of speech to explain belief acquisition, thus completing an account of the nature of judgement. The result of employing teleological notions is that our judgements are no longer viewed as mental artifacts, or products of mental activity, which serve to represent the world. Our judgements about the world, both what we say and what we think, are, rather, the achievements we reach through the coordinated exercise of skills acquired in a community setting.

Success in an inquiry depends upon skill in judging, but it also requires an opportunity adequate for bringing us to knowledge or justified belief. Perception provides that opportunity. An account of perception is defended which emphasizes its role in providing opportunities for

taking discriminatory action towards the world of public particulars. A consequence of the accounts of belief acquisition and language learning is that perception is non-epistemic.

The theories of judgement and perception presented appear to create special difficulties for understanding how perception contributes towards our arriving at opinions which we are justified in holding to be true. The conclusion defended, contrary to traditional empiricist theories of knowledge, is that perception does not contribute towards the evaluation of judgements. The role of perception in inquiry is remote and causal; it makes it possible for us to make judgements about the world, but does not contribute to our getting things right. Skepticism is blocked, and our beliefs are justified, by the work of linguistic communities in preserving semantic conventions for speech and thought. Perception belongs to the natural world, but the adjudication of contrary opinions and the conduct of inquiries belongs to the social world.

PREFACE

What follows is an attempt to provide an account of the role which perception plays in the acquisition of justified beliefs about the world. There is more than one role which it plays, and the way to completing an account of these roles traverses the interface of the philosophy of language, philosophical psychology and epistemology. The prime motivation for the particular account worked out here is the conviction that perception is nonepistemic - that is, that to perceive something is to do business outside any linguistic or conceptual framework. Perception does not, despite what has been widely held since Descartes, hand over belief, knowledge or ideas; it neither results in concepts or beliefs, nor does it presuppose them; nor, as I will argue, is any class of beliefs or of statements 'based on', 'derived from', 'verified by', 'confirmed by' or 'subject to the test of' perception. Our conceptual and epistemic resources, and the merits of our epistemic position in the world, are the result of the employment of what has traditionally been called the faculty of judgement.

It will be argued here that the operation of the faculties of perception and judgement are independent - that neither provides the materials or resources of the other; in particular, I will be concerned to argue that perception does not do the job traditionally attributed to it by empiricists - it does not maintain quality control (episte-

mically speaking) over the work of the faculty of judgement. But if, as I will maintain, the work of perception is neither causally nor logically involved in the work of judgement, then special difficulties arise in understanding how perception of the world can assist in the negotiations between conflicting empirical claims. The conclusion I will be defending is that perception does not assist in those negotiations; it is not complicitous in our decisions about what to say or think about the world. Neither what we see, nor the seeing of it, favours saying or thinking one thing rather than another.

The conclusion is not, however, of either the skeptical or the anti-realist persuasion. It is not skeptical because, although our beliefs cannot be tried by experience, they must face another tribunal - one which has the legitimate powers to give or withhold epistemic sanction. It is not anti-realist because, although perception neither aids nor hinders the acquisition of belief or the bestowal of epistemic rank on beliefs, the faculty which does do this work is effective only with the causal sustenance of a thoroughly realist form of perception.

The facts do not speak for themselves, but we speak on their behalf, and the prescriptive pressures exerted on what we say - the pressure to say one thing rather than another, and to modify and revise what has been said - are not exerted by perception, but by the linguistic communities to which we belong.

I have chosen a division into chapters which is designed to make the joints seem as natural as possible, rather than take the alternative course of aiming to keep the chapters uniform in length or as manageable for the reader as possible. As a result, the chapter-sized portions are, except for the first, too large to be consumed in one sitting. To compensate for this, I have divided the chapters into bite-sized sections. These sections are separated at natural divisions in the line of thought and, though a few are very short, tend toward a size suitable for a single sitting.

The argument proceeds in five stages. The first chapter argues for adopting a teleological framework for talking about human action in general; the thesis depends upon the application of this point to linguistic activity in particular. The position as a whole gets an anti-Cartesian flavor from this move, with the hope that our epistemic enterprises can be explicated within a framework sufficiently general to accommodate a wide variety of human activities. Several key notions needed for employing the teleological framework are explicated.

The second chapter discusses the conventionality of language and argues for the possibility of acquiring and using a conventional language while at the same time rejecting traditional abstractionist and innatist mechanisms. In this chapter, sections A - C give a minimally contentious characterization of what a natural language is and

then raise the main problem which arises in saying how a natural language is acquired. Sections D - F criticize two traditional solutions to the problem and then offer an alternative solution. In this chapter linguistic communities are given a role in providing and maintaining linguistic conventions of a sort needed to underwrite the positive theory of justification arrived at in the last chapter.

The third chapter provides an account of the role of a conventional language in the formation of belief, thus completing the theory of judgement. Sections A and B are concerned with the possible grounds we might have for attributing beliefs to others on the basis of their behavior. Out of those sections comes a need to shift attention to the puzzling resistance which belief attributions have to truth-preserving substitutions. Through an examination of the failure of substitution in belief reports in sections C - E, an account of the role played by language belief acquisition is arrived at. In the last section of this chapter, a difficulty which seems more acute than usual, because of the theory of language underway, is answered - the difficulty concerning how it is that our talk can be said to be about the world.

With the theory of judgement completed, Chapter Four addresses the problems of perception which bear on reaching an understanding of the relationships between perception and judgement. Sections A - C are concerned to say something about what sorts of things we can be said to per-

ceive. The answer provided, that we perceive only public particulars, makes the account of perception realist in nature. This eliminates certain skeptical doubts that might later be raised to trouble the theory of justification. Sections D - G provide an analysis of perception in one of its modes - seeing. The analysis given, when seen in light of the theory of judgement, immediately suggests a possible relationship between perception and judgement which can be stated in terms of concepts of sufficient generality as to be applicable to most of human activity.

Chapter Five then draws together the accounts of judgement and perception. In sections A - C general considerations about the relationship between perception and judgement are dealt with, using the accounts of each already provided. But it is found there that the accounts seem, thus far, to lack the resources for accommodating any justificatory relationship between the two. In sections D - H a relationship between justified beliefs or reports and perception is considered which ties the two terms of the relationship to a third - linguistic competence. An evaluation of three different views about justification in section H leads to an alternative theory in section I, but it is seen there that a problem in the theory of meaning blocks the use of the alternative theory. This problem the theory of meaning is dealt with in sections J - L, with the result that empiricist theories of justification are found to come into conflict with the view that languages

are conventional. Section M then proceeds to sketch the view of justification which preserves the thesis that languages are conventional. The theory of justification relies heavily on the account of conventionality provided in the first chapter.

What is hoped for is that, by the end, a general theory about the relationships between our epistemic position in the world and our many perceptual positions is completed in enough detail so that it can be seen how the account would go on to address other problems in the theory of knowledge. Three important problems are steered away from, though how they would be approached should be evident from what is completed - nothing is said about knowledge, but only belief; skeptical problems are not confronted head on, and so no attempt is made to give an account of certainty; and, finally, the view of language which is worked out here looks, by the end, as if it must be called upon to say something about the notion of truth.

ACKNOWLEDGEMENT

My list of philosophical creditors is long. The largest debt is owed to the guidance of Richard Bosley, whose influence has been exerted through ten years of rewarding conversations and can be seen on every page. Marshall Hopkins and Arlis Davis are also substantial shareholders in the ideas expressed here. The members of my supervisory and examining committees, Mohan Matthen, Leo Mos and Fred Dretske, and especially Allen Carlson and Bruce Hunter, provided the necessary goading and the right mixture of criticism and encouragement. Stephen DeHaven supplied these things (especially the goading) long after the others had gone to bed. Sue Campbell and Linda Hopkins provided the inspiration to dispell excessive doubts and the ballast for my flights of fancy.

TABLE OF CONTENTS

| CHAPTER | | PAGE |
|---------|---|------|
| I | THE TELEOLOGICAL GROUNDWORK | 1 |
| | Notes | 18 |
| II | JUDGEMENT AND CONVENTION | 19 |
| | Convention and Ostension | 19 |
| | Precedents and Conventions | 27 |
| | The Problem of Judgement | 42 |
| | Abstractionism | 46 |
| | Innatism | 64 |
| | The Possibility of Judgement | 76 |
| | Notes | 87 |
| III | PROPOSITIONAL STATES | 90 |
| | Behavior and the Underdetermination of Belief | 90 |
| | Behavioral Evidence for Propositional States | 101 |
| | Indirect Discourse: A Model for Propositional States | 106 |
| | Intensionality and Propositional States | 134 |
| | The Role of Language in the Acquisition of Beliefs and Knowledge | 148 |
| | Language and the World: The Case of Names | 160 |
| | Notes | 189 |

| CHAPTER | | PAGE |
|---------|---|------|
| IV | PERCEPTION | 194 |
| | Perception and Propositional States | 194 |
| | The Possible Objects of Perception: Publicity | 199 |
| | The Possible Objects of Perception: Particularity | 215 |
| | Causal Theories of Perception | 225 |
| | Seeing As An Achievement | 236 |
| | Visual Differentiation | 242 |
| | Perception and Discriminatory Action | 249 |
| | Notes | 279 |
| V | PERCEPTION AND JUSTIFICATION | 285 |
| | Perception and Judgement | 285 |
| | Perception and Judgements About the World | 296 |
| | The Inquirer and the Apple-Picker | 300 |
| | Perception and Verbal-Error | 307 |
| | Adequate Perceptual Positions | 313 |
| | Adequate Perceptual Positions for Judging | 315 |
| | Perception, Justification and Verbal Error | 328 |
| | Three Theories of Justification | 334 |
| | Firth: Verbal Errors and Ideolects | 344 |
| | Quine: Verbal Errors and the Interdeterminacy of Translation | 354 |
| | Eliminating Indeterminacy | 359 |
| | Empiricism and the Conventionality of Language | 374 |

CHAPTER

PAGE

| | | |
|---|--|-----|
| V | PERCEPTION AND JUSTIFICATION | |
| | The Justification of Perceptual Judgements | 381 |
| | Conclusion | 391 |
| | Notes | 396 |
| | BIBLIOGRAPHY | 400 |

CHAPTER ONE

THE TELEOLOGICAL GROUNDWORK

There are sometimes situations in which our reaching some goal, or managing some achievement, is only made possible by a conjunction of two very different sorts of factors: the exercise of a skill or an ability, and an opportunity. We may have all the requisite skills for doing something, and so for reaching a given objective, but lack the opportunity in which the exercise of those skills will take us to our objective. Likewise, there are many who have a special opportunity to do something, but who lack the skill to bring it off. In this essay I mean to show that reaching the objectives of human inquiry requires a conjunction of just such factors; that the relationships between perceiving, judging and acquiring justified beliefs are, respectively, just the relationships between having an opportunity, exercising a skill and reaching an objective.

The first part of my task will be to give an account of the nature of judgement, the resources needed for making judgements, and the propositional states of knowing and believing which we sometimes reach by using those resources. Secondly, I will provide an account of perception in which to perceive something is to have one kind of opportunity for performing one class of actions. The final task will be to show how the acquisition of justified perceptual beliefs requires the exercise of a special class of skills while

having a special kind of opportunity. It will be necessary to discuss and criticize parts of traditional theories in order to demonstrate the need for the account I offer. But the criticism will always be subordinated to the construction of an alternative theory.

Since the discussion will be carried out within the framework of a blatantly teleological theory of action, I will use this introductory chapter to defend the choice of a teleological theory over a causal theory, and to lay out two teleological notions which will play important roles in the following chapters. First, for the matter of defense.

My dissatisfaction with strictly causal accounts of action is this: They necessarily fail to answer a minimal demand which must be made of any philosophical theory of action - that it provide us with a way of making a distinction, among events involving agents, between those events which are actions and those which are not. A causal theory of action is one which takes that distinction to be based on a distinction between the causes of the events. If an event has the right sort of cause, then it is an action; if not, it is a mere movement of the body - a twitch and not a wink, or a stumble and not a slide to home plate. But I believe that it can be shown that no account of the distinction between events which are actions and those which are not can be given solely by reference to a distinction between the causes of those events. Many events (or many event-descriptions) which are clearly not actions (or

action-descriptions) are causally indistinguishable from many events which clearly are actions.

If on waving my hand I get the waiter's attention, I may also alert a comrade to start in motion the revolution, and in addition I may inadvertently disturb the air and the billiards player, frighten a fly, and startle the man seated behind me. A "fine-grained theory of act individuation" might take all these events to be properly called actions of mine, in which case such a theory should not be bothered in the same way by the argument I will give. But even fine-grained theories of act individuation recognize an important distinction among these events, namely that some of them are intentional actions of mine, and others are not. In the event of a fine-grained causal theory of actions, the object of criticism against which the argument to follow is directed becomes the possibility of making the intentional/unintentional distinction. Either way, a distinction basic to the theory of action will evade us if we look for the distinction only among the causes of various events.

Every action I perform has numerous unforeseen, unexpected, and even undesirable events which are causally coterminous with or causally consequent upon it. By virtue of being events which are causally consequent upon my action they are, eo ipso, events which are causally consequent upon whatever events caused my action. They are, therefore, causally consequent upon those events (whatever

they might be in a given version of the theory) which are said to make a consequent event into an action. All those events for which my action is causally efficacious ought to be actions of mine.

If we call "C" whatever it is which is thought to be causally sufficient for making some consequent event (E), an action, then we may put the point in this way: If C is causally sufficient for E₁, then E₁ is causally sufficient for E₂ ... E_n, and if C is causally sufficient for E₁, then C is causally sufficient for E₂ ... E_n. Therefore, E₂ ... E_n are actions. There are two obvious ways of avoiding this proliferation of individual actions into a potentially infinite series of actions. We might require a specific degree of "directness" of the causal efficacy of C, or we might add an additional, noncausal criterion for the class of actions.

Consider the first possibility. Call any event which is too indirectly caused by C an "indirect result" of C, but not an action. This move has two serious limitations. First, there are an indefinite number of events we can individuate which lie spatially and temporally between C and E₁, and on the causal path from C to E₁, which are not themselves actions; plus, there are an indefinite number of events on the same causal path but lying spatially and temporally beyond the occurrence of any movement of my body which might be said to be an action of mine or part of an action of mine. We introduce the intractable difficulty

that any specification of degree of casual directness will include some events which are not actions of mine and exclude some events which are actions of mine. There are actions (think of wiggling your finger, getting someone's attention by wiggling your finger, and starting the revolution by getting your comrade's attention by wiggling your finger) which are of radically different causal distances from C. Secondly, because of branching lines of causation there are always some events which are not actions but which are related to C by the same degree of causal directness as E.

Consider, then, the second way of avoiding an indefinite proliferation of actions within a causal model - by adding a further criterion which is met only by those events for which C is causally efficacious and which are actions. The standard way of doing this is to suppose that only those events are actions which are the intentional objects of certain beliefs and/or desires of the agent. C may or may not, on a given version of the theory, turn out to be identical with those beliefs and/or desires. So, if C causes E₁ ... E_n, but E₁ is the object of my beliefs and desires and E₂ ... E_n are not, then only E₁ an action of mine.

In the case where C is identical with the beliefs and desires of the agent, this theory remains open to two interpretations. Either E₁ is an action because it is caused by the beliefs and desires, or it is an action not

because it is caused by them, but because it is their intentional object (because it is what is believed will happen or be done and it is what is desired). If the causal efficacy of the beliefs and desires is relevant to E's counting as an action, then it is not this causal efficacy which results in the distinction between actions and nonactions, because the beliefs and desires are also causally efficacious in bringing about E ... E which are not, or not all, actions. So it must be the fact that E is what was believed done and was desired by the agent, which makes it relevantly different from E ... E; in which case the theory requires, in effect, that the agent's goals be the relevant criterion for distinguishing his actions from other events. The theory then becomes blatantly teleological. The final possibility left open is that C may not turn out to be identical with the beliefs and desires of the agent, though E alone counts as an action of his because it is both caused by C and is the object of his desires and beliefs about what he is doing. The same considerations weigh against this alternative as against the last. The need for the second conjunct in this sketchy account makes the theory into a teleological theory by placing the burden of drawing the distinction on citing the agents goals or objectives.

These considerations are, I think, sufficient to justify my decision to use a teleological framework when discussing actions in the following chapters. In the remainder of

this introduction I will introduce two teleological notions which will play crucial roles in our discussion of epistemological matters. First we will need a distinction between acts and achievements, and second we will need a notion of "subordination", which is one kind of relationship that distinguishable acts and achievements can have to each other.

The teleological theory used will be kept to a minimum. The notion of an action will be tied logically to the notion of an objective or goal, so that the difference between my arms moving towards the apple and my reaching for the apple is that in the latter case one of my objectives must be having the apple, grasping it, or eating it. The difference between my making the sounds normally made by someone who says "Please shut the door" and my saying "Please shut the door" is that in the latter case I perform an act of speech, say, making a request, and so one of my objectives is to get you to know what it is that I would like you to do - to get the door closed, or to deaden the sounds from the hall. In order to understand the structure of activities having any degree of complexity we must first become clear about whether any particular part of that activity is an act or an achievement. The distinction between acts and achievements had a relatively short philosophical history, more because of doubts about the usefulness of the distinction than because of doubts about its substance. The distinction is generally made by offering

examples, but since I will make a somewhat wider use of it than is usual, I will provide some criteria by which verbs which indicate the performance of an action are distinguished from verbs which indicate the management of an achievement. The primary distinction is not, however, between two kinds of verbs, but between two kinds of events. In many cases the same verb will function sometimes to indicate an action and sometimes to indicate an achievement.

The act/achievement distinction appears in various guises in the work of Gilbert Ryle,² who often contrasts activity, task and process verbs with achievement, success and got-it verbs, or verbs of starting and finishing.

There is an undeniable difference between running a race, reaching for an apple, looking for John, waving one's hand at the waiter, which are all actions or activities, and the achievements of winning a race, reaching the apple, finding John and catching the waiter's attention. The central metaphysical difference is that acts and activities are all events which have some temporal duration, but no achievement lasts for any length of time, however short. That some events have no temporal extension should come as no surprise since all events with a finite duration have beginnings and endings, and beginnings and endings are events which have no temporal extension. Events which last for any length of time can be said to have parts which last for shorter lengths of time. Consequently, actions can be

said to have temporal parts, and those parts have durations - running has as a part lifting one leg and as a later part putting it down in a different position. But since the winning of a race has no temporal extension, it has no temporal parts either.

Not every verb is either an act or an achievement verb - "broke" in "The window broke in the storm." does not serve to say what action was performed or what achievement was managed. Only when there is a verb of purposive behavior is it reasonable to ask whether it is an act or an achievement verb. The metaphysical point that achievements have no temporal extension will allow us to distinguish achievement verbs from among the purposive behavior verbs. I offer the following five criteria for achievement verbs; they are, perhaps, not independent of each other, but the only claim I make on their behalf is that any verb which has all of these features is an achievement verb, or that any event having these features (or which is indicated by a verb which has these features) is an achievement.

- 1) Achievements, unlike actions, cannot be left half finished - they are either reached or they are not. This follows from the metaphysical point that achievements are durationless events, and so, have no temporally extended parts. Taking a swing at someone is something I can start to do and then stop - in mid-swing - before I have finished; but I cannot hit him and then stop

before I have finished hitting him. Hitting, unlike swinging, is an achievement. This feature of achievements yields a grammatical test for the achievement verb:

2) If the main verb of a sentence is an achievement verb, then as soon as some simple present tense form of the sentence is true, a present perfect tense form of the same sentence is true. As Ryle points out, "I can say 'I have seen it' as soon as I can say 'I see it'".³

3) Because achievements lack a duration, and we cannot begin them, continue them, and at a later time finish them; we lack a felicitous use for achievement verbs in the present continuous tense. "I am reaching the top" only means that I am about to reach the top; "I am hitting him" only means that I have hit him, and I am repeating my past achievement. But "I am walking" reports an activity which is underway, not an event which is pending or being repeated.

4) Achievements have a way of disappearing if we take a count of the things we have done. This is never true of our actions. To run a race and to win it is not to do two things, running and something else in addition to running. But to mix a cake and to bake it is to do two things - first to mix it, and then to go on and do something else in addition. There is nothing that an athlete can do in addition to running a race in order

to win it. It is true that winning the race is not the same as running it - one must run it faster than the others, one's running must meet certain conditions specified in the rules of the competition. But no way of mixing the cake, and no conditions which the mixing meets, will bake the cake.

- 5) This notion of achievement includes, so far as I can tell, everything which J.L. Austin was prepared to call a "perlocutionary act". We may use one of his distinguishing marks of the perlocutionary as a distinguishing mark of achievements. Austin writes, "If we mention both a B act (illocution) and a C act (perlocution) we shall say 'by B-ing he C-ed' rather than 'in B-ing'.⁴ Extending this beyond acts of speech to a general view of action we can observe that we manage some achievement by performing some action, but the achievement cannot consist in performing that action. Winning does not consist in doing this, that, and the other thing in the way in which running does consist in moving one's legs in a certain way, but if we run by moving our legs in a certain way that is because running consists in moving our legs that way. With achievements it is a different matter - although we win by running, our winning does not consist in running.

None of this adds up to a theory about what an achievement is; these are only marks of achievements which will, jointly, suffice to distinguish them from actions. While a

verb which satisfies all five conditions will be an achievement verb, there is considerable room for debate as to whether or not a given verb does in fact satisfy them, so a case will have to be made for each verb. In general, nothing will be gained for my purposes by doing this, but in particular cases where the fact that a verb is an achievement verb bears important philosophical weight, I will present a case for taking it to be an achievement, and not an act verb.

Although achievements are one species of event belonging to the same general class as beginnings and endings, achievements are not the beginnings and endings of actions. They are not beginnings because an achievement presupposes the completion of an act; that is, we manage an achievement by, or as a result of, engaging in action. Even where an achievement stands temporally at the beginning of one activity, it is only an achievement qua being the culmination of another activity. For example, there is one sense of "getting the motor started" (not identical with pumping the gas pedal and turning the key, but the result of doing those things) which stands at the beginning of an activity of driving to Calgary. Though this sense marks the beginning of the drive to Calgary, it counts as an achievement only by virtue of marking the culmination of the activity of pumping the gas and turning the key. Nor are achievements identical with the endings of actions. When I pump the gas and turn the key I may yet fail to get the motor

started and, after a few tries, give up. The action is over but the achievement of getting the motor started is not managed, nor is there some other achievement which I have managed.

Although achievements do not consist in actions, they are importantly related in that achievements are the results got, in one sense, by performing actions. I say "in one sense" because, for all the interesting similarities, the act/achievement dichotomy is not the same as the process/product dichotomy. Building is a process or activity, the product of which is a building. But a building is not an event, let alone a durationless one. Completing the building (product) can be an achievement, but the building cannot. We may, then, say that an achievement is a result got by performing an action, though it is not the product of the action. The result of quickly turning my head to the left is that I catch a glimpse of the ball, but neither the catching of the glimpse, nor the glimpse, will be a product I have produced by turning my head.

Finally, it should be emphasized that the notion of an achievement, as used here, does not coincide very well with the ordinary use of the word "achievement". Taking fourth place is as much an achievement as winning the race and so, for that matter, is losing. There is no presupposition made that the achievement got by performing an act was intended or desired. By turning my head I may catch a glimpse of something to my complete surprise; it may be the

last thing on earth that I expected, intended or wanted to see. These sorts of achievements are, in ordinary parlance, more accurately said to be only the results of our actions since they have no obvious relation to our objectives. But, as I have said, the more technical notion of achievement that we are working with just is one kind of result.

Together with the act/achievement distinction we will need to keep before us the fact that doing one thing can be subordinated to doing another. Performing one act can be subordinated to winning it, and managing one achievement can be subordinated to managing another, as winning this race can be subordinate to qualifying for the Olympics. In general, if A is done in order to do B (or simply "in order to B"), in which case B is done by A-ing, then A is subordinated to B. But A is not subordinated to B if doing B consists in doing A. In the latter case there is a relation of identity between doing A and doing B, but not a relation of subordination. If winning all of the first, second and third races is identical with qualifying for the Olympics, then winning the first race is subordinated to qualifying, winning the second race is subordinated to qualifying, and winning the third race is subordinated to qualifying, but winning all of them is not.

We can understand the notion of subordination by viewing relatively complex activities as structures of less complex acts and achievements. I may turn the steering

wheel slightly in order to make a turn, but this turn and many others may all be made in order to arrive at the grocery store. I arrive at the grocery store in order to purchase a load of bread, and I am purchasing the bread in order to make sandwiches in order to gain weight for my next bout, and I am trying to gain weight for my next bout in order to win the title. Driving to the store is then subordinated to buying bread, but it is also subordinated to winning the title. Let us say that driving to the store is more deeply subordinated to winning the title than is, say, winning the next bout. In general, if doing A is subordinated to B, and B is subordinated to C, then doing A is subordinated to C, though more deeply subordinated than B.

Notice that both acts and achievements can be subordinated to a given achievement. I can wave my arm (act) in order to get the waiter's attention (achievement) and I can get the waiter's attention in order to get my bill (achievement). I may ponder and search (act) for the right word, remember it (achievement), then utter the word (act) in order to tell you something (achievement). If we are skilled enough in any particular activity, especially if we have become skilled by practicing and following the example of others, we need not be cognizant of the subordinate acts we perform and the subordinate achievements we manage in engaging in that activity. I might, through practice, have become skillful enough to hit the bull's-eye nine times out

of ten and only discover later, if at all, that I was able to manage that feat by holding the elbow of my shooting arm parallel to the ground and pointing my forward foot directly at the target. Nonetheless, each time I get into a position to shoot, I move into that posture in order to hit the target - I raise my elbow and turn my foot in order to get into that posture. Few, if any, could begin to say exactly how they walk, talk or ride a bicycle, but when they do walk they swing their arms and legs in a certain rhythm, they shift their centre of gravity in appropriate ways, and numerous other things, all in order to get from one place to the next.

The structure of subordination needed to manage some achievement must be uncovered in order to give a complete account of how something gets done, but reflection on and knowledge of that structure is not a prerequisite for performing the action and managing the achievement; it could even get in the way. Giving such an account is one of our tasks, and understanding how a certain achievement gets managed is one of our objectives, not the agent's.

With this admittedly sketchy account of the need for, and structure of a teleological framework for action, we can proceed to the particular task at hand. The teleology will not be a glaring presence in most of what follows, and shows itself only in relatively innocuous ways. But its use is crucial in understanding what the relationship

between justified perceptual belief and judgement is, and at that point its presence will be very conspicuous.

The result of the use of a teleological framework throughout the following chapters will be to stake out a position from which to take a stand against the Cartesian spirit in most of epistemology. When Descartes tells us, in an autobiographical anecdote, that

The setting in of winter detained me in a quarter where, since I found no society to divert me, while fortunately I had also no cares or passions to trouble me, I remained the whole day shut up alone in a stove heated room, where I had complete leisure to occupy myself with my own thoughts ...5

He illustrates well the form into which his influence cast the problems with which it is thought to be the business of epistemologists to concern themselves. The Cartesian spectre of a mind shut up alone to occupy itself with its own thoughts can be seen haunting the remark of a contemporary epistemologist, Keith Lehrer, who writes that

... there is no exit from the circle of one's own beliefs from which one can sally forth to find some exquisite tool to measure the merits of what lies within the circle of subjectivity.⁶

The difficulty for epistemology since Descartes has been to find a way out of the stove heated room, or to accustom itself, as Lehrer's image suggests, to there being no way out. The Discourse On Method makes clear that Descartes views the society, from whose diversions he withdraws to the stove heated room, as the source of the preju-

dices from which his method is meant to free us, so that we may undertake a project of pure, unsullied inquiry. The conclusion I argue for in the final chapter is that society provides the only appropriate standard against which we can measure the success or failure of our inquiries - the "exquisite tool to measure the merits of what lies within the circle of subjectivity".

By understanding inquiry as a complex social activity which involves the use of social tools and skills, and by understanding the results of inquiry as epistemic achievements, the appropriate objects for investigation by epistemologists are seen to be the complex arrays of subordinate acts and achievements which make up the project of inquiring. The relevant acts and achievements are linguistic, and it is to the examination of these that we now turn.

Notes for Chapter One

1. As in A.I. Goldman, A Theory of Human Action, Prentice-Hall, Englewood Cliffs, New Jersey, 1970.
2. Especially in Dilemmas, Chapter VII, Cambridge University Press, London, 1954, and The Concept of Mind, Penguin Books, Ltd., Harmondsworth, Middlesex, England, 1949.
3. Dilemmas, Cambridge University Press, London, 1954, p. 102.
4. How To Do Things With Words, Oxford University Press, New York, 1962, p. 107.
5. Discourse on Method; trans. by E.S. Haldand and G.R.T. Ross, Cambridge University Press, London, 1972, p. 87.
6. Knowledge, Oxford University Press, Oxford, 1974, pp. 17-18.

CHAPTER TWO

JUDGEMENT AND CONVENTION

A. Conventionality and Ostension

In this chapter I will discuss what earlier generations would commonly have referred to as a theory of judgement. Traditionally, the term 'judgement' has been used equivocally, referring sometimes to an act of judging and sometimes to a product of such an act - an artifact devised by the work of a mind to represent to itself the world. As I will use the term, judgements are sometimes mental acts and sometimes mental achievements; this much ambiguity is unavoidable, but I eschew entirely the use of 'judgement' to refer to a product of mental or linguistic activity. Hopefully, this will reduce the temptation to wonder, at a later stage, about the relationship between judgements and the states of affairs about which we make judgements. The metaphor in common use - 'making a judgement' - will be retained; but it ought to incline us no more towards a view of 'judgements made' as mental or linguistic products or artifacts than talk of making an effort inclines us towards viewing the effort made as an artifact or product left over from the activity of making it.

The theory of judgement has two parts - an account of the skills and resources we use in judging and an account of the use of those skills and resources in making judge-

ments. In another mode of expression, the two parts are a theory of concept formation or acquisition, and an account of the use of those concepts in acquiring belief.

I will use an assumption about natural languages which will be left undefended (if it should even be thought to need any defense). This assumption sets the boundary for the theory of judgement and the epistemology which follows; conflict with it will be taken as sufficient reason for rejecting other positions. My assumption is that natural languages are conventional. I take this to mean that, first, the use of an expression in a language is, counterfactually, arbitrary, but noncounterfactually, nonarbitrary. Thus, if things had gone otherwise, the expression "red" could have been used in the way in which we now use the expression "green", and vice versa, but as things happen to have gone, the expression "red" cannot be used in the way in which we now use the expression "green". Second, the use of an expression which would otherwise be arbitrary is rendered nonarbitrary by demands for conformity between the usage of one speaker and the usage of other speakers in a community. The actual usage of a community provides the conventions or precedents with which an individual's usage must accord. The sense to be given to the crucial notion of arbitrariness used here will be clarified shortly. My assumption that the use of an expression is conventional is meant to contrast with the view that there is a natural, nonarbitrary use for each expression which is

independent of the actual linguistic practices of a linguistic community.

This assumption sets certain limits on what will count as using a language and on what will count as learning a language. The second part of the thesis that language is conventional is, perhaps, an empirical thesis. Once the use of an expression is fixed or established, there are restrictions on what will count as the correct use of the expression by a speaker of the language. The relevant restrictions on the use of a word in a language consist in demands made for conformity to actual linguistic practice; demands which are enforced in the language learning situation. The need to maintain conformity of word usage within a given linguistic community is satisfied by teaching continuity of linguistic practices between each speaker and his linguistic peers. There will no doubt be many breaks in continuity within any community of speakers, but the language learning process is never completed, and we have corrective procedures for the mature speaker as well as for the novice. Where a linguistic community is lax in enforcing its restrictions and maintaining linguistic continuity, the consequence is rarely anything worse than the offspring of language we know as idiolects; but even their survival depends upon maintaining relatively strict conformity of usage within the linguistic subcommunity. That we satisfy the need to preserve conformity of usage within our linguistic communities by applying corrective procedures in

language learning situations, and so converge the restrictions on the use of a word by using the word, may be an accidental fact about our languages. Perhaps the same result could be obtained by surgery, and so language acquisition need not be language learning. But it is not an accidental fact that what our procedures are designed to do is secure conformity of usage between one speaker and others.

It will be useful on occasion to resort to a baptismal model when talking about how an ongoing linguistic practice, or convention, gets started. It does not matter whether the precedents for the use of a word are a result of tacit or explicit agreement, or of one person being in a position of power or authority adequate for imposing his preferences on others. David Lewis¹ has argued that conventions can be established by example in a setting which involves neither agreement nor authority. But these historical matters can be set aside. The fact that no quasi-baptismal ceremony need ever be performed does not rule out the possibility of putting an expression to use in that way. What should be kept in mind is that the possibility of introducing an expression into our language in this way shows that whatever conditions are necessary for establishing a convention are met by the baptismal story. No logical injustice will be done by simplifying the following discussion with the baptismal model for the establishment of a linguistic convention. At any rate, the only expressions whose use I

will be concerned with are those which can be introduced in that way, and so those expressions which, if only their histories were known, might have been introduced by some baptismal act.

The first important point which follows from the conventional nature of language is that there is a large and important class of expressions, the coining and teaching of which is only made possible by ostensive acts of reference. One condition which we must insist be satisfied if the expressions in our language are to be given their use by convention is that it be possible to make clear with respect to what each convention for speech is to be established. An object, with respect to which a convention is established for saying what it is, must be individuated for the initiates into the conventional practice, before it becomes possible for them to follow the practice of saying what it is by uttering a certain expression on the right occasion. This is the requirement that such an object be individuated prior to the establishment of the convention. We cannot teach a child the use of the expression "ball", making clear to him what object the expression is to be used for speaking of, by saying to the child "every ball". Nor can the child be taught the use of the expression by being given a description of the objects which it is to be used for speaking of. The general requirement for the establishment of conventions for saying what things are must be that an important class of objects be individuated

for the initiates into a set of conventions by nonlinguistic means. Failing to satisfy this requirement would necessarily result in the failure of the initiate to learn the use of the expression and would necessarily lead an account of the possibility of language into circularity. Rather than tell what the expression "ball" can be used to refer to, we must show him. Only when the linguistic resources of a child increase as he progresses in his use of a language does it become possible for him to begin to learn the use of other expressions by being told how to use them. The same must be said of the original development of the linguistic framework of a community of speakers - only as their linguistic resources develop does it become possible for them to add new expressions to their vocabulary by saying how they are to be used.

Fortunately, the prelinguistic skills for individuating objects with respect to which a convention is to be established are not hard to find. We can hold up, look at, point towards, and otherwise direct attention to a very large class of individuals, without relying on prior linguistic skills. So it is this general requirement that it be possible to show nonverbally how, or with respect to what object, an expression is to be used, which leads me to say that the conventionality of our language presupposes the possibility of demonstrating the use of at least some expressions by prelinguistic means. Since the account of judgement to follow will be limited to an account of the

use of expressions for saying of something what it is, or what perceptual features it has, I will hereafter refer to the relevant prelinguistic resources presupposed in both the institution of such conventions and the teaching of them, as ostensive or directional acts.²

It is not necessary to assume that the class of expressions I am concerned with are the first expressions introduced in any language - indeed, there may be good reasons for thinking it likely that expressions involved in the performance of certain basic social acts like commanding, warning, and appeasing would have come first in the order of genesis. An imperative verb like "fetch" or "bring" need not presuppose the prior ability to use expressions like "wood" or "water" for saying what is to be fetched, since that can be made clear by coordinating the utterance of "fetch" with an act of pointing. So it is not impossible that the introduction of an expression for saying what some object of ostension is to be called involves coordination between a directional act of pointing or looking and the utterance of a demonstrative expression, in order to succeed in getting someone's attention directed to the relevant object.

The expression "this" in the sentence "This is an X" differs significantly from expressions put to the service of saying of something what it is, but it does not for that reason fail to be an expression which gets its use as a result of convention. We get our listener's attention

directed toward something by coordinating the utterance of "this" with an ostensive act, but we could just as well have achieved that end with the practice of, say, snapping our fingers and then pointing towards an object. There is no presupposition that a person already be able to say, or even to think, what the object to which he is directing our attention is, in order for him to be able to point at it and say "this" or for him to be able to say "fetch this". This is part of what is so far-fetched in Russell's characterizing the demonstratives as the perfect proper names.³ The use of demonstratives in no way helps to say what an object of our acquaintance is - they only help to show which object we are then concerned with.

A child, mimicking our behavior of uttering "this" while pointing, can point to any object other than the one we pointed at and still be using the expression just as we do - that is, he would be following that convention correctly. What is more, he may also get the result of directing someone's attention to that object. Other ways of going wrong become possible when an expression is used for saying what something is. It follows that something more must be done to establish such a convention or to teach one. It is necessary, but not sufficient, for us to direct an initiate's attention to an object with respect to which a convention is to be either established or taught. He must then be shown how to go on from there - what to say about it. Hence, "this" is followed by "X" or "is an X".

Although it is not necessary for the establishment or teaching of a convention for saying what something is that both a directional act and a linguistic act be performed in making clear which thing is to be named, I will generally use examples in which someone points and says, "This is (an) X". It is only necessary to perform both directional and linguistic acts when we advance beyond showing which thing is to be named and proceed to show what to say in order to say what the thing is - this we can show only by saying "X", and thereby setting or teaching a linguistic precedent for the use of the expression "X". In showing someone which thing we mean to give a name for we need only perform nonlinguistic acts (pointing, holding up, etc.) and we manage, when we succeed, nonlinguistic achievements (getting someone to look at something, or getting someone's attention directed and fixed). These are resources we have prior to language, but we can subordinate them to instituting, teaching, and using a language.

B. Precedents and Conventions

I will be engaged in laying out some of the support needed in language acquisition as we proceed. This account is meant to apply both to the acquisition of language by children born into a community that already has a language and to the acquisition of language by a society of creatures who at first have no language at all. In order to

keep the second sort of case before us, I will frequently use the notion of a precedent for speech, rather than the notion of a convention for speech. The notion of a precedent is the pivotal notion for my account of judgements. I use the expression "precedent" in giving this account because it brings with it a distinction between setting a precedent and following one. If a speaker sets a precedent for saying what something is by saying "This is an X", it thereby becomes possible for others to follow his precedent and say of that thing "This is an X". We have, then, two steps for introducing into our language an expression for saying what something is, or for teaching an initiate the use of such an expression. First, the object must be individuated for the initiate, and at a certain stage this must be done by nonlinguistic means. Second, a precedent must be set for using a certain expression in saying what the object is.

A critical feature of precedents for speech is that they are immune to one form of criticism. That is to say, in setting a precedent for speech an agent performs an act of uttering a word and his action cannot be criticized in at least one way in which the act of utterance of an agent who is following the precedent can be criticized. This is the substance of the assumption that our languages are conventional and not natural. In setting a precedent, an agent utters a word which can then be used for, say, saying who this child is, saying what this is, or saying what color

this thing is. But what word the speaker chooses to utter is in one sense arbitrary, since until he sets an example for us, our community has no other word for saying those things. Since there is, at the time of the setting of the precedent, no word for saying what this is, and since there is no standard imposed by the thing itself which the utterance must meet, a speaker who sets a precedent cannot utter the wrong word. This trivial point will prove to have important epistemological consequences.

There are many ways in which a speaker who sets a precedent can be criticized. He may be criticized on aesthetic grounds because the word he invented for us to use has a harsh sound, or because he chose a pretty sounding word for an ugly animal. He may also be criticized on pragmatic grounds because the word he chose for his son is too long or too hard to pronounce in a short space of time, or too hard to distinguish from his sister's name. He may even be criticized because we already have a precedent for speaking about just those things, and a synonym would only confuse us, or because we already use the same word for another job, and a homonym would only confuse us. Then again, if we are not easily confused about these sorts of things we might allow the synonym or the homonym to stand, just as we might allow the very long word to stand, or the ugly word.

We might criticize a man for naming all five of his sons John, but we could not criticize him on the ground

that that is not their real name, unless they have already been baptized differently. The one kind of criticism the precedent setter is not open to, when he says "This is X", is the criticism that what he said is false; that it is not really an X.

A warning must be given that the notion of a convention for speech is not the same as the notion of a precedent for speech. I will sometimes talk about precedents rather than conventions, because the notion of a precedent keeps clearly before us the baptismal model, and so the sense of arbitrariness and the inapplicability of notions of correctness and incorrectness. But the simplifications afforded by the model bring distortions as well. Conventions must replace precedents in discussion of language use among speakers who are far removed from the baptismal fount.

It is a matter of convention in a linguistic community to use a particular expression for saying who this man is, only if the expression has a history of use in that community for doing that job. Precedents and conventions are not the same because, though every case of following a convention is a case of following a precedent, not every case of following a precedent is a case of following a convention. A convention exists only when a precedent has become entrenched in practice by regular use. Setting and following a precedent could be a one shot affair; only when the use of a particular precedent for saying what something is becomes a relatively established tradition in a commu-

nity of speakers do we have a convention for saying what that thing is. For the most part the precedents of a community are lost in its history, but the conventions are displayed in present speech practices. Our children learn to follow our conventions, by observing and imitating our behavior, when they learn to speak our language. Whenever we follow a convention for the use of a word, any word, our act of utterance stands in a relationship with an earlier act of utterance, or group of utterances, which counts as the precedent for the use of that word. Saul Kripke, writing about proper names, describes just this relationship with the image of a chain. After the initial baptism of Feynman, that is, after the setting of the precedent of using "Feynman" in order to say who this particular person is,

Through various sorts of talk the name is spread from link to link as if by a chain A certain passage of communication reaching ultimately to the man himself does reach the speaker . . . a chain of communication going back to Feynman himself has been established, by virtue of his membership in a community which passed the name on from link to link.⁴

Kripke's concerns are very different from mine, as are his conclusions. One difference needs to be noted for future reference. Kripke's use of the image of a chain of communicative transmission, in which the name 'Feynman' reaches back to Feynman himself, differs from the chain I have described, which relates one speaker's use of an expression with the use made of it by other speakers. The

relationship of 'following' a linguistic precedent or convention is a relationship which can obtain only between acts of speech. So, the relationship of which Kripke speaks, which obtains between uses of the name 'Feynman' and Feynman himself, cannot be of the same sort.

In general, the relationship, if any, between a speaker's use of a word and the item about which he speaks, must be a relation of a very different sort than the relationship between one speaker's use of an expression and the use made of that expression by other speakers. Kripke's discussion of the chain of transmission obscures that difference. The chain of which I speak extends only from one speaker to others, but cannot extend outside the community of speakers to link their speech to the nonlinguistic world.

It is useful to note the radically different sorts of relationships which Kripke's use of the image of the chain obscures, and to note the end which the obfuscation serves. By imitating other speakers, an initiate into the language puts his usage into line with the usage of others - his 'following' their convention consists in him imitating their practices. If his linguistic peers and elders are in turn mimics of the practices of an earlier, or more extended generation of speakers, then there is a univocal relation (of copying, mimicking or imitating) which the initiate's utterances bear to the utterances of the more remote speakers. But that relation, which I am calling 'follow-

ing', does not (without the aid of a mimetic theory of meaning or reference) stand between the utterances of any speaker - even the 'name-giver' - and the items about which he speaks. So my use of a name, say "Aristotle", does not, by virtue alone of its position in a successfully maintained chain of transmission, stand in relationship to Aristotle. The 'name-giver' is at the beginning of the chain, but Aristotle is not. A speaker remote from the baptismal fount is related to the name-giver by acts of following; the name-giver, however, could only be 'related' to Aristotle by acts of attention, perception, or description. By obscuring the difference between these, Kripke's metaphor gives the illusion that all members of a linguistic community are, by virtue of their position in such causal chains, in a relationship with the same items of the nonlinguistic world.

I will continue to use the metaphor of a chain of transmission, along which the use of expressions is passed from speaker to speaker, and which serves to bind individuals into a linguistic community. But the chain of communication which I will be speaking about puts speakers in relationships with each other and not with the items about which they speak.

There are many kinds of variations and even breaks possible in the chain. The pronunciation of a word may vary slightly but regularly over a long time. Our pronunciation of "Socrates" need bear little resemblance to his

own pronunciation of his name. At some point a precedent might be overthrown and a new word instituted into the language. That is why, in the long run, original precedents do not matter. Initially, a child does not learn to speak by discovering the history of a language; he merely follows the example of his contemporaries. If what he says bears some relationship to speakers in the distant past, that is because his community has maintained continuity with them. His own training only requires that he learn to maintain continuity with his own community.

The importance, for our purposes, of the distinction between a precedent and a convention for speech is that an initiate into a conventional practice is trained and educated by his contemporaries and does not need continuity in his practices with the history of his language. Perhaps the community does maintain continuity between their practices and the original precedents by means of this training and education. But in the transmission of language from one generation to the next no deliberate effort is, or need be, made to maintain continuity between present practices and the far linguistic past of the community. That there is some such continuity, in spite of the shifts and modifications, and even reversals of original precedents is, epistemologically, an unimportant consequence of the historical transmission of community practices.

Though it is a feature of incidental interest that a community maintains continuity with its linguistic ances-

tors, it is of the utmost importance that the individuals in the community maintain continuity with the rest of the community. It should be stressed that the continuity needed is not a continuity of opinions but a continuity of practices. If there is not already agreement in practice, then there is no support for meaningful disagreement or agreement in opinions. Thus a community of carpenters could have no meaningful dispute about whether a given board is more or less than one foot long if each carpenter used a standard foot rule of a different length. An agreement as to which standard to use for "one foot" is not an agreement of opinion as to a matter of fact, but only an agreement in using "one foot", and a standard rod, in a certain way. An apprentice entering the community is not told the opinion of the others, rather he is trained in a set of practices which include the use of the expression "one foot" and the use of a certain rod as a measuring device.

To follow a convention for speech is to maintain continuity of one's own linguistic resources and practices with those of one's immediate community. This sort of continuity is unavoidable, since the community passes those resources on to the individual and trains him in their use. It is because, in following the linguistic practices of our community, we are thereby following the conventions for speech that the notion of a precedent can yield, together with its simplifications, to the notion of a convention.

All that we lose by replacing one notion with the other is an historical model. Conventions take with them the most important feature of precedents, namely that they are immune to certain sorts of criticism. They are, in the same sense, arbitrary, and so are open to criticism on aesthetic or pragmatic grounds, but they are not open to the criticism of being false. In so far as our community has a conventional practice of uttering a certain expression for doing a certain job, it is not open to the charge that it might be getting the facts wrong. In preserving such a practice, the community only provides the resources for doing certain jobs, for performing certain acts and for managing certain achievements, but it cannot share the blame with a member of the community who uses those resources and bungles the job.

We have seen that the act of setting precedents for speech is not characterized by one notion of correctness. There is no possibility of uttering the wrong word, or the right word, when we lay the foundations of a language by setting precedents. For the same reasons, conventions do not admit of truth or falsity. If our community uses the word "red" for saying what color ripe tomatoes happen, as a matter of fact, to be, it is not subject to criticisms of correctness and incorrectness for having that convention.

Once, however, a precedent is set or a convention established, it becomes possible to follow the precedent or convention and, ipso facto, it becomes possible to do

something wrong which we could not have done wrong in setting a precedent. One can follow a precedent correctly, or one can fail to follow it correctly. Like any other convention, a convention for speech is followed only if one does that which it is a convention to do. If there is a convention of kneeling in front of a certain statue and I do not kneel, but curtsy, when I am in front of the statue, then I fail to follow the convention. To curtsy is to do the wrong thing. Following a precedent or a convention is a matter of doing the same thing that was done by whoever set the precedent or whoever maintains the convention in their practices. A fortiori to follow a precedent for speech is to do something that was done by the person who set the precedent, and to follow a convention for speech is to do something that a particular community, in which there is such a convention, does. In particular, if someone sets a precedent for saying what this is by saying "X", I follow the precedent by saying "X". To follow a precedent for speech is to utter the same word as the person who set the precedent.

To utter the same word is not yet enough, however, to utter either the right or the wrong word. If there is a precedent for saying who this man is by saying "John" and a precedent for saying who that man is by saying "Frank", and if I utter the name "John" in order to say who that man is, then I have uttered the wrong name, even though I have followed a precedent - namely, the precedent for saying who

this man is. If my objective is to say who this man is, then I must say "John", because that is the precedent for saying who he is. It is a truism with respect to anything that we might be doing that whether or not we are doing it correctly depends upon exactly what it is that we are trying to do, or what it is that we are trying to get done.

It seems then, that to utter the right or the wrong word, two things are presupposed:

- 1) That there is a precedent set or a convention established for the use of the word that is uttered; and
- 2) That there is some objective aimed at in uttering the word.

The failure of either of these conditions is sufficient for the categories of correctness and incorrectness, or rightness and wrongness, to fail to be applicable to the act of utterance or the word uttered and, I suggest, the satisfaction of both of these conditions is sufficient for the applicability of these categories to an utterance.

Together these conditions provide us with a sense in which the conventions for speech provide us with standards against which to measure our success or failure when we engage in certain acts of speech - namely, acts of uttering a word, as when we say who this man is, or what the color of this is. If we say whatever it is conventional to say in order to do such and such, then we have correctly followed the convention for doing that, and we can succeed in

doing that. But the notion of correctly following conventions for speech is not the same as the notion of speaking the truth. Correctly following conventions is necessary for performing acts of speech to which the notion of truth is not applicable, such as giving a particular command, or asking a particular question (one may not correctly follow a convention and still succeed in putting a question, but not the right question). But correctly following conventions for speech is at least a necessary condition for the truth of any particular claim or assertion. If we say something in order to say what color some red object is, and we begin by saying "The color of this is", and then go on to utter a word for the use of which there is a convention, then unless we utter the same word for which there is a convention in use for saying what that color is, we necessarily fail to say something which is true. What we said could only have been true ~~if we had uttered the word~~ "red"; but to have uttered that word is to have correctly followed the convention for saying what that color is. Failure to correctly follow any precedent or convention in making an assertion or a claim will guarantee either the falsity, or the lack of the applicability of any truth-value to what is said.

If there is a convention for saying what this is by saying "X", then to say that this is an X is to follow that convention correctly. But if this is an X, then it is true to say that it is an X. But what is it to say "This is an

X"? Minimally, it is to follow the convention for saying what it is. Contrariwise, if this is not an X, then it is false to say that it is an X, and to say that it is an X is necessarily to fail to follow the convention for saying what this is. In J.L. Austin's defense of a particular theory of truth, he writes

The truth of statements remains still a matter, as it was with the most rudimentary languages, of the words used being the ones conventionally appointed for situations of the type to which that referred to belongs.⁵

Austin is here writing about the relationship between the words used in their "conventionally appointed" roles and certain "situations of the type to which that referred to belongs". But this relationship, which Austin takes to be the important 'truth-making' relation, is not the one we are talking about here. If the words are used in their conventionally appointed roles, then each utterance of a word has a relationship to another utterance or another group of utterances - namely, the relationship of following a precedent or of following a convention for the use of that word. It is this relationship which I am claiming is necessary for truth, and the lack of which is sufficient for failing to say something which is true.

A community of speakers maintains continuity in its practices by training and educating initiates into the language in the following of conventions for speech. In doing so the community provides its members with the re-

sources for speaking and trains them in the use of those resources; in addition it provides them with standards which a speaker's utterances must meet (that is, conventions which they must follow in speaking) if they are to be correct, or say what is true.

Finally, it should be noted that if to speak a language is, at one deep level of subordination, to utter words, and if to utter words is to follow the conventions of a community, then speaking is one of those activities which requires that we make an effort to do correctly. It is true that most adult speakers need make no more effort to talk than they need make to ride a bicycle or dance. Perhaps talking might be one of those things which might be better said to require an effort to learn to do, but no longer demands effort once we have learned it. But what cannot be ignored is that maintaining continuity with the practices of a community requires that we have and use a rich variety of skills and faculties - everything from controlling the timing of our utterances and the pronunciation of our words, to having and exercising the faculty of memory in order to make the right sound, and so utter the correct word, on a given occasion. Only when our training and skills fail us, as when we are stuck for the right word, or can't recall the name of an acquaintance, do the operations we perform regularly become conspicuous. Talking is an activity which involves operations and mechanisms at a deep level which deserve to be rather sharply contrasted with

such events as the conditioned salivation of Pavlov's dogs. They neither need to, nor could, make an effort to salivate when the bell rang; there could have been no occasion on which they had to make an extra effort in order to remember whether or not this is the right time to salivate ("Was that the bell, or just a passing train?"), in the way in which we sometimes must make that extra effort to maintain continuity between our behavior and the practices of our community ("Is his name 'John', or is it 'Joseph'?"). A dog's failing to salivate at a certain time might indicate that something is wrong with the dog, or that something is wrong with the theory of conditioned responses, but would never indicate that he has made a mistake.

C. The Problem of Judgement

At the beginning of the previous section I pointed out that there are two very different sorts of things which must be done, or conditions which must be met, for the introduction into our language of an expression for saying what something is, or what color it is, or for teaching an initiate the use of such expressions. It must be possible to nonlinguistically individuate some object for an initiate, and an expression must be introduced. Both are necessary conditions for someone's being able to follow a precedent, and hence for later saying of something what it is by saying "This is (an) X". The failure to meet either

of these conditions necessarily results in the failure of the initiate to learn the use of the expression. This much is, I believe, relatively uncontentious - it is all part of what one is committed to in accepting the assumption that the use of an expression in a language is a matter of convention. Important divisions among philosophers begin to show themselves when an attempt is made to say what is presupposed by one's ability to follow a precedent for speech once it is allowed that they can be set. Rationalists can be divided from the empiricists by what they say that we bring with us to the language learning situation in order to account for our ability to follow a precedent for speech. The former hold that we must bring prior linguistic skills or dispositions to the learning situation, whereas the latter hold that the learning situation itself is sufficient to account for our using language in the way that we do.

The problems will differ when the precedent introduces a proper name, or name for an individual, and when it introduces a common name, or name common to members of a kind. Further complications arise in accounting for the possibility of learning and following a precedent for saying of an individual that it has a certain property or feature which; we assume, other individuals might share with it. If a proper name (hence, "P") is given, for saying of an individual who (or which) he (or it) is at some baptismal ceremony, an initiate must have and exercise an ability to

perhaps recognize that individual again and an ability to recall the expression "P" for saying who (which) he (it) is. Failure to have and exercise such skills will result in an inability to say what something is. Without begging the question as to which skills (recognition, recollection, memory) are in fact necessary for following any particular precedent for speaking, I will introduce the notion of "judging" in order to talk about what the initiate is doing in exercising those skills in order to follow a precedent for speech. We judge that this is the person for whom the precedent of saying "P" was set for saying who he is. In telling you who it is that is standing in the corner of the room I must follow the precedent of saying "P", and in order to do that I must judge that the person standing in the corner of the room is that person of whom I learned that what to say in order to say who he is, is "P".

Complications arise for property terms and kind terms (hence, "F" and "X", respectively) just because we are not taught those conventions by being shown all X's and all things that are F, nor can it be a requirement for the establishment of the convention that all instances of a given kind or color first be individuated ostensively. The possibility of introducing such expressions into our language with ostensive acts rests on the possibility of selecting paradigms, or examples, and saying of them "This is (an) X" or "This X is F". The presupposition is then made that the initiate will be able to "generalize" from

these paradigms in the sense that he will be able to proceed to things other than the paradigms and say of them "This is an X" or "This X is F". In other words, the precedents set and followed in these cases are not merely precedents for saying of the paradigm what it is, or what color it is, but for saying of all things of the same kind or color as the paradigm what they are, or what color they are. Being able to judge that this thing is one of the same kind of things as the paradigm is a necessary condition for mastering the use of the expression "This is an X", and the ability to judge that the color of this thing is the same as the color of the paradigm is a necessary condition for learning the use of some color expression.

The major problem which confronts any theory of judgment can now be stated: How is it possible for an initiate into a language to acquire, from ostensive individuation of the paradigms, the ability to proceed correctly to individuals of the same kind, or having the same property, as that paradigm? The solution to this problem has traditionally taken one of two forms. Rationalists account for this ability by citing prior or "innate" ideas, concepts, or linguistic knowledge. Empiricists give an account of our ability to use these expressions by viewing the language learning situation as a testing ground for a series of increasingly better hypotheses which the language learner makes as to the exact property, or properties, defining a class of objects, with respect to which a given precedent

for speech is being set or taught. The account to be given partial development in the remainder of this chapter is best seen, in ways to be clarified in the final section, as an alternative to both the innatism of the rationalists and the abstractionism of the empirists. Since section E is largely devoted to expounding and defending a principle incompatible with the view that there are any prior linguistic resources, I will devote the next section of this chapter to stating my reasons for seeking an alternative to abstractionism. Although the final defense of my charges against abstractionism must wait until the fourth chapter, this is the place to introduce those charges.

D. Abstractionism

Abstractionism, a theory of concept formation, and hence a theory of language acquisition, is the view that the skills acquired and used in coming to follow precedents for speech are a product of the formation and testing of hypotheses in the language learning situation itself. The first difficulty that abstractionism confronts is that the very possibility of forming and testing hypotheses might be thought to presuppose the use of a language of some sort. That is, abstractionism can be charged with presupposing the prior operation of an innate language, and hence the truth of a doctrine to which it is meant to be an alternative. It will be the argument of the next chapter that complex mental activities of the sort involved in hypo-

thesis formation and testing do, in fact, presuppose the prior use of a language. But since in a later section of this chapter it will be argued that presupposition of an innate language is neither necessary nor sufficient to account for the ability to follow precedents for speech and to apply them to new cases, I will concentrate in this section on different sorts of difficulties that abstractionism must face.

Let us first remove the temptation to think that abstractionist doctrines should be separated from the use of ostension. Peter Geach cites, with obvious approval, H.H. Price as an example of an empiricist philosopher who wisely dissociates his account of abstraction from "ceremonious ostensive definition"⁷ and quotes the following passage from the latter's Thinking and Experience:

The common features, e.g. "cat", in these otherwise unlike utterances is gradually correlated with a common factor in observed environmental situations which are otherwise unlike. Similarly "black" gradually sorts itself out from another range of utterances which are otherwise unlike, and is correlated with a visible quality experienced in unlike situations.⁸

I would first draw attention to the fact that even Price requires that the use of the expressions "cat" and "black" be learned in some "environmental situation" in which cats and blackness are present. This is hardly an accident. Unless the initiate has mastery of a large part of the language, and so could be told how the expressions are to be used, he would necessarily fail to learn their use. But it is not sufficient that the cat merely be

present in his "environmental situation". If, for example, the cat was present but not in sight, the initiate would still fail to learn the use of the expression "cat". Clearly the advantage of having cats present in the language learning situation is that the initiate can then notice them. What, then, is the point of leaving the initiate to his own devices, to abstract the utterance of the word "cat" from the total utterance and the cat from the other items in his field of vision? It would seem an act of neglect on the instructor's part not to point to a cat, hold one up, or otherwise direct the initiate's attention to the cat. Pointing to a cat or holding one up and saying "This is a cat" would not necessarily give the initiate the ability to apply such a precedent to further instances of cats, but it would make unnecessary the double process of abstraction that Price requires the language learner to go through.

Next, I would draw attention to the fact that Price, in this passage, does not seem to be one of those philosophers who takes common names to be fixed in use by description. To teach the use of the word cat by, say, offering a definition, would presuppose that the initiate already had a grasp of the property terms in which the description is couched. Price treats abstraction as a means of learning both common names and property terms independently. But then what is the language learner supposed to abstract from his environment when he is being taught the use of a common

name? The "common factor", in Price's example, is just the cat. But if that is so, then ostensive acts remove the need for any sort of abstraction on the initiate's part when learning the use of a common name. In particular, a child who has already learned to look where one is pointing, or whose attention can be caught by placing the animal before him, is saved by just those resources from having to go through the process of abstraction. I conclude that ostension should not be separated from the language learning situation, and what is more, once ostension is granted it altogether removes the need for abstraction when common names are being taught. This point ought to be so obvious that we must take some care to see why it is not universally taken for granted.

Perhaps the first point to notice is that the method of ostension can at best single out for the language learner individual instances of a kind, or individual instantiations of a color, and so provides the language learner with no skills for going on to apply the precedent he may then acquire to future cats and future instantiations of that color. In general, it would be unreasonable to require that in the learning situation all the skills needed to use the language are acquired. First, it is patently false that they are acquired in that situation. A child may grow to fear dogs, and fear dogs encountered in the future, it can recognize its mother, and do many other such things before learning to say "dog" or "mama" in order to say what

or who those things are. A child brings all of this background with it to the language learning situation. It will certainly be argued by some that all of these things that a child can do prior to learning a public language are managed only because the child has already gone through the process of abstraction. But to grant even this much is to remove the need for abstraction per se in the language learning situation itself. Ostensively bringing an initiate's attention to some object and then training him to follow a convention for speech is not a defective method just because those acts cannot provide the initiate with the means to go on and apply the convention in future cases. In none of this is there a reason to ostracize ostension from the language learning situation.

There are, however, two more serious considerations which are meant to support both the claim that abstraction is needed for learning language and the claim that ostension could not be a means for replacing the process of abstraction. My general line of objection to abstractionism as a theory to account for the possibility of following, with a certain amount of consistency, the precedents and conventions of language, is that it greatly misrepresents the language learning situation. It is a doctrine designed to solve problems which could not confront the language learner. I will distinguish two kinds of difficulties the doctrine is meant to resolve. — There is sometimes thought to be a metaphysical complexity among the objects

that an initiate encounters, or there is sometimes thought to be a complexity of thought within the language learner, which leaves him unable to eliminate various alternative interpretations that he might give to his instructor's utterances. I will argue that neither form of complexity exists in that situation; and that without one or both of these there is no problem left for abstractionism to solve.

First I will consider the motive of metaphysical complexity. It must be held by abstractionists with such a motive that it is both possible for one person to point to a property of an object and for another to perceive the property to which the first points. This would facilitate language learning but for the fact (so the thesis goes) that every object has (or even is) a collection of properties. The learning of a language is then hindered by the fact that in pointing to the color of an object one is also pointing to its shape, to the object itself and to any number of other things. On this view, although it is possible to point to and perceive the properties of objects, this is never sufficient for us to be able to individuate for an initiate into the language a particular property with respect to which a precedent is being set or taught for saying what it is. Hence, it is only with the presentation of a large number of examples, and the formation and testing of various hypotheses, that the use of an expression is thought to be mastered.

Ostension is of little help in this situation because in pointing to the color of the object I am also pointing to its shape, and many other properties that the object might have. What is worse, my pointing to the color cannot be distinguished from my pointing to the object itself; not only do we fail to make clear to the initiate with respect to which property a precedent is being set or taught, but we fail to make clear to him whether a property expression or a common name is being taught.

My dissatisfaction with this problem rests on rejecting the assumption that the language learner is confronted with a plethora of objects, among which he has to sort and sift before discovering for which a convention for speech is being taught. This is a form of complexity which exists only if both individuals and their properties are possible objects of ostensive acts and of perception. But properties are not objects of ostension or perception. In the way in which we can point to and see an object, we cannot point to and see its color. This is a claim whose defense will have to wait for the fourth chapter, where the foundation for this abstractionist problem can be removed. There is, however, a different sort of objection which can be raised at this time. Before raising it, however, let us note that if properties are neither objects of ostension nor objects of perception, then the language learner confronts a much simpler problem - the problem as to which of the individuals before him is actually being pointed at.

But that is a problem which the language teacher can remove by exercising sufficient caution in the act of pointing or showing.

The assumption of metaphysical complexity has the consequence that it destroys the distinction between the topic and the comment, the act of reference and the act of predication, what we say something about and what we say about it, in any act of speech with the complexity of a claim, assertion, statement or description. Without this distinction it becomes impossible to say anything about something. It is never enough to make a claim or an assertion to bring our listener to know of what we mean to speak, or to get his attention directed to that of which we mean to speak. For such acts of speech, we also need to be able to go on and say something about it. But on the assumption of metaphysical complexity simple descriptions of things, which involve the individuation of the thing to be described by a directional act and the utterance of a demonstrative, become impossible. "This is red" cannot be used to give a description if

- 1) the demonstrative act individuates for the listener both an individual and the color of an individual (this assumes that we can point to the color and see it); and
- 2) the same sentence could be uttered in saying of the color what it is.

If both 1) and 2) are true, then statements of the form "This is F" could not count as descriptions - any more than

could the statement "This is John" or the statement "This is a tomato".

The argument for this is as follows: If a precedent for the use of the word "red" is set and taught by pointing to the color (it doesn't matter how the initiate finds out that it is the color which is being individuated, so we may suppose that he does it by abstraction) and saying "This is red", then what the initiate is learning to do as he learns to follow that precedent, is to say "red" in order to say what that (the color) is. But if the assumption of metaphysical complexity is true, and abstraction is the procedure by which we learn to follow the precedent or the convention for the use of a property term, then it will follow that certain kinds of simple descriptions are impossible to give. A description of the sort I have in mind would be given if I pointed at or held up a ripe tomato and said "This is red". If I am to describe the tomato to you, and you are to be able to come to know that it is the tomato that I am describing, and not something else, then I must be indicating the tomato when I say "this". But if the tomato is what I am talking about, the 'topic' of my comment, then I fail to follow the conventions of the community when I say "is red". This is because the community we are imagining is one in which the language users are confronted with metaphysical complexities, and have established and passed on a convention for the use of the word "red" by pointing to something else; not a tomato, but

a color. To say "red" is to follow the procedure for saying what that is (the color), but not for saying what this is (the tomato), or for saying what color the tomato is. It would then follow that to point to the tomato and say "red" would be to fail to follow the convention for the use of that word. It would be to make the same mistake as saying "tomato" in order to say what that color is.

There are certain inferences that we often make - and make correctly - which look like this:

- i) This is a tomato.
- ii) This is red.
- iii) Therefore, this is a red tomato.

But if we allow both the assumption that language is conventional and the assumption that there is metaphysical complexity which requires abstraction to learn to follow those conventions, then such inferences would never go through, because the referent of "this" could not be the same in i) and ii). If the referent of both was the tomato, then ii) would be false; if the referent of both was the color, then i) would be false. But such inferences are sometimes sound. Therefore, the assumption of metaphysical complexity is false.

To make room in our language for such things as descriptions we must allow that when I point to a tomato and say "This is red", I am not, in uttering the word "red", uttering the name of one of many things before me. If I point to a tomato and say "This is X" it is at least logically

possible for you to misunderstand me by taking me to be telling you what color the tomato is - that is, by taking X to be a color term. But even if you take X to be a color term (or any other property expression) then the entire expression "This is X" acquires for you the status of a description of the tomato. But the only way it can be a description of the tomato is for the demonstrative "this", and therefore also the act of pointing which accompanies the utterance of "this", to be indicating the tomato and not the color. If it were possible to point to colors, then the word "red" could be uttered to set a precedent for saying what that is to which one is pointing. But the result of that would be that no expression of the form "This is F", where "F" would be the name of a property, could be used to describe an object. Such expressions would have the same status as expressions of the form "This is X", and properties of things would be just another group of things. I have already suggested that properties cannot be objects of ostension, so we can now add to that that property expressions are not names of things. The mistake of thinking of property terms as names for saying what something is goes hand in hand with the Platonist picture of colors as detachable from things, floating alongside. The simple trick of seeing all expressions of the form "This is F", whenever F is a property term, as descriptions of what is indicated in coordination with the utterance of "this", makes it impossible to think of the ostensive act

as having a property for its object. The object of the ostensive act must be whatever is being described. So there can be no ambiguity as to which thing we are pointing at - the tomato or its color. The ambiguity can, at best, be one as to whether an expression is being put to use to say what something is, or if it is being put to use to describe it.

Notice, finally, that that ambiguity can be removed from the language learning situation, as it can from any discourse, by adding to the situation the utterance of the expression already taught for saying what the object in question is, before introducing the property term. Hence, a child who already knows that this is called a "ball" and that that is called a "tomato" would not be open to the confusion between a name and a description, if we told him "This ball is red" and "This tomato is red".

To say this much is to concede, happily, the point that ostensive acts coupled with the utterance of a property term are not sufficient for teaching the use of the property term - at least not to a relative newcomer to language use. But this only forces upon us an order for the teaching of radically different sorts of expressions in our language. A child who understands our use of demonstrative expressions has sufficient resources for learning the use of names, since we can indicate to him with respect to what a precedent is being taught for speaking, by pointing to or holding up the relevant object. A child who has already

mastered the use of some names would be in a position to learn how to describe those things, since we can hold up or point to a paradigm, say what it is, and then say something about it. If he already follows us in the performance of the first two acts, then he has part of what is needed for acquiring the third skill - of describing the object, and all that is needed for avoiding the confusion between the use of a name and the use of a property term.

All that remains to be done is to provide an account of how an initiate into our language is to learn with respect to which property of an object an expression is being instituted or taught for describing that object - how the initiate is to learn what to say in order to say what color something is, rather than what shape or size it is.

I will follow traditional metaphysicians in distinguishing between determinables (such expressions as "shape" and "color") and determinants (such expressions as "round", "square", "red" and "green"). Expressions of the second kind can be uttered in giving an answer to questions like "What color is this?" or "What shape is this?", formed with expressions of the first kind. It is clear that if an initiate could first master the use of determinable expressions, and we add to that a mastery of some common and proper names, he would then have sufficient resources for going into a language learning situation and quickly mastering the use of determinant expressions. The combined use of directional acts for individuating a paradigm and the utter-

ance of "This ball is round in shape" or "The shape of this ball is round" could make clear to the neophyte how the expressions "round" is to be used. Unfortunately, this requires means for accounting for the learning of determinable expressions. Contrariwise, if the neophyte could be supposed to first master the use of determinant expressions, these could be used in teaching the use of determinables, by saying "Red, green and yellow are colors" or "Square and round are shapes".

It appears that the prior mastery of either set of expressions would make possible learning of the other set of expressions and would eliminate this last area of possible confusion for the language learner. But given the inadequacy of the abstractionist's characterization of the language learning situation, we do not yet have a foothold which permits us to see how either set of expressions could be taught or instituted. Providing a detailed account of how this problem is answered lies beyond the horizon of our present discussion. When an alternative account of the possibility of judgement is provided in the final section of this chapter, I will show briefly how a solution would arise with the support of that account.

My dissatisfaction, thus far, with abstractionism has been a dissatisfaction with the metaphysical realism or platonism which accompanies it. But if it can be stripped of the metaphysical accoutrements which motivate it, abstractionism would misrepresent the learning situation in

an additional respect. An account of coming to know with respect to which property of a thing a precedent is being set or taught for speaking, which would still admit of the label "abstractionism", would be one which makes the presupposition that the language learner finds himself presented, if not with a multitude of objects of ostension or perception, with at least a multitude of alternatives in thought. That is, the language learner would be portrayed not as a creature confused as to which "object" is being named, but confused as to which hypothesis about the use of the expression is the correct one. The view that the initiate can formulate and test various hypotheses accompanies the metaphysical presuppositions in the abstractionist account, but can be separated from those presuppositions and might be thought to leave a plausible part of the abstractionist doctrine intact. I will give reasons in the next chapter for rejecting even this version of abstractionism, on the ground that it attributes to the language learner a capacity for thought - that is, for thinking that the precedent might be one for speaking of this property, or for that property, or that one, etc. It will be argued there that the ability to think, and hence for this confusion of thought, is incommensurate with the support available to a prelinguistic creature. It follows that if holding any belief presupposes the use of precedents for speech, then a prelinguistic creature cannot entertain an hypothesis as to which property of a thing a precedent for speech is being

set, let alone entertain a number of conflicting hypotheses. The difficulties that a language learner might encounter in grasping and using the earliest expressions he is exposed to cannot be those that abstractionism sets out to resolve.

The arguments of Chapter Three will have to bear the burden of this charge, but the general force of the argument can be seen by examining an example.

The mistake of the abstractionist is, on this account, to attribute to the language learner a complexity of thought which is incommensurate with the resources for thought available to him. A glaring example of this sort of incommensurability is W.V.C. Quine's claim⁹ that a significant amount of hypothesis formation and testing is necessary for translating native utterances of the expression "gavagai", even though the natives might always respond, when asked what a gavagai is, by pointing to a rabbit. The reason that ostension cannot settle the matter, and that even abstractionism can only take us part way towards an adequate translation, is that there are too many competing, and equally good hypotheses available as to the actual object of ostension, or as to which aspect of our stimulus environment might be the referent of the term. Among the hypotheses Quine discusses as a possible translation of the expression "gavagai" in the circumstances described is "rabbit stages" or "brief temporal segments of rabbits". That is, the natives might all along have been pointing to temporal slices of rabbits, or that temporal part of the

rabbit extending from t_1 to t_2 . The linguist, Quine argues, would be making a grave mistake in settling too quickly for the hypothesis "rabbit". Even given the assumption that a temporal part of a rabbit is a possible object of ostension, there are other interesting assumptions which, when given up, would render Quine's translation ineligible. First, notice that we, or Quine, could not have entertained the hypothesis that the object of the native's reference might be a temporal rabbit slice unless we already had the notion of a four-dimensional object, and so unless we have some understanding of the geometry of n-dimensions. Without this prior knowledge, this could not be one of the hypotheses that we, as linguists, could entertain. Secondly, notice that in so entertaining this hypothesis, we are necessarily presupposing that the native speakers have a knowledge of four-dimensional objects, and so of the geometry of n-dimensions. Unless we are willing to make this presupposition we cannot reasonably consider that hypothesis. It then goes without saying that this would not be one of the hypotheses a child must test when learning the use of an expression. He could not even entertain it.

On the support of the arguments of Chapter Three, we can extend this point in a general way. Although we might become confused as to whether a new expression is being taught us for saying what color or what shape an unusual object is, because we can entertain the hypothesis "He might be talking about that odd color, or that unusual

shape, or a temporal slice of that object", this sort of confusion cannot infect the learning situation for a newcomer to the language. If we, as mature speakers, become confused, that is only possible because we already have the resources for becoming so confused, we are already familiar with a number of alternative things one might be saying when one says "This X is F". But those same resources are adequate for extricating us from our confusion - we can ask "Do you mean the color?".

I conclude that abstractionism is untenable, because it is a theory which tries to solve a problem in the language learning situation which does not exist. It is a theory which tries to explain how we manage to avoid certain confusions which cannot arise. Because of its general untenability, it cannot be put forth as an account of how we are able to make judgements as to what to say in order to say what some new object of our acquaintance is, or what color it is. We are now in a position to return to our main business of offering a partial account of how judgement is possible, but before we do that it will be necessary to say a few things about the other leading contender for an account of the possibility of judgement - innatism. Exploring the conflict between innatism and my general assumption of the conventionality of language will lead us to a clearer statement of the use to which the expression "judgement" is to be put.

E. Innatism

The major focus of interest on what I have labelled "judgement" is on the problem of exactly what sort of abilities or faculties are presupposed in judging of some new object of acquaintance that it is an X or that it is F in color. Given the overriding assumption of this discussion - that language is conventional - it will follow that none of the abilities presupposed in judging can themselves be or presuppose linguistic skills. Nor, as I shall argue in Chapter Three, can it be a presupposition of judging that something is an X that we think that it is, because thinking that it is an X presupposes having judged that it is. The limitations imposed on the cognitive capacities of prelinguistic creatures, which all language users are at some stage, are such that we can rule out of an account of the possibility of judgement any cognitive or conceptual capacities which themselves can be shown to presuppose that one is following a precedent for speech. The demand that an account of the possibility of judgement not require any linguistic support follows from our earlier assumption of the conventionality of language. That is a general structure, henceforth to be referred to as R, which cannot be violated without introducing circularity into the account.

We can state this demand, R, in the following way:

Any skills, faculties, or dispositions of a creature cited in giving an account of what makes it possible for that creature to apply an expression taught by the use of paradigms to new objects of acquaintance, cannot, in the initial stages of

language learning, be skills, faculties, or dispositions which themselves presuppose that the creature has or uses a language. Such accounts must be given by citing the nonlinguistic resources of the initiate.

Restriction R might be thought to beg the question against all innatist theories of language learning. If in fact it does exclude reference in one's account to innate linguistic faculties or dispositions, that is only because the claim that there are any such faculties or dispositions conflicts with my original assumption that languages are conventional. I do not believe that the principle excludes outright the view that there are innate grammars, in the way in which it does exclude the Cartesian view that there are innate ideas of, say, color or figure. But this is hardly a matter of begging the question against them. It merely shifts the onus, where it ought to be, to the Cartesians for saying how they come to have the language they use in putting forth their theories.

If we accept Descartes' reduction of innate ideas and principles to certain sorts of dispositions or propensities,

"... in the same sense we say that in some families generosity is innate, in others certain diseases like gout or gravel ... because they are born with a certain disposition or propensity for contracting them ..."¹⁰,

then we can perhaps remove the conflict between Descartes' position and the requirement R. But the cost of doing this is to take away from the Cartesians the right to say that our innate dispositions are linguistic dispositions or

propensities and, I think, we take the substance out of the Cartesian claim that we have or need innate ideas. What is left innate in us could no more be linguistic, and no more be an idea, than is a propensity for gout or gravel.

Attempts to explain the ability of an individual to learn and use a public language by citing the individual's innate resources do not violate restriction R unless the innate resources turn out to be peculiarly linguistic resources. The doctrine of innate languages is as clear a violation as can be imagined, providing we take that doctrine to be claiming that the ability to follow a given convention for speech, and so for applying the publicly learned conventional expression to new cases that we encounter, can be explained by reference to an innate language that we already apply when we enter the learning situation. We must be clear that this doctrine, which is centrally positioned in rationalist theories of knowledge, is only in violation of restriction R, providing that the innate resources are linguistic in nature, and not like dispositions to gout or gravel. Restriction R was introduced solely to exclude circular accounts of language acquisition, but innatism means to avoid simply being circular by making the innate language of a very different kind than the public conventional language. But if it is to remain in violation of restriction R it must do this while clearly being a language. It seems ad hoc to introduce the notion of an innate language to account for the learning of a

public language, but there are, in fact, other problems which have motivated philosophers and psychologists to hypothesize the prior operation of such a language - such as accounting for necessity, or explaining the universal presence of certain structures in languages. Whether or not the hypothesis of an innate language is a good hypothesis for explaining those things is not something that we need to consider. It is only in explaining our ability to follow precedents or conventions, to apply individual words correctly in novel situations, that the hypothesis could conflict with restriction R.

The argument I will give is only meant to show that the hypothesis is, in an important respect, not necessary to explain following and applying - since the hypothesis includes details which are essential to the hypothesis but which play no part in the explanation of this phenomenon, and without which the doctrine would not violate restriction R. And secondly, the argument will show that the doctrine of innatism is not sufficient as an explanation of following and applying linguistic conventions, because the part of the doctrine which appears to make it work as an explanation is merely taken for granted by the account, and not taken to be the actual explanatory factor.

The rationalist account of language learning is a case of over-explanation. Let us be clear that what we need to account for is our ability to make the transition from a paradigm object, for which the expression "X" is given for

saying what it is, to other objects which are of the same kind and say of them what they are by saying "X". What needs to be accounted for is our ability to follow precedents. The only restriction we put on an acceptable account is that it not presuppose the prior mastery of a public conventional language. The skills we require, whatever they may be, need not be infallible - we do make mistakes, and so only want an account which will explain our ability to follow precedents for the most part correctly; that is, we need only an account which would explain how it is possible for us to do this.

The rationalist explains our ability to employ a public conventional language by giving our use of the conventional language the support of a prior natural language. The support for explaining our ability to follow precedents in a conventional language cannot be the support of a further conventional language, because the question arises with every conventional language as to what resources make it possible to follow the conventions of that language with some consistency. So if an innate language is hypothesized to explain this ability, it cannot be a conventional language, for that would require a further language for its support, and so on, ad infinitum. To see what my charge of over-explanation amounts to we shall consider three different hypotheses as to the nature of our putative "prior linguistic resources". A final point to make before proceeding to do this is that I will be somewhat adamant about using the expression prior linguistic resources, or prior language,

rather than innate linguistic resources, or innate language. The phenomenon to be accounted for is the ability to follow precedents set for speech, and those are not precedents to which we are introduced in the capacity of a student of that language immediately at birth. The hypothesis must be only that we bring certain linguistic resources with us to the language learning situation, that those resources are prior to learning the public language. There is no reason to suppose that those resources are with us at birth and not acquired sometime between birth and the time we begin to learn to speak. So I will use the expression prior to capture any resources which we acquire up to the time of our introduction to language, including those we bring with us from birth.

Let us consider, then, three hypotheses for explaining language learning by reference to prior linguistic resources.

Hypothesis One: We bring with us to the language learning situation a prior "vocabulary"; that is, analogs to public expressions, or a representational system.

Clearly, these prior resources are not sufficient for accounting for our ability to follow precedents with some consistency. We might have the analog to an expression, a representation (call it "I"), in our prior vocabulary, which has as its proper extension just the class of objects for speaking of which we have set a precedent of saying "X", but lack the ability to correctly apply "I" with any

consistency. To learn the public conventional language with the support of a prior language the initiate must manage a translation from "X" to "I". But unless he successfully judges that the paradigm is an "I", he will not succeed. In addition, to successfully follow the precedent for using "X" with some consistency, he must judge with the same amount of consistency that a newly encountered object is in the extension of "I". If we have the vocabulary, but lack the ability to correctly use it, we should then lack the requisite resources for following the precedents of our public conventional language. Likewise, having innate rules of grammar would not explain our putting together novel grammatical sentences - we would also need the skills to follow those rules. So Hypothesis One does not give us, by itself, the resources for following precedents for the use of words. Suppose that we add some additional resources to those of Hypothesis One, and call the resulting hypothesis

Hypothesis Two: We bring with us to the language learning situation

- (a) a prior vocabulary or representational system, and
- (b) the prior ability to, for the most part, correctly apply or use that vocabulary.

Parts (a) and (b) of Hypothesis Two provide us, jointly, with what some philosophers would be willing to call prior concepts, whereas Hypothesis One does not provide us with all that is necessary for being said to have a concept.

Surely Hypothesis Two attributes to the language learner resources which are adequate for explaining his ability to follow precedents for speech. When a paradigm is shown, the language learner has the ability to think that it is an "I". The expression "X" is then given for public use. The initiate remembers the expression "X" and that "X" means I. Since the initiate already has the ability to apply "I" across the board, even to objects never before encountered, when he finds another object of the same kind as the paradigm he can already judge that it is an I and, knowing the public convention, he can translate into the public language and say "X".

Hypothesis Two works where Hypothesis One failed because Hypothesis Two claims as prior resources not just a vocabulary or representational system, but the ability to correctly apply those representations and so to form reliable beliefs as to the nature of new objects of our acquaintance. Let us then be clear about what it is that the second hypothesis attributes to the language learner at the time of his initiation into the conventions of a society for saying what something is. He must have something which is the analog of our set of public expressions and he must have the ability to proceed from paradigms to other objects and to know when one of his analog expressions does or does not apply to those objects - that is, to know when an object falls in the extension of one of his analog expressions. But then it is clear that Hypothesis Two says more than is

necessary to explain the phenomenon of language learning and use. To see why I say this, consider a further hypothesis to explain the same phenomena.

Hypothesis Three: We bring with us to the language learning situation the prior ability to, for the most part, correctly apply or use a vocabulary.

Hypothesis Three does not attribute any vocabulary to the language learner, only the ability to apply and use one if he should acquire one. So it does not attribute prior concepts or a prior language, only the skills needed for acquiring concepts or applying a language. These skills, if there are any such, must not themselves presuppose the employment of another language. But that is true of the skills referred to in both Hypotheses Two and Three. They must then be nonlinguistic skills which, if they are to do the work demanded of them in the second and third hypotheses, can be subordinated to the job of following precedents for speech or thought when we have those precedents or when we have a prior representational system. But if these skills are such as to make it possible to employ a representational system, and if we must attribute these skills to explain the ability to learn a language, then the prior vocabulary of Hypothesis One is an unnecessary part of Hypothesis Two. We can defer to those nonlinguistic skills alone in order to explain the ability to follow the precedents of a public conventional language.

After all, if an innate language and the skill to employ it is sufficient to explain our ability to learn and employ a public conventional language, then the skills required for correctly applying the expression analogs of the prior language will be sufficient to account for our correctly applying the expressions of the public conventional language. The real explanatory power of the innate language hypothesis comes from the assumed skills we must have for using the innate language; by adding analogs of expressions to the explanation, we in no way help to explain the acquisition of a public language. But when the representational system drops out of the explanation, and we are left with only the skills for using a language, the explanation no longer can be properly said to involve an innate or prior language. Nor does the explanation given involve prior or innate concepts, unless having a concept is taken to be identical with having certain nonlinguistic skills or dispositions.

I have already shown how, with restriction R, we can divide our mental acts and achievements, our skills and dispositions, into two groups depending upon whether or not the performance of the mental act, the management of the mental achievement, or the exercise of the skill or disposition, presupposes the prior setting of a precedent for speech. I am now in a position to claim with some reason that the innate language hypothesis does not rule out there being any nonlinguistic resources, since it also assumed

them when it assumes that we can employ an innate language. Merely having an innate language would not presuppose such resources, but neither would that explain our being able to acquire a public language. So, innate languages notwithstanding, if we are to explain language learning and use, we must explain it in terms of those dispositions, skills, acts and achievements which satisfy restriction R.

I have argued that abstractionism misrepresents the language learning situation in a number of ways, and so creates artificial problems for the language learner to solve. Innatism is caught between being circular qua account of language acquisition, or avoiding circularity by packing all of its explanatory punch into skills, faculties, or dispositions which are nonlinguistic. This seems to be the place where we must search for the skills required in using and learning a language if circularity is to be avoided.

It should be obvious that not every mental act we perform or mental achievement we manage involves actually using language or presupposes having made some use of language. We must hear a word uttered, or a sound made, if we are to either follow or set a precedent for speech. We must see an object with respect to which a name is to be introduced or taught. We must in some way be sensitive to differences in objects if we are to ever learn the use of a color word or a shape expression. We must have and exercise a faculty of memory if we are to follow any precedent

- that is, we must be able (with practice perhaps) to recall what sound to make or what word to utter. We must have the ability to recognize an individual at a later time if we are ever to learn the use of proper names (this is not to be confused with recognizing an individual as so and so, or as such and such). We can feel pains or itches; be startled by lightning, or frightened by a bear. In addition we can grow thirsty and so want water, or anticipate being fed whenever a bell is heard. In short, the mental lives of nonlinguistic or prelinguistic beings can be relatively rich. It may be arguable whether any particular mental act or achievement is possible for a nonlinguistic creature - philosophers have given reasons for thinking that seeing, wanting, and fearing presuppose the use of language - but it cannot be argued that such creatures have no form of mental life without destroying the very possibility of a conventional language. Every creature who has and uses a language at some earlier time had none. In the next chapter I will argue for one criterion for telling whether or not a given mental act or achievement presupposes the use of a language. Although it is only one criterion, it will prove very useful, since it effectively removes beliefs and knowledge from the prelinguistic and nonlinguistic world.

Since judging is just the mental activity, whatever it may turn out to be in a particular case, involved in applying a word in a future case on the basis of learning from

past cases, restriction R permits us to give a more specific definition: Judgement is the exercise of our prelinguistic faculties or skills, that is, the performance of nonlinguistic acts and the management of nonlinguistic achievements, in subordination to the performance of acts of language or the management of linguistic achievements. We use, in judging, skills which we already have and exercise, but in judgement we are putting them into the service of speech and (propositional) thought.

F. The Possibility of Judgement

If our language is conventional, then we must presuppose that creatures with a language have certain mental faculties and skills prior to learning a language. Without these they could not set precedents for speech, follow others in their practices, and so learn to speak a language. Our final task in this chapter will be to show how our prelinguistic life supports us in learning to follow and apply the conventions for speech to new cases, after having been taught those conventions only by example.

Such an account can begin by keeping sight of a distinction between what we, in giving an account of the development of language, are able to think or say about regularities and consistencies in the world and in behavior, and what initiates or neophytes in our language are able to think or say. Neglect of this distinction leads one to think that an account of how a creature manages to achieve

certain forms of consistency in its behavior cannot be given without at the same time attributing to the creature rather sophisticated levels of cognizance of the regularity in nature. Rather than explain the creature's behavior in terms of the regularity in nature, the temptation is to explain the behavior in terms of the creature's cognizance of those regularities. "How else", we are led to ask when succumbing to the lure of innatism, "can any regularity in the world matter to the creature's behavior?" In reaction to this over-rationalization we are sometimes pushed to the other extreme of viewing them as behavioral automotons pushed and shoved into regular patterns of behavior by regularities in the pattern of stimulus. But to go to that excess is to fail to appreciate the question being asked - the question to which over-rationalization is meant to be one answer. Our question is not what makes these creatures behave in certain ways, but what makes it possible for them to do so. It is the difference between asking how I was able to steal the apple (I used a stepladder) and asking what drove me to do it (I'm a kleptomaniac).

The accounts we construct to answer that question, like the one underway, are given from our position deep within a linguistic framework that is already rich with distinctions and classifications. Hence there is nothing to prevent us from noting regularities among things when giving our account. Because of restriction R, what we must avoid doing is attributing to initiates/into our linguistic framework any form of awareness of those regularities which

presupposes that they have already mastered a language.

We can now begin to see how an initiate, having been shown only a few examples and given an expression "X" for saying what they are, is able to proceed with a great deal of consistency and only a minimal amount of correction, to another X (to what we, from within our linguistic framework, think and speak of as another X), how he is able to recall the right word to utter in order to say what the new object is, and so follow the conventions of his community. The mysteriousness of this ability ought to dissipate when we observe that creatures that do not have, and will never acquire a language, already have and exercise the skills we need in order to make correct judgements. If an animal eats only, say, certain kinds of fruit, then he is able to proceed from one piece to another of the same kind without having to form the belief that this is an apple, and without having to sample everything that lies in its path. But just as he does not have to think that what he is eating is an apple, he does not have to think that it is one of those things (whatever they are) that he ate before. It will be enough if, upon seeing the piece of fruit he anticipates or expects a juicy meal, whereas upon seeing a stone he has no such expectation - the stone arouses none of the same desires, the creature has none of the same expectations, and so exhibits none of the same behavior towards stones as he does towards apples.

Creatures of all levels of complexity show considerable

regularity in their behavior where there is no temptation to suppose that they can do this only because they have sound opinions. Shrimp larvae move towards the light and not into the dark; a hungry dog accepts food but shows no interest in chasing sticks; a kitten chases small moving objects but runs from sudden loud noises. All animals have selective feeding and mating habits, and a few have peculiar hunting and migration patterns. Some of this behavior is simple and stereotyped, but all of it need not be. Some of it is the result of conditioning, training or practice.

The first assumption we must make to account for regularity in behavior is that a creature's makeup is such that it has a relatively consistent response or reaction to relatively similar stimuli. This is not part of the creature's cognitive makeup, only part of his physical makeup. Although there is a vast difference in degree of complexity between the 'behavior' of a quantity of water relative to variations in temperature, and the 'behavior' of the rods of my retina relative to variations in the intensity of incident light, there is no interesting conceptual difference. Both exhibit regular changes correlating with regular changes in the stimulus environment. But the consistent differential response of my retina to differences in incident radiation makes possible consistency in my use of color expressions. Consistent differential responses, however, provide only the minimal support for regularity in purposive behavior. If a naturally aggressive animal, say

a bull, responds to red objects by charging them, we could begin to explain his behavior by discovering, for example, that red light hurts his sensitive eyes. The natural response of an aggressive, large animal might be to charge at things which bother him. There should be no temptation to suppose that he must first think that an object is red, or think that it is irritating his eyes, before he will charge at it. That it irritates his eyes, and that he is a large aggressive creature is sufficient to explain his behavior. Perhaps we would find ourselves discussing the bull's genetic background, or the way in which he was nurtured, if we wanted an explanation of his surly disposition, but citing his surly disposition is sufficient to explain why he exhibits that kind of behavior when he is irritated, and why he does not sulk or flee instead. The irritation of his sensitive eyes sufficiently explains the consistency of his behavior with respect to objects of a certain color.

Suppose that we then conduct a simple experiment in classical conditioning and give the bull an electric shock, which produces a similar sort of irritation, each time that the bull looks at a yellow object. On the condition that the bull's eyes are complex enough to respond in a consistent way when and only when stimulated by the light from bright yellow objects, we can begin to condition a particular response. (Pavlov's dogs could not have been conditioned to salivate at the ringing of a bell if they had all

been deaf, or if they had the same response in their inner ear to every sound.) We might then discover that the bull consistently charges at yellow objects. Still no attribution of thought is called for to explain his behavior; but, in spite of the conditioning model, we are still required to explain his behavior by citing various aspects of his mental life - he feels pain, he has an aggressive disposition, he is angered, and so on. If the bull continues to charge yellow objects even when the artificial irritation is stopped, we would then have to add another ingredient to explain his behavior, say, that he continues for some time to expect or anticipate the irritation when he sees a yellow object. The bull shows all the regularity and consistency in behavior needed to account for the possibility of following correctly precedents or conventions for speech - if he only had the structure for the other sorts of behavior needed to speak. There ought to be no temptation to account for the regularity in his behavior by supposing that some form of thought or knowledge is involved. But then there ought to be no temptation to account for a child's ability to follow precedents for speech by presupposing any different sorts of skills or disposition.

The key to the relationship between a creature's already having a consistent set of reactions to stimuli and regular forms of behavior and its being able, thereby, to judge in consistent ways, lies in the relationship that is established in the language learning situation between an already

well entrenched reaction, or a practiced form of behavior, and the utterance of a certain expression. Differential responses to stimuli make possible discriminatory patterns of behavior, and those prelinguistic discriminatory patterns make possible the following of conventions for speech once they are established and taught.

A creature that naturally, or through training, will eat only X's, already has the prelinguistic skills and faculties for behavioral feats which go into an account of the possibility of judgement. Our having a prelinguistic, shared form of behavior accounts for the possibility of passing from a given paradigm or set of examples to a number of other individuals which we can only now, from our vantage point deep within a language, think and speak of as individuals of the same kind as the paradigm. We cannot, without circularity, account for the coherence of our prelinguistic behavior by attributing ideas or thought to prelinguistic creatures. So the temptation of innatism can be resisted at this level.

It should be no mystery that all those individuals which we call by the same name might have been, say, eaten by our prelinguistic ancestors. It is this prelinguistic grouping which makes possible the linguistic grouping. Our prelinguistic life is a nexus of natural conditioned and learned patterns of consistent behavior. Our preconceptual or prelinguistic behavior does not just happen to involve dealing in discriminatory ways with classes of individuals

which, miraculously and inexplicably, correspond more or less accurately to the classes of individuals we later demarcate within our language by assigning common names and property terms. The correspondence is unavoidable because the classifications we make within our framework of speech are made possible by the prelinguistic forms of behavior - our conceptual and linguistic framework is constructed on the ground of a prelinguistic and preconceptual form of life. The language we create takes advantage of our behavioral practices as we coordinate the making of particular sounds with engaging in particular kinds of behavior, and subordinate the making of certain sounds to the managing of certain achievements.

A child cries when hungry or in pain. But if the child's pain or hunger is regularly attended to when it cries it quickly begins to expect that attention when it cries. By developing enough control over his own behavior, the child can subordinate crying to getting attention, sympathy or help. But crying is only one way of doing that. We very quickly learn that there are conventional ways of doing the same thing - by saying "ouch", or "My foot hurts!" and so on. All the things that we could have managed by crying, moaning, or exaggerating a limp can be done in ways which involve following conventions for speech. It is not significantly more difficult to learn to utter a certain word or words when in pain than it is to cry a little louder when attention is wanted. Asking how a

small child can learn to say "pain" in the right cases is like asking how he learns to say "ouch" in the right cases, and that is as pointless as asking how the baby knows when to cry. Likewise, saying of things of the same kind what they are by saying, e.g. "fruit", when one wants to say what they are, does not differ in important respects from always exhibiting some other form of behavior (like eating) towards objects of the same kind when one wants to eat.

These beginnings provide us with enough of a base to make possible the building up of forms of linguistic behavior not as directly tied to our prelinguistic behavior. If we eat a certain fruit only after it turns red, which we might learn to do only after many sour experiences with the green fruit, or if our community harvests a red flower for ceremonies, then we learn a form of behavior, or learn to make a distinction, which will make it possible for us to institute and use the word "red". The same practices would have been of no help in instituting and using the word "colored", because all fruit and all flowers have some color. But once we have enough determinant expressions, we can use them in order to institute determinables; we can say "Red, green and blue are colors" and "Round, square and oblong are shapes."

I said much earlier that the general account to be given of the possibility of judgement would be an alternative to both the abstractionist and the rationalist accounts. The abstractionist misrepresents the language learning

situation and so over-rationalizes what must be done in that situation. Our account agrees with the innatist in looking for the explanation of the possibility of learning and using a conventional language in the resources that are brought by the initiate into the language learning situation, but denies that any of what we bring into the learning situation is already linguistic in nature. Our account is committed to two theses about language acquisition: First, that in the initial stages of learning or developing a language our nonlinguistic skills, faculties and dispositions are sufficient for enabling us to set and follow precedents for the use of certain elementary classes of expressions. Second, that the process of learning and establishing a language involves a certain stratification of expressions, such that the mastery of one class of expressions will be an additional skill that we can add to our repertoire and which then helps to account for our ability to learn the use of another class of expressions. I have suggested that one such stratification would be that the learning of determinable expressions presupposes the mastery of determinant expressions, and that the learning of determinant expressions presupposes the mastery of some proper or common names. Even though there are a priori considerations which suggest this particular stratification (roughly in the form of transcendental argument from the fact of the use of one class of expressions to the support needed to acquire that use), the stratification ought to be

subject to empirical tests in language learning situations.

The results, thus far, can be summarized as follows.

To speak a language is to follow conventional practices in a community. Following conventions puts our acts of speech into a relationship with acts of speech of others in our linguistic community, but does not thereby strike any relationship between our utterances and the world of which we speak. The notion of "following" must, at bottom, be explicated in terms of notions of natural similarity - mimicry and imitation. But our acts of utterances, taken singly or considered together, cannot reasonably be thought to resemble the world about which we speak - we do not, in that sense, construct a picture of the world by speaking about it.

I have argued for the possibility of learning and using a language without prior cognitive achievements of any high level of sophistication. The ability to proceed beyond the language learning situation and apply expressions learned to novel situations is not an ability which is acquired in the language learning situation itself. The necessary ability is a feature of our prelinguistic lives, which serves us in activities other than language use, and which we learn, in the language learning situation, to subordinate to following conventions for speech. We have nonlinguistic ways of proceeding in which we navigate with considerable regularity of behavior through the world. In using a language we follow these already well-worn paths. Whether

those paths are part of our native constitution, or whether they are cut into us by trial and error is a matter of indifference to the business of learning to use a language. Our ability to follow procedures depends upon the prior establishment of nonlinguistic procedures, but is indifferent to how the latter are established. Those nonlinguistic practices and procedures already have the continuity which our linguistic practices then acquire.

This completes the first part of a theory of judgement - an account of the skills and resources we use in making judgements. The next chapter is concerned with the second part of a theory of judgement - the use of those resources and skills in acquiring beliefs. I will argue that to believe something to be the case is to be in a state - a propositional state, the reaching of which requires the use of our conventional linguistic practices.

Notes for Chapter Two

1. Lewis, David K. Convention: A Philosophical Study, Harvard University Press, Cambridge, Mass., 1969.
2. I will reserve until Chapter Four arguments to show that the class of possible objects of ostension is coextensive with the class of possible objects of perception. Hence, this limitation imposed on my account of judgement will not restrict the account of perceptual judgement to be constructed in the final chapter.

3. Pears, David (ed). Russell's Logical Atomism, Fontana/Collins, London, 1972, p. 56.
4. Kripke, Saul. "Naming and Necessity", in Davidson and Harman (eds.) Semantics of Natural Language, Synthese Library, Vol. 40, D. Reidel Publishing Co., Dordrecht-Holland, 1972, pp. 298-299.
5. Austin, J.L. "Truth", in Philosophical Papers. Oxford University Press, London, 1970, pp. 125-126.
6. Twentieth century rationalism separates rather sharply from its seventeenth century ancestors on this matter. The complaint of the classical rationalist was that the language learning or idea acquiring situation cannot itself account for necessary truths, but at best generalizations. See, for example, Leibniz's Preface to the New Essays, or Descartes' Notes Against A Program. Perhaps because of the greater success of empiricists in dealing with this problem in our century (via the notion of analyticity) contemporary rationalists complain instead that the language learning situation cannot account for the so-called "creative aspect" of language. So rather than innate ideas and principles, they argue for an innate grammar.
7. Geach, Peter. Mental Acts, Routledge & Kegan Paul, London, 1957, p. 34.

8. Price, H.H. Thinking and Experience, Hutchinson, 1953, p. 215.
9. Quine, W.V.O. Word and Object, M.I.T. and Wiley, New York, 1963, Section #12.
10. Descartes. Notes Against A Program, from Haldane and Ross (eds.), The Philosophical Works of Descartes, Cambridge University Press, Cambridge, 1911, p. 442.
11. The three hypotheses are couched in the terminology of "representations" and "representational systems", rather than "expressions" and "languages", following Jerry Fodor in The Language of Thought, which I take to be the paradigm of the view I am criticizing here.

CHAPTER THREE

PROPOSITIONAL STATES

A. Behavior and the Underdetermination of Belief

In the last chapter I argued that any account given of the acquisition and use of a conventional language must be given by citing those skills, faculties or dispositions that an initiate into a language brings with him when he begins to learn the language. In my discussion of abstractionism, in particular the assumption made by abstractionist theories that a prelinguistic creature is capable of entertaining and testing hypotheses, I waived the presentation of an argument to support the claim that pre- and nonlinguistic creatures do not have the resources for coming to hold an opinion, or for acquiring knowledge. In this chapter I will supply some of the considerations which lead me to say that the acquisition of knowledge and beliefs presupposes the use of a language. Doing so will support my criticisms of abstractionism, but will also allow us to lay out more of the machinery which is necessary for giving an account of perceptual knowledge. Judging that the tomato is red is necessary for either coming to know or to believe that it is red, so understanding what must be done (on the judgement side) to acquire beliefs will complete one part of the account of justified perceptual belief.

Occasionally, when someone is said to be thinking something or other, what is meant is that he engaged in some activity - figuring, calculating, deliberating and so on. But sometimes to think something is to have a certain belief or hold a certain opinion. Using this second sense of "thinks", I will assume that "P thinks that S", where "S" can be replaced by a complete sentence, is synonymous with "P believes that S" and "P is of the opinion that S". That to think something in this sense is not to engage in any activity is clear from examples of opinions that are held for extended periods of time. I do not cease to be of a certain opinion merely because I am asleep, or unconscious, and I have held many of the same opinions for a number of years, even during those moments when I was totally absorbed in other activities. It follows from the distinction between thinking as an activity and thinking which is not an activity that whatever the relationship between using a language (even in interior monologue) and having a belief, to have a belief is not to use a language. To use a language is always to engage in some activity, but to have a belief is not to engage in any activity.

We can begin to investigate the relationship between using a language and having a belief by mentioning some of the considerations which motivate the claim that thought is possible without language. Descartes writes, in his reply to the third set of objections, that

in reasoning we unite not names but the things signified by names; and I marvel that the opposite

can occur to anyone. For who doubts that a Frenchman and a German are able to reason in exactly the same way about the same things, though they yet conceive the words in an entirely diverse way?!

It should be noticed first of all that Descartes is talking about the activity of reasoning, and not the holding of an opinion. But his point can easily be applied to the claim that the use of a language is essential for holding an opinion. If a Frenchman and a German use a different language, how can they hold the same opinion on any given matter? But this consideration does not, in fact, carry any weight. In one sense of the expression "what a person says", a German and a Frenchman can say the same thing, but no one supposes for that reason that their saying something does not presuppose that they use a language. If we only report what a man says in direct quotations, then we can never say that a man speaking in French and a man speaking in German have said the same thing. What the German said was "Es regnet", but what the Frenchman said was "Il pleut". But if we report what they said in indirect discourse, we can report that what they both said was that it is raining. If it were not possible to say in one language what was said by someone else in another language there could be no such thing as a translation from one language to another. We must be able to say the same thing, but in a different language. The fact that two people who speak only different languages cannot use the same language but can hold the same beliefs, no more shows the independence

of belief and language than the fact that they can say the same things shows the independence of talking and having a language. It is useful to notice that the only natural way we have of reporting what someone believes is to use the indirect method of reporting. We say that a person believes that ..., and not, without considerable artificiality, that he believes "...".

Another consideration that can be offered in favor of saying that thought and language are independent is the simple fact that we often do say that, e.g., a dog or a cat thinks something, but we do not suppose for that reason that dogs and cats have a language. This is a relatively uninteresting argument for two different sorts of reasons. First, if the fact that we do say something or other is to carry any weight at all, that must be because the cases in which we say it are taken to be the very paradigms for the use of such an expression. But if dogs and cats are said to think, that can hardly be because we take something that they do to be the very paradigm for the application of the expression "thinking". Secondly, if we are to give ordinary language any weight, it cannot be used to settle the matter, because although we often do say of nonlinguistic creatures that they think, it is much less natural to say that they have beliefs, and even less natural to say that they hold opinions. But the sense of "thinking" that we are here concerned with is the sense in which to think just is to have a belief or hold an opinion on some matter; I do not

mean to deny that there may be a good many other senses in which nonlinguistic creatures think.

Certainly the strongest considerations to be given in favor of saying of some nonlinguistic creature that it thinks that something or other is the case are given by citing the sorts of behavior that those creatures are capable of exhibiting. What I intend to argue next is that no nonlinguistic behavior whatsoever counts as evidence for attributing a belief to such creatures. Furthermore, if such behavior was used as evidence for the attribution of a particular belief, that would have the consequence of destroying our grounds for attributing other (nonpropositional) mental acts or states to nonlinguistic creatures.

I can begin to defend these claims by mentioning one of the sorts of examples which has led a philosopher to say that nonlinguistic behavior does provide us with sufficient grounds for the attribution of thought. Norman Malcolm² gives the example of a dog, chasing a cat around the corner of a house, barking and scratching up an oak tree after the cat has, unseen by the dog, run up a different tree. We are permitted, according to Malcolm, to explain the dog's behavior by saying that he thinks that the cat went up that tree (the oak). The example is not filled out in enough detail to give a very good explanation of the behavior of the dog, but it has been filled out enough, I suggest, for us to see that the attribution of a belief to the dog is

neither useful, nor necessary, for explaining his barking up that particular tree.

The verb "thinks", when followed by a that-clause, is an intensional verb³ and, unlike "hit", "barks up" or "scratches", creates a context in which the substitution of co-referential expressions will not preserve the truth of the original sentence. The principle that one of the terms of a true identity statement can be substituted for the other in any statement which is true and produce a true statement does not apply for statements which report or attribute either belief or knowledge. If the dog barked up the oak, then it is true that the dog barked up a beech, because an oak is one kind of beech, and it is true that the dog barked up the Lithocarpus, because an oak is a Lithocarpus. But if I think that this is an oak, it does not follow that I think that it is a Lithocarpus, though it may happen that I do believe that. Because it does not follow from the fact that I think that this is an oak that I think that this is a Lithocarpus, we are forced to hold that the belief that this is an oak and the belief that this is a Lithocarpus are two distinct beliefs - that is, one can hold one of those beliefs and not hold the other, or one can hold one of them and hold the contrary of the other. The question naturally arises, for Malcolm's case, as to which of these beliefs the dog has. Does he think that the cat went up the oak, or that the cat went up the Lithocarpus, or does he think that the cat went up the

tree? Donald Davidson notes this sort of difficulty and comments on it as follows:

These considerations will probably be less persuasive to dog lovers than to others, but in any case they do not constitute an argument. At best what we have shown, or claimed, is that unless there is behavior that can be interpreted as speech, the evidence will not be adequate to justify the fine distinctions we are used to making in the attribution of thoughts.⁴

I think that the failure of substitutivity in belief reports permits us to make a much stronger claim than the one Davidson is willing to venture. What the failure of substitutivity shows is that when we substitute one expression for another in the appropriate contexts in a belief report we run the risk of giving an incorrect belief report, because we may then be attributing to the agent a belief other than the one that he in fact has. As a result, when we introduce new expressions into our language, it becomes possible to hold beliefs which, at an earlier time, it was not possible to hold. Our beliefs can be at least as rich as our language; an introduction of a new expression (of the right sort - one which can stand as one of the terms of an identity statement) makes it possible for us to have a new belief. But what is important for us to note at this stage is that our possible beliefs about an object can greatly outnumber the possible nonlinguistic forms of behavior which can have the same thing as their 'object'. No piece of behavior directed at a certain tree can provide us with evidence for saying that the agent has this belief, rather than some other belief, unless the agent also exhi-

bits some linguistic behavior. Although which belief is held can (perhaps only with a great deal of effort) be said, it cannot be shown nonlinguistically. So if a particular belief is attributed to the dog, no evidence can be produced in the dog's behavior, which gives us reason to suppose that that is what he thinks. But the problem is substantially worse than this. In Malcolm's example, the dog was meant to have a false belief; but what reason is there for thinking that? There could be no behavioral grounds for preferring to say that he thinks that the cat went up the oak, rather than saying that he thinks that he was chasing a bird, and that the bird went up the other tree, and that he flew from the top of the other tree to the top of the oak. How many false or true beliefs are actually consistent with the dog's behavior will, of course, depend upon how much of the dog's behavior we care to take into account - how much of it we write into the example. But no matter how much we fill in, so long as the behavior is nonlinguistic, a little bit of ingenuity and imagination will make it possible for us to find a large number of different beliefs, some of them true and some of them false, equally evidenced by the behavior. It is logically impossible for us to ever come to know, or even to have good reason for thinking, that a nonlinguistic creature has any particular belief.

The story differs only slightly when we try to determine on the basis of behavioral evidence that a nonlinguistic

creature knows that something or other is the case. The expression "knows that" also creates a context in which substitution of co-referential expressions fails to preserve truth. Modifying Malcolm's example so that the dog barks up the tree in which the cat is in fact hiding, we cannot find behavioral evidence for claiming that the dog knows one, rather than another thing.

There are two different ways in which nonlinguistic behavior underdetermines the attribution of any particular belief or bit of knowledge. On the one hand, the underdetermination involves a certain amount of deviousness, simply because with enough ingenuity we can think up a number of radically different, but extensionally non-equivalent belief and knowledge attributions which are equally consistent with, and supported by, the behavior. Perhaps some principle of charity, which involved giving preference to the most consistent, simplest, and the most correct set of beliefs, would eliminate most of the alternative attributions. But then it would be our methodological principle, and not the creature's behavior which would make the attribution of a particular set of beliefs look plausible. To apply the principle is, of course, to presuppose that the creatures have knowledge and beliefs. On the other hand, the failure of substitution in reports of belief and knowledge results in a stronger sense in which nonlinguistic behavior underdetermines the attribution of any particular belief. No deviousness is needed to magnify

the problem, and no principle of charity can eliminate the alternative attributions that are available.

Finally, both of these ways in which behavior underdetermines belief and knowledge attributions, also underdetermine the attribution of a distinction between knowing something, and merely believing it. What evidence could possibly be found, in a creature's nonlinguistic behavior, for preferring to say that it knows, rather than just believes something? There are two sorts of cases in which we would not be tempted to think that dog's propositional attitude might turn out, on the behavioral evidence, to be indeterminate between knowledge and belief. Malcolm attributes a false belief to the dog that barks up the wrong tree, so no substitution into the contexts provided by the belief report would produce an attribution which could count as an attribution of knowledge. Any expression used for identifying the tree which the dog is, in fact, barking up, so long as it is drawn from a true identity statement in which the other term refers to that tree, will also refer to the wrong tree. On the other hand, when the dog is barking up the right tree, and plainly has the cat in sight, it would not be plausible to say that the dog believes, but does not know, that the cat is up that tree. But between these two extremes lie cases in which we could not have grounds for preferring the attribution of belief to attributions of knowledge. If the dog is barking up the right tree, but the cat is concealed among the leaves,

there could be no reason to attribute belief rather than knowledge. The same problem does not arise for creatures with a language - upon asking them their grounds for claiming to know something, we can find reasons for saying that they only believe, and do not know, what they claim to know. Even though the dog does not, in a particular case, see the cat, his sense of smell might put him in a good, or better position for knowing that the cat is up the tree; but we cannot come to know what the dog's position is in a particular case, or what his grounds are in a particular case. In such a case it is impossible, and not just difficult, to know whether the dog knows or merely believes something.

It should be noticed that if we explain the dog's behavior by attributing to him nonintensional states, it is no longer logically impossible for us to come to know which particular state he is in at a given time - which is not to say that it becomes an easy matter. But if there are mental states, the attribution of which are adequate for explaining his behavior, and the attribution of which do not create contexts in which the substitution of identicals fails to preserve the truth of the attribution, then such attributions would not be logically underdetermined by his behavior. So if "wanting", "expecting", "desiring", "fearing", "seeing", etc. have nonintensional uses, then the dog's seeing the tree, say, would be one and the same thing as his seeing the oak and his seeing the Lithocarpus,

and there could be no uncertainty as to which of these he sees. We could still disagree as to whether the dog sees the cat or smells him, but not as to whether he sees the cat or the calico, if the cat happens to be a calico. Further behavior could give us strong evidence for saying that he sees, but does not smell the cat, but no further behavior will give us reason to say that he thinks that the cat went up the oak, rather than that the cat went up the Lithocarpus.

None of this is an argument for saying that nonlinguistic creatures do not have intensional states, but only nonintensional states; but it is an argument for saying that an account of their behavior in terms of their nonintensional states would be a better account of their behavior than one given in terms of intensional states. The question as to which belief a nonlinguistic creature has can never be answered.

B. Behavioral Evidence for Propositional States

There is a further difficulty created by attributing beliefs to nonlinguistic creatures, one which ought to cause more concern. Suppose, first of all, that it is true that behavioral evidence is the only evidence we have for the attribution of mental states to nonlinguistic creatures. Whether we take the creature's behavior to be evidence, or the criterion, for attributing to it any sort of 'inner

life', there will have to be two distinguishable mental states or acts. A horse becomes frightened when a bear comes into the clearing in which it is grazing, but does not become frightened when a rabbit comes into the clearing. We can explain this difference in the horse's behavior in two very different sorts of ways. On the one hand, we can explain the difference by saying simply that the horse is frightened by bears but not frightened by rabbits. This description covers his behavior well enough, but it does not help us to understand why the horse is frightened of one sort of thing and not another. Simply by virtue of attributing to the horse 'fear behavior', by taking the nostrils' flaring, the whinnying, etc. to be the horse's way of showing fear, we thereby attribute to him one sort of mental state. We may then want to go on and explain his fearing bears, but not rabbits, by discovering that horses have a natural fear response to things with certain sorts of odors, or by discovering that this particular horse was previously mauled by a bear. In that way we might find an explanation of why this horse was frightened by this bear. On the other hand we might attempt a very different sort of explanation of the horse's different response to the bear and the rabbit. We might include, as one ingredient in our explanation, the claim that the horse thinks that the bear is dangerous, or that he thinks that the bear will harm him. We might then want to explain the horse's thinking that by showing that he was previously mauled by a bear, or

by hypothesizing that horse's have this particular innate belief about bears.

On the assumption that our grounds for attributing mental states and acts to nonlinguistic creatures are overt nonlinguistic acts, the second mode of explanation raises a serious problem. What would the evidence, or the criteria, be for saying that this horse thinks that the bear is dangerous, or will cause harm? Presumably, if the horse looked up as the bear came into the clearing but showed no fear behavior and went on grazing as the bear lumbered by, there would be no evidence for saying that the horse thought that the bear was dangerous. The evidence that he does think this must, then, be the fear behavior. But if the fear behavior is the evidence for saying that the horse has a certain belief, what is the evidence for saying that he feels fear? Why is this particular kind of behavior fear behavior, and not "belief-of-a-certain-kind" behavior? If the fear and the belief really are two different mental states of the horse, and the grounds for the attribution of any mental state is the overt behavior of the horse, then we must look for two very different pieces of overt behavior to justify our attribution of two very different sorts of mental states. Supposing that all the horse does when it spots the bear coming into the clearing is to whinny, flare its nostrils, open its eyes wide and run, we must decide which subset of this collection of pieces of behavior counts as evidence that the horse is frightened, and which subset counts as evidence that it has a certain

belief. If we take all of this behavior to be fear behavior then we will have no behavioral evidence for saying that the horse has a belief. If we take all the behavior to be evidence that he has a certain thought, then we will have no behavioral evidence for saying that the horse feels fear. On the model in which our support for claiming that a nonlinguistic creature has mental states is all to be given by citing a creature's overt nonlinguistic behavior, we do not enrich the creatures' mental lives by attributing beliefs to them, rather we take away our grounds for saying that they have other sorts of mental states. Once we begin to do this, there is nothing to stop us from denying that they even feel pain - after all, a dog's whine could be taken to count as evidence that he believes that the thorn in his paw is harmful, or might cause an infection. But if we do count his whining and limping as showing that he has such a belief, then there is no longer anything left that we can count as showing that he is in pain.

Perhaps it will be objected that the case has been drastically overstated. That the horse feels fear is evidenced by the behavior we have labelled "fear behavior". But that he has the belief that the bear is harmful is evidenced not by the fear behavior, but by his feeling fear upon seeing the bear. After all, he did not show fear behavior, and so did not feel fear, when he saw the rabbit. That difference must be accounted for - the difference between his feeling fear in the one case but not in the

other. In explaining that difference we can cite his different beliefs about the two cases.

This reply makes an important concession. It no longer takes the explanandum to be the horse's behavior, nor does it take the horse's behavior to be the evidence for saying that it has a certain belief. The behavior is only evidence for saying that it feels fear; the evidence that it has a certain belief is meant to be the fact that the horse feels fear in the one case and not in the other. But this kind of account is, I suggest, idle. No more of the horse's behavior is explained or understood by saying that he has a certain belief about the bear than by saying that he was frightened by the bear, but not by the rabbit. Nor do we give a better explanation of the horse's feeling fear in the one case by citing his beliefs (and then explaining his having those beliefs by looking to his innate makeup or his past encounters with bears), than by supposing that his fear of bears is innate or the result of past encounters. It is by no means clear what role is left for an ingredient of belief in any explanation of the behavior of nonlinguistic creatures. To say this is not to suggest that they have no mental life, but rather that their mental life is sufficiently complex to explain their behavior without the attribution of beliefs and the problems that those attributions bring with them.

In the previous section, I argued that nonlinguistic behavior does not provide us with grounds for attributing

one propositional state rather than another. Nonlinguistic behavior is simply not rich enough to correlate with the subtle distinctions we can make in speech, but the subtle distinctions made in speech can individuate distinguishable mental states. In this section, I have tried to show that it is at best unclear what role the attribution of propositional states to nonlinguistic creatures is meant to play in understanding or explaining their behavior.

C. Indirect Discourse: A Model for Propositional States

My discussion thus far has been negative. I have argued that the fact that having a different language does not conflict with having the same belief does not rule out the need for language in having beliefs. I have argued that the hypothesis that nonlinguistic creatures have beliefs is a bad empirical hypothesis if it is meant to help explain their behavior. The hypothesis has two faults: First, it is meant to explain particular pieces of behavior, but must concede that no nonlinguistic behavior can point to the presence of particular beliefs; secondly, the hypothesis has the consequence of exhausting our behavioral grounds for attributing nonpropositional states to such creatures - the hypothesis could not, I believe, in any ad hoc way, justify withholding belief attributions in explaining every piece of overt behavior that a nonlinguistic creature exhibits. Finally, the hypothesis is an idle hypothesis qua explanation of overt behavior. In reporting the behavior of a

language user, some of our reports will necessarily be intentional if some of the behavior reported is speech behavior; but none of the acts that a nonlinguistic creature performs require that we report them in ways which are intentional - that is, have contexts which do not permit of free substitution. I say that the hypothesis that nonlinguistic creatures have beliefs is idle because it explains behavior by attributing intentional states to them, where the attribution of nonintentional states suffices. By attributing less we can explain as much. I have not argued that the attribution of nonintentional states would suffice to explain any piece of behavior that any nonlinguistic creature might exhibit. I do not believe that a general argument to establish that could be constructed, simply because of the variety of behavior and the different mental states needed to explain it. Rather, one would have to show that it could be done for each particular test case offered. What should be clear, however, is that explanations of behavior in terms of nonintentional attributions would not encounter the crucial difficulty of being logically underdetermined by the behavior. It should be more reasonable to prefer a theory, when one is available, which does not support particular claims which are underdetermined by any possible evidence.

There is one other methodological consideration which ought to be given considerable weight, and on which I will rest my case. The very fact that reports of belief and

knowledge fail to allow for free substitution is a relatively mysterious phenomenon. It is difficult, if not impossible, to give an account of this phenomenon on the assumption that nonlinguistic creatures have beliefs and knowledge. On the other hand, as I will try to show next, the phenomenon can be accounted for with relative ease on the assumption that the use of a language is essential to coming to have a belief or in coming to know something. This, I believe, is the strongest reason for thinking that belief and knowledge (or intensional mental states in general) presuppose the use of language. In order to establish this, it will be necessary to discuss in some detail the phenomenon of intensionality. Doing this will provide us with some insight into exactly what role language plays in belief and knowledge.

Cases of direct quotation of another speaker provide us with very clear cases of the failure of substitution to preserve truth, and with a ready explanation of why it fails to preserve truth. If John says "I am tall", and I report what he says in direct quotation, I must say "John said 'I am tall'", and not "John said 'John is tall'". The reason I must use the expression "I" in the subject place of the imbedded sentence is simply that that is the expression that John used. The failure of substitution to yield a correct report, even though I used an expression which refers to the same person that John referred to when he said "I", is explained by saying that that is not the

expression John used. I suggested earlier that there is no natural analog to direct quotation for reporting a belief; but there is a perfect analog to belief reports in indirect quotations of what people say. Indirect quotations and belief reports exhibit exactly the same structure with respect to the phenomenon of substitution.

First, indirect quotes and belief reports are indifferent, for their correctness, to the language in which they are given. I can say in English what Descartes said on some topic, and I can say what he believed.

Secondly, both indirect quotes and belief reports are of two very different kinds with respect to the phenomenon of substitution. On the one hand, there are reports and quotations which permit of, and even require for their correctness, substitution into the subject position of the sentence following the that-clause, but do not permit of substitution into the predicate position.⁵ On the other hand, there are reports and quotations which do not admit of substitution into either subject or predicate position. I will discuss in detail the first sort of case, and discuss only briefly the second case.

Let us focus on cases of indirect quotation of claims, assertions or statements which were originally made by uttering simple declarative sentences. Logically, the simplest cases will be those that permit substitution in the subject position, but not in the predicate position. Understanding how reports of a class of speech acts work

will require that we first understand how the original speech act works.

To make a claim, statement or assertion we must perform acts which are subordinated to two very different sorts of objectives. In saying this I am assuming that at one deep level of subordination, to utter the same word twice is to perform the same act twice. But the achievements to which two distinct utterances of the same word are subordinated may be very different. Thus, when I say on the one hand "'Blue' is my old dog's name", and on the other hand "The sky is blue", I am at one deep level of subordination performing the same act twice - uttering the word "blue". But at another level, the two performances of the act are subordinated to two very different objectives: Informing you of my dog's name, and telling you something about the weather. When I make a claim, statement or assertion, I must manage two very different sorts of achievements if I am to be understood. I must get you to know of what I am speaking, and I must get you to understand what I have said about it. Any act of communication requires, for its success, a certain amount of 'uptake' on the part of the audience. Although part of the uptake is the responsibility of the audience - they must be paying attention - part of it is also the responsibility of the speaker - he must speak loud enough, pronounce his words properly, and so on. The speaker's acts are subordinated to managing some general communicative achievement like warning or informing some-

one. Doing this requires the speaker to manage two very different sorts of achievements in the cases of claiming, asserting and stating. In successfully managing a complex communicative achievement, like informing you of the color of a certain ball, the distinction between the topic of which I speak (the ball) and the comment I make on the topic (that it is red), has a parallel distinction which we can locate in the subordinate acts I perform and the subordinate achievements I manage.

When I point to the ball and say "This ...", I coordinate those two acts for managing one achievement - getting your attention drawn to the ball. I will follow the traditional practice of calling the act an act of reference, but will label the achievement I aim at as that of individuating the ball. I use the word "individuation" to emphasize that what the speaker aims at doing in referring is to make clear to his listener which thing he is speaking about; his objective is not, in so far as he is performing the subordinate act of referring, to make something known about that to which he refers. If my acts of reference succeed, then the result will be that you 'know' of which thing or person I mean to say something.

I use the word "know", in the previous sentence, with reservations. For you to know of what I speak you need have no propositional knowledge about it. You need not know that it is a ball, that it is red, or any such thing. To suppose that you must know something about it would be

to suppose that you could not pick up that object and sincerely ask "Is this an X?" or "Is this F?" for some X or F, since on that supposition you already know that it is an X, or that it is F. You ought to be able to ask questions like "Is this X an F?". Although you need not know anything about the object for me to individuate it, you might, and I may use what you already know to aid me in managing that achievement. The cases can stretch from the extreme at which absolutely nothing is known, as when we instruct a small child in the use of a word by getting his attention fixed on something by saying "This ..." and pointing, to the extreme at which I get you to know of whom I speak by saying "The Arabic philosopher who wrote The Incoherence of the Incoherence ...". At the one extreme nothing need be known about the thing I point at for me to individuate it; at the other extreme I will not manage the achievement unless you have some special knowledge about Averroes.

An important reason for keeping sight of the distinction between the act of referring and the achievement of individuating is that one and the same achievement (say of individuating this ball) can be managed by performing any number of different acts. I may point to the ball and say "This", I may say "The ball in the corner", or "John's ball". The performance of these different acts could result in the same achievement if (but not only if) this ball is John's ball, and is the ball in the corner. It is no

different than our being able to win the free-style competition by using any one of different kinds of strokes.

Finally, it should be noticed that if I succeed in individuating something for someone, by saying "X ...", it does not follow that the thing in question is an X. We may both think, mistakenly, that the man seated across from us is called "John", and I may whisper something about him to you - "John is wearing a garrish tie-tonight". When I utter the word "John" in referring to the man seated across from us, I can succeed in individuating that man for you, even though that man is not, in fact, John. My objective is the important thing, and it is not the objective of saying who the man is, but of remarking on his tie; but I can reach that objective while being mistaken about who he is.

If I want to make a claim, statement or assertion after performing the first act I must go on to perform another, like saying "is red", or saying "also wrote Kulliyat". At a deep level of subordination, I may here be performing the same act (of uttering a particular word) which on another occasion could have been done in referring to something. But viewed as part of the larger activity of, say, making a statement, the act I perform here will be part of saying something about the ball; it will be part of my comment on a particular topic. I will follow standard usage and refer to this act as an act of predication. If I perform an act of reference, and succeed in individuat-

ing something for my listener, and I also perform an act of predication, it will not follow, even if my listener is paying attention, that I have succeeded in informing him of something. My listener must be able to hear the words I utter, he must understand them, and he must understand what I have said about something. I will call this the achievement of predication, to distinguish it from the acts I perform which may or may not have that result. The achievement of predication is that part of my listener's total uptake to which my act of predication is subordinated. It is, roughly, my achievement of getting my comment on something across to the listener. Managing both an achievement of individuation and an achievement of predication is necessary for me to inform you of something. I can manage either of these achievements while failing to manage the other, and the result of either failure will be the failure to inform you.

It does not follow, of course, that because my acts succeed and both achievements are managed that I get my listener to know something about the object of which I speak, because my statement may have been false. But my managing these earlier achievements is at least a necessary condition of getting my listener to know, in this way, something about the object.

We are now in a position to say how indirect quotations and belief reports form the two important classes mentioned earlier. In our account, we will need to keep separate the

two subordinate acts and the two subordinate achievements a speaker must perform and reach, if he is to have done enough in talking to provide a second person with the needed material for reporting to another what was said. In giving a report to another person, in indirect quotation, let us suppose that the second speaker utters a sentence of the form "P said that S", where P is the name of the original speaker, and S is, like P's original utterance, a simple declarative sentence. Our concern, then, will be only with the acts and achievements of the second speaker which underwrite his utterance of the sentence S; nothing will be said about his use of the words "P said that".

First, it is clear that in some cases substitution into the subject position is permissible. If Frank says to me "Our good friend John is in town", and I report to you "Frank says that John is in town", I have given a correct report; and if you in turn report to someone else "Frank says that Snorkle is in town", because you know that the other person only knows John by his nickname "Snorkle", you have also given a correct report. Your report can be correct even though Frank does not know that John's nickname is Snorkle, and even if Frank believes that Snorkle is someone other than John. Secondly, there are also cases where substitution is required in the subject position. This is most evident when the subject position is occupied by a pronoun. If Frank says to me "I am gaining weight", I cannot give a correct report to another person except by

substituting another expression for Frank's use of "I". I can give a correct report by saying "Frank says that he is gaining weight". In such cases, it would be difficult to convey to another person what it was that Frank said if I had to use the same expression he did in referring to himself, since I cannot report correctly what he said by saying "Frank said that I am gaining weight". The same two sorts of cases can be found in reports of beliefs. If you hold some belief about a particular person, then I may report that belief about that person by referring to him, in the subject position of the imbedded sentence, with a name by which you do not know him, and certainly with a pronoun that you did not use (suppose that your belief is about me); and if you hold some belief about yourself, I can report that belief by saying "P thinks that he ...", whereas you would report the same belief by saying "I think that I ...". By trying to understand why substitution is permitted on some occasions we will be able to see why it is not permitted on all occasions.

Why, in these sorts of cases, is substitution permitted in indirect discourse, but never permitted in direct quotation? Notice that in reporting something that John said by saying "John said that the ball is red", I do not myself say anything about the ball other than that John said something about it; in addition, I report what it was that he said about it. But in doing this, I do not perform an act of predication which has as its object the ball. If I did

undertake to perform the act of predication which I attribute to John, then it would be a fair criticism to level against my report to say that the ball is not, in fact, red. So "John said that the ball is red, but it is not red" would be a contradiction; but it is not. Notice also that it is no criticism of the correctness of my direct quotation "John said 'the ball is red'", to point out that the ball is not red. But it would be if I were undertaking to say of the ball what color it is when I uttered the words "is red". At one level I am performing the same act that John performed - namely, I am uttering the words "is red"; but since I do not, in uttering those words, aim at the same objective as John (the objective of saying what color the ball is), my utterance is not subordinated to the same end. So, although at a very deep level of subordination John and I may be said to have performed some of the same acts, in light of our very different objectives, we have not performed the same acts - he has performed an act of predication but I have not. The fact that I have uttered the same words that he uttered in performing an act of predication, no more means that I have performed an act of predication, than the fact that I wave my hand in the same way that John did when he signaled the waiter, means that I was performing an act of signaling - I was just chasing away a fly. It is an essential characteristic of both direct discourse and indirect discourse that the reporter does not perform all the same acts, nor undertake to manage all the same achievements, as

the speaker reported. If he did, then his act of speech would be open to all the same criticisms as the original speaker's act. Only at one deep level of subordination does the reporter undertake to perform any of the same acts - namely, he may utter the same words. But to utter a word, any word, is not necessarily to either refer or predicate. We can give a report of what another person said by performing all or some of the same acts that he performed at a deep level of subordination, but we cannot give a report by performing all the same acts and managing all the same achievements at every level of subordination.

In this fact we can see the difference between direct and indirect discourse, and the explanation of why substitution is permissible in the sorts of cases now under consideration. In direct discourse, the reporter neither refers nor predicates, and so cannot be held responsible if the predication is false, or if the referring expression fails to individuate. He only aims to utter the same words as the original speaker, and so only tries to imitate his actions at one deep level of subordination. But where substitution is permissible in the subject position, that is because the reporter is himself performing the act of referring, and is responsible (if a successful report is to be given) for managing the achievement of individuation. In indirect discourse, the speaker himself can (in just the sorts of cases we are now considering) refer to that of which the original speaker said something, and in order to succeed in

informing us of what the original speaker said something, the reporter must manage the achievement of individuating for us that of which the speaker said something. Viewed in this way, the permissibility of substitution, and even the requirement of it in some cases, is no mystery. We have already seen that there are many ways of managing the same achievement - many ways of referring which can individuate the same thing. The considerations which govern our choice of words when we refer to something are not semantic considerations - we may refer to the man by uttering the name "John" even though the man's name is not John. The consideration we must take into account in choosing an expression to use are pragmatic, since we aim at managing a certain achievement (of individuating for our listener) and we will succeed only if we use an expression that does that job for the particular audience we happen to have. For different audiences we may have to use a different expression in referring, since it is that audience that we want to get to know of which thing we are speaking.

We can substitute into the subject place in giving an indirect report because the responsibility for managing a certain communicative achievement is ours. Because doing this is the reporter's responsibility, his is also the freedom to use whatever means he judges will be adequate to bring it off. But a balance must be struck between the pragmatic consideration of getting something across to someone else, and the semantic consideration of giving a

correct report. To do that, the reporter must individuate the same object as the original speaker; but that is only to say that he can substitute only "coreferential" expressions into the subject position.

The same considerations show why substitution is not freely permissible, in such cases, in the predicate position. The reporter, we have seen, is not predicating something or other of the object - he is only telling us what someone else said about it. But because managing the achievement of predication is not part of his job, he has no say in how it is done. His job is to inform us of what someone else said about it, not to claim or assert the same thing as the other person. Pragmatic considerations force the reporter to take a certain amount of leeway here - since he does, at least, take responsibility for informing us of what another said about something. But here the pragmatic considerations operate at an entirely different level than before. Where his job is simply to get us to know of what something was said, he has, as it were, the legitimate right to choose his own tools, but where his job is to get us to know what was said about something, then with every substitution he makes, he takes a liberty, and runs the risk of being fairly criticized. If P says "The sea is chartreuse", I may report this to someone who does not know what color chartreuse is by saying "He said that the sea is yellowish-green", when my primary objective is to get my listener to know what color the sea is (suppose I know my listener to be

interested in finding out if the tide contains a particular kind of algae). But in doing this I run the risk of being fairly accused of giving a false report. The original speaker may interrupt to say that chartreuse is not just any shade of yellow-green, but a particularly brilliant shade of yellow-green. Whether or not the original speaker would object to a particular restatement of his utterance will depend upon his evaluation of the use of the alternative expression chosen by the reporter for saying what he originally aimed at saying in using the expression that he chose. So, although some restatement may be permitted, it is permitted on pragmatic grounds which are held to override the need to give a correct report. The important semantical point still holds - correctness is not guaranteed by the substitution of coextensional terms. It is this fact which marks the radical distinction between verbs like "hit" or "scratched", and verbs like "said" and "believes". To avoid the risk of giving an incorrect report (if the correctness of my report is a matter of importance in some speech context), I may only be able to solve my pragmatic problem of informing a listener with a defective color vocabulary by saying something like "He said that the sea is chartreuse, which is, by the way, a shade of brilliant yellow-green".

The semantic considerations which govern the giving of a correct report are, with respect to the subject expression, that I individuate the same object as the original speaker. I cannot be held to informing you as to what he said (what

words he uttered) in order to manage that achievement. But in the predicate position, informing an audience as to what another person said is all that the reporter is doing when he utters the words which occupy that position. That is why substitution into that position cannot guarantee the correctness of the report; to utter a different word would be to fail to report what was said.

It is crucial not to confuse this sense of "what was said" with the very different sense in which what was said in German might be the very same as what was said in English.

Let us call that sense of the expression "what P said", in which what he said could have been said in another language, the fact sense of what P said. In this sense of "what is said" it at least makes sense to say that what P said was a fact. Nothing technical is meant by the word "fact" - only that we can say about what someone has just said "That's a fact!". In saying this we only mean to express agreement with what was said, and not to say what its ontological status is. This, then, is the sense of "what P said" in which what P said could be either true or false. On the other hand, the sense in which what P said, in performing the act of predication was "ist weiss", what he said is neither true nor false, and what he said is the sort of thing which it does not make sense to express agreement with - whether by saying "That's a fact" or otherwise. Nor is what he said in that sense something which could be said in

another language. Let us call this second sense of "what P said" the utterance sense.

The only way that what P said in the utterance sense can be said is, for example, by saying "Der Mann ist weiss". So when it comes to giving an indirect report of what another said, and the reporter comes to uttering the words which will fit the predicate position of the imbedded sentence, if the reporter means to say what it was that the other person said about something, he must say what the other person said in the utterance sense. If he were to fill the predicate position by saying what the other person said in the fact sense, then he would have to fill that position with "P said that the man is white". The resulting report would then be "P said that the man P said that the man is white", which is, of course, nonsense. The nonsense is produced by forgetting that the reporter has already individuated for his audience that of which the original speaker spoke; he did this when he said "P said that the man ...". To have done that much is already to have done part of what is necessary for informing us as to what the original speaker said in the fact sense; it is to have gotten us to know of what or whom the original speaker said something. But since to have done only that much is not to have done enough to inform us as to what was said about the man, in the fact sense, the reporter must do more - he must get us to know what it was that the original speaker had to say about the man. To tell us that is simply to tell us what he said in

the utterance sense. Substitution fails in the predicate position because to utter a different word would be to fail to report what was said about something, where "what was said" can only have the utterance sense.

Earlier, a distinction was drawn between referential acts and individuating achievements, and between predictive acts and predictive achievements. The murky issue of indirect discourse can be clarified by noting that making a claim or a statement is to reach an achievement - the achievement of saying something (in what has been called the fact sense). The distinction between the fact and utterance senses of 'what is said' is an instance of the act/achievement distinction. 'What P said', in the fact sense, satisfies all five criteria for achievements given in Chapter One, but 'what P said', in the utterance sense, satisfies none of them. First, when what P says consists in uttering the words "It is red", what he says can be left half-finished - it has the temporally extended parts of uttering "it", uttering "is", and uttering "red". But what P says when he says that something is red cannot be left half-finished. If he only utters the words "It is", and then is interrupted, he has not yet said anything at all in the fact sense. The problem is not that he has not yet finished saying what color it is, but that he has not yet done that at all. Second, as soon as he says what color something is, in the fact sense, he has said what color it is. Until he has said what color it is, he has not said

anything at all in the fact sense, but in the utterance sense he may have said something (e.g. "It is") before he has said what color it is. Third, there is no felicitous use of the present continuous tense for talking about what someone says in the fact sense. "He is saying something about the weather" can be appropriately said of someone before they have even completed a sentence - before, that is, there is anything said in the fact sense. But if the same report is to be understood as a report of what someone is saying in the fact sense, then it must be understood to mean that he is about to say something about the weather, or that he has just said some things about the weather and is about to say some other things about it. Fourth, to say "It is red" and to say that it is red is not to do two things - to utter a sentence and to do something else in addition. Fifth, and last, we say that something is red by uttering the sentence "It is red", but our saying the former does not consist in our uttering that sentence. This is made evident by the fact that two speakers, speaking different languages, may both say the same thing by uttering two very different sentences. But if what they said consisted in the sentences they uttered, then the same thing would consist of both of the two different sentences. While it may make sense to suppose that one thing could consist of either of two different things, it does not make sense to suppose that it could consist of both of them - unless they are both parts or ingredients of the one thing;

but in that case, neither of the speakers will have said anything in the fact sense until the other has also said something in the utterance sense. So we must utter the sentence in order to say that something is red, and we do the latter by doing the former, though the one does not consist in the other.

Thus to say, as I have done, that what a person said could have been said in another language, is only to say that the same achievement could have been managed by uttering words of another tongue. The reason that what a person said (in the utterance sense) could not have been said in another language is just the accidental fact that to utter those words is not to utter words of the other language. It is, in this same sense, possible that the same thing could have been said in another language, providing that those words happen also to be words of the other language. In the utterance sense of what is said, the same thing could be said in two different languages only if the one language contains foreign homophones of the other.

In recognizing that to say something (e.g., to say what color something is) is to manage an achievement, we simplify the accounts of direct and indirect discourse, and the accounts of why substitution fails in each. In direct discourse we report only what someone has said in the utterance sense (i.e., we report someone's acts of utterance). In indirect discourse, on the other hand, we report what someone has said in the fact sense (i.e., we report an

achievement managed by the original speaker). The achievement of saying what color something is is a complex achievement; we have already seen that the two subordinate achievements of individuation and predication are each necessary for saying what color something is. But the two achievements are also jointly sufficient for managing the higher achievement - that is, if a speaker succeeds in individuating a topic for a listener, and in addition succeeds in predicating something of that topic, then he has succeeded in saying something about something. A speaker who says, for example, that a certain book is red, performs two acts and manages two achievements, but when the two acts (of uttering "That book" and uttering "is red") are co-ordinated with each other, then they support a single complex achievement. If the two acts of utterance are separated by a period of two weeks, then there are prima facie grounds for not treating them as a single, but complex, achievement. In reporting what someone says, we must never manage the complex achievement ourselves. To do so is to open ourselves up to all the same forms of criticism as the original speaker, because our reaching the complex achievement is for us to have said that the book is red. A report given in direct discourse is only a report of the acts of utterance performed by the original speaker; although our report is then given by imitating the same acts of utterance, we do not subordinate our own utterances to any of the same achievements. A report given in indirect discourse is a

report of the achievements managed by the original speaker; in giving such a report, we manage the individuating achievement on our own; the choice of the relevant linguistic resources for managing that achievement is the responsibility of the reporter, so substitution can, and sometimes must be, made in the subject position. Having managed the individuating achievement on his own, the reporter must then relinquish responsibility for any predicative achievement to the original speaker - so he then tells his listener only what the original speaker's predicative acts of utterance were. By not undertaking to manage the predicative achievement on his own, the reporter fails to reach the complex achievement of saying what color something is; precisely that failure saves the reporter from the epistemic responsibilities which the original speaker incurs. Our refusal to permit substitution in the predicate position when giving a report is, then, our way of abdicating responsibility for the predicative achievement.

It should be noticed that this characterization of the failure of substitution in the predicate position does not prevent reports, in English, of what someone said in another language. It does not, that is, require that we report a saying of "Der mann ist weiss" by saying "P said that the man ist weiss" or "P said that the man 'ist weiss'". If we were trying to report in German what P said (in the fact sense), when we came to the predicate part of the imbedded sentence, we would have to report what he said (in the

utterance sense) in the predicate part of his original utterance, by using the expression "ist weiss". On the other hand, if we are trying to report in English what P said (in the fact sense), when we come to the part where we report what he said (in the utterance sense) in the predicate position, we would have to utter the expression "is white". But the reason we would have to use the expression "is white" is not that that is the expression P used, since it isn't, but because we are trying to say what he said (fact sense) in English. What we are providing is a translation, and by giving a translation we can say the same thing that P said in the fact sense, but we cannot say the same thing that P said in the utterance sense. It would be a mistake to think that when we give a report, in another language, of what someone said (in the fact sense) that we are doing something which is incompatible with the general claim that substitution fails in the predicate position of reports given in indirect discourse. We are not substituting "is white" for P's "ist weiss". Translation is not substitution. However translation is to be finally understood, syntactical differences between languages rule out the possibility of construing translation as substitution of words or phrases for words or phrases.

Finally, something must be said about those cases in which substitution fails in the subject position of the sentence imbedded in the that-clause of an indirect report.

I will not attempt a detailed discussion of these cases, but bring them up only to show that the preceding account does not shipwreck upon them, and to set them aside. The sorts of cases that our general project is concerned with (those to be given an essential role in sentences like "P sees that X is F" in the primary sense of "sees that") all turn out to be open to substitution in the subject position.

An example of an indirect report in which substitution fails in the subject position would be the report "President Kennedy said (in 1960) that the first man to walk on the moon would be an American". To be consistent with the preceding discussion, we must maintain that the reporter does not take it upon himself to manage the achievement of individuating that person of whom Kennedy spoke. If this is in fact correct, then we would have to be able to explain why this remains a case of indirect quotation and does not simply reduce to direct quotation. We do not have far to search for the answer. The reporter does not, in these cases, waive responsibility for managing the achievement of individuation on his own because he has a completely different task underway - as in direct quotation, where his task is to utter the same words as the original speaker. Rather, in these cases, the reporter cannot give a correct report if he individuates any particular person or thing, simply because the original speaker did not manage that achievement either. There was no particular person of whom

Kennedy was saying that he would be an American. Although

1) JFK said that the first man to walk on the moon would be an American;

and

2) Neil Armstrong is the first man to walk on the moon

are both true,

3) JFK said that Neil Armstrong would be an American

is clearly false. But 3) is false not because 1) is a direct report, or because substitution fails in general in the subject position, but because when JFK said what was attributed to him in 1) he did not individuate some man for his audience. The reporter who utters 1) does not attribute to him such an achievement, but the reporter who utters 3) does, since he himself individuates a particular man. JFK is not reported in 1) as doing the sort of thing which could have been done by pointing at someone and saying "He" or "this man", even if Neil Armstrong happened to have been standing right in front of him. So when the reporter who utters 3) does just that sort of thing, he does not give us a correct report of what JFK said.

As has already been said, I bring up these cases only to set them aside. Defending a particular account of the failure of substitution in the subject position will only lead us away from our objective. It will be sufficient here to nod agreement with the general approach undertaken

by Keith Donnellan, who writes of these sorts of cases

It seems to me we must say ... that, in the so-called 'intensional' context, the person is not really referred to, but that when we do refer to him, the rest of the sentence says something about him.⁸

Donnellan does not, of course, have the distinction between acts and achievements in mind, and so says "refer" where I would have said "individuate". But I would also have to allow that reference is deleted in these cases - since no acts are performed in subordination to individuation, there is no act of reference. There is, I believe, much promise in Donnellan's positive suggestion that "when we do refer to him, the rest of the sentence says something about him". We can, the suggestion is, take JFK's original utterance of "the first man to walk on the moon" to state a condition, the meeting of which, by some individual, entitles us to apply the remainder of the sentence in talking about that individual. Thus, if the first man to walk on the moon had been a Russian, we could have pointed to him, or uttered his name, and said "If JFK was right, he is an American". In doing that, we individuate someone, but do not suppose that JFK had that man in mind. There is one strong point in favor of Donnellan's suggestion, which we will raise. Cases where substitution is permitted in the subject position differ in an important way from cases where it is not, in this respect: In the former cases, when the subject term is not correctly applicable to an individual who is

the topic of the whole utterance, then the predication is still good. In our earlier example, I said that John had a garrish tie, and I managed the achievement of predication, even though the man I individuated by saying "John" was not, in fact, John. But in cases where substitution is not permissible in the subject position, when the subject term is not correctly applicable to some individual, then the predication is cancelled. In JFK's prediction, the predication is not meant to go through unless it is true to say of some man that is individuated "He is the first man to walk on the moon". The role of the expression "the first man to walk on the moon" is essential, not accidental, to what JFK said in the fact sense, and so for any report of what he said to be correct, it must use that expression.

But this feature is captured by Donnellan's suggestion, on which we might read JFK's claim that the first man to walk on the moon would be an American as, roughly, "Whoever it is true to say of 'He is the first man to walk on the moon', it will be true to say of 'He is an American'". On this account, substitution does not guarantee correctness in either the subject or the predicate place, but indirect discourse remains distinct from direct discourse.

Our extended excursion into the topic of indirect discourse has been essential to reaching our final goal. It can now be seen in a relatively brief space how this discussion is relevant to understanding the intensionality of

belief and knowledge, and to the relationship between language and propositional states.

D. Intensionality and Propositional States

We may draw our conclusion from the last section as follows: The phenomenon of the intensionality of reports of what a person says is an expected consequence of the fact that to say something, whether in the fact or the utterance sense, is to make a use of language. The only way that we can report someone else's characterization of something, without ourselves characterizing it, is for us to report what resources the other person used in order to characterize it. When what we report P as doing is saying something, then the resources he used were expressions. To use a different expression in reporting what was said is, in the intensional context, to report the use of a different expression, and so to give an incorrect report.

But to use a different expression than the original speaker in giving an indirect report of what he said in the fact sense would not produce an incorrect report unless the original speaker's expression was essential to saying just what he said - if he could have said the same thing in the fact sense with a different expression, then we could report the same thing with a different expression than the one he in fact used. What the failure of substitution shows is that the individuation of what is said in the fact sense is contingent upon what is said in the utter-

ance sense. By uttering a different expression the original speaker can say something different, and by uttering a different expression the reporter reports something different said. Consequently, using a particular expression in giving a report is essential to giving a report of what was said.

A crucial point to recognize for our considerations of belief is that to say something in the fact sense is not identical with using a language, but nonetheless using a language - saying something in the utterance sense - is ~~necessary~~ for saying something in the fact sense. The relationship between what is said (fact sense) and the expressions uttered is one of subordination, not identity. I utter the sentence "The ball is red" in order to say that the ball is red; but saying that is not identical with uttering a sentence because I could have said exactly the same thing, about exactly the same ball, by uttering a different sentence (i. e., "This ball is red"). In the fact sense I have said the same thing, but in the utterance sense I have said something different. If we distinguish the two senses of saying something, then the claim that I can say two different things in order to say the same thing, is not paradoxical. We have already seen that to believe that something or other is the case is not to be engaged in any activity, and a fortiori is not to be using a language subvocally. Further, we have noted that there is no natural sense to be given to the notion of giving a

direct report of what another believes; all belief reports and attributions have the form of indirect reports or indirect quotations. When we say what someone believes we say what he believes, to broaden the original notation slightly, in the fact sense. First, what one believes is true or false; second, what one believes can be reported in different languages; third, "what one believes" is not a use of language; fourth, what one believes can be reported only in a form matching the form of an indirect quotation; and fifth, reports of what one believes have exactly the same intensional features as what one says in the fact sense.

My final defense of the thesis that propositional states presuppose a use of language rests on these similarities between what is said and what is believed, in particular the fifth point of similarity mentioned. My claim is simply that the intensionality of belief reports can be explained on the hypothesis that coming to believe something presupposes making a use of language, but it is difficult, if not impossible, to explain intensionality on the assumption that beliefs are pre- or nonlinguistic. This consideration ought to make it more reasonable than not to hold that belief presupposes language, and ought to shift the burden to those who hold that belief is not linguistic to produce an account of the phenomenon of intensionality.

I will defend this claim in two steps. First, I will restate the problem in a way which brings out sharply the crucial point of resemblance between saying something in the fact sense and believing something, and which locates precisely that feature of beliefs which I claim no theory of belief as a prelinguistic state can account for. Secondly, I will criticize three ways in which philosophers in one tradition have tried to break down the essential connection between belief and language.

The problem presented by the failure of substitution in belief contexts can be seen by contrasting two arguments:

- A. (1) P hit the oak
 (2) Oak = Lithocarpus
 (3) P hit the Lithocarpus

and

- B. (1') P thinks that this is an oak
 (2') Oak = Lithocarpus
 (3') P thinks that this is a Lithocarpus

Argument A is valid because (3) follows from (1) and (2), but argument B, in spite of its similarities to A, is invalid because (3') does not follow from (1') and (2'). The reason that B is invalid is that substituting "Lithocarpus" for "Oak" produces not merely a different belief report, but also a report of a different belief. (1') and (3') are different belief reports by virtue of the different expressions used in giving the report. But by virtue of

what are there different beliefs reported? (1) and (3) are different hitting reports by virtue of the different expressions used in giving the reports; but although the use of a different expression produces a different hitting report, it does not have the consequence that a different hit is reported. By using a different expression (in the relevant context) we get a different belief report and a different belief reported. The remarkable feature which distinguishes the two arguments is that the use of a particular expression is essential for reporting a particular belief, but the use of any particular expression is accidental to giving a report of a particular hitting. To borrow an expression from Quine's discussion of the same phenomenon in model contexts, giving a correct report of a particular belief "means adopting an invidious attitude towards certain ways of uniquely specifying"⁹ the belief; or from elsewhere, "this means adopting a frankly inequalitarian attitude toward the various ways of specifying"¹⁰ the belief. This invidious attitude is readily understandable in quotation or reports of what is said since the only way we can report someone else's characterization of something without characterizing it ourselves is for us to report what expression the other person used in order to characterize it. To report what expression was used we must utter the expression in giving our report, hence the invidious attitude. But the obstacle to substitution would be no more mysterious and our attitude towards the use of a particular expression

no more invidious, in reports of beliefs, if we acknowledge that believing something also requires a use of language. Barring such an acknowledgement, this attitude remains mysterious, and truly invidious. Those who hold that belief is nonlinguistic must be prepared to meet the challenge of explaining why the use of a particular expression is essential, and not just accidental, for reporting a particular belief.

Philosophers who have attempted to explain the failure of substitution in belief contexts have, in various ways, been led to the position that to have a belief is to have some relationship or other with an expression or sentence. There is, however, a tradition following Frege which hopes to provide an alternative account of the phenomenon, so I shall examine briefly three ways in which that tradition hopes to break down the essential connection between a particular belief and a particular expression.

On the Fregean alternative,¹² the failure of a context to allow for free substitution of coextensional expressions is thought to indicate that the expression in that context does not have its usual referent; it shifts its reference when placed in an intensional context. Thus, argument B above is held to be invalid because the expression "oak" in (1') refers to something different than the same expression in (2'). On Frege's own formulation, "oak" in (2') refers to the actual oak or oaks, but "oak" in (1') refers to its sense or meaning. The whole sentence in the that-clause,

by virtue of a change in the reference of some or all of its component expressions, no longer designates the same thing either. Outside of an intensional context the sentence is thought by Frege to denote its truth-value, but inside that context it is thought to denote a proposition. The proposition is, in turn, held by Frege to be identical with a thought. The thought is merely expressed by the sentence in the that-clause, and does not itself require for its existence either the sentence or any other expression. As Frege writes in one essay:

The thought, in itself immaterial, clothes itself in the material garment of a sentence and thereby becomes comprehensible to us. We say a sentence expresses a thought.¹³

Our invidious attitude towards the use of a particular expression is then thought to be explained by saying that to use a different expression is to use an expression which, in the intensional context, designates something different, and so does not give us a report of the right thought. There is, however, one very serious flaw with Frege's analysis. On his account, we are not allowed to substitute coreferential expressions in intensional contexts because, though they have the same referent, they can have different meanings or senses - "the morning star" does not have the same meaning as "the evening star", even though both expressions are used to refer to the same planet. On the other hand, the Fregean account does not bar the substitution of one expression for another in an

intensional context if the new expression has the same meaning as the original expression. If the two expressions have the same meaning, then in an intensional context in which the referent of an expression is thought to be its own meaning, the referent of the two expressions will be the same, and so substitution will be possible. ¹⁴ In other words, Frege is committed to allowing substitution providing that the substituted term is a synonym of the term it is substituted for. But clearly, substitution of synonyms in belief reports does not guarantee that the correctness of the report is preserved. Suppose that there is a class of geometrical figures - call them "hexagons" - which are defined as "closed plane figures with six equilateral sides and interior angles equal to 720 degrees". Support further that P does not know the second conjunct of the definition but does know the first. As it turns out, his ignorance is no problem for him, since every and any closed plane figure with six equilateral sides has 720 degrees in its interior angle, and no other figure does. P's curiosity, coupled with his ineptness, leads him one day to calculate that every hexangle has interior angles equal to 740 degrees. Of the next hexangle that P draws we can say

(i) P believes that this is a hexangle
is true. In addition,

(ii) "hexangle" is synonymous with "closed plane figure with six equilateral sides and interior angles equal to 720 degrees"

is true. But in spite of the truth of (i) and (ii),

(iii) P believes that this is a closed plane figure with six equilateral sides and interior angles equal to 720 degrees

is false. But on Frege's analysis, it could not be false. Hence, Frege's account is unacceptable. The unacceptable consequence of Frege's account is that it makes it impossible for persons to fail to have true beliefs about the meaning of the words they use.

The general interesting feature of the Fregean account is that it hopes to account for the invidious attitude that Quine mentions by supposing that the referent of an expression in an intensional context has changed. But the general weakness of such an account is that no specification of the new referent of an expression in such a context will explain why a particular expression is essential for reporting the presence of that referent. David Kaplan¹⁵ shows how a Fregean account can be modified to account for the invidious attitude - by supposing that the new referent of a term when it appears in an intensional context is the expression itself. Thus, in "P believes that this is an oak", the expression "oak" is taken to refer to itself - though perhaps to a different instance or token of itself. This

modification has problems of its own, but they need not concern us, since the modified account concedes the point we want - that what is reported when we report a belief includes an expression, or a use of an expression. As Kaplan writes, "having a name of X, I shall ... take to be essential to having a belief about X". The invidious attitude towards using a particular expression for reporting a particular belief is accounted for in the framework of a theory which connects having beliefs with having a language.

But not all philosophers working in the Fregean tradition have found it useful to allow that to have a belief is to have anything to do with a language. Alonzo Church, for example, takes belief reports to be "statements about certain abstract entities which we shall call 'propositions'" and argues that there is "an insuperable objection against alternative analyses that undertake to do away with propositions in favour of such more concrete things as sentences". Church's argument is directed against Carnap's analysis of belief reports as reports of a relationship between a person and a sentence. The argument against Carnap's analysis of belief is also meant to be an argument against his analysis of reports of what is said, since Carnap also takes those reports to be reports about the relationship between a person and a sentence (one that the person utters). The argument against the latter is as follows: Consider the example

(1) Seneca said that man is a rational animal.

It is clearly false that Seneca wrote or uttered the English sentence "Man is a rational animal". So (1) must report some relationship between Seneca and some Latin sentence, the English translation of which would be "Man is a rational animal". Thus, (1) is analyzed as "Seneca wrote words whose translation from Latin into English is 'Man is a rational animal'", or some variant of this. Church rejects any such attempt at analysis on, roughly, the ground that any variant of the analysis will have an occurrence of a word, like "English", which indicates the language in which the analysis is being given. But that has the unacceptable consequence that a different analysis must be given of the same sentence for every different language in which the analysis is stated. Thus, if the analysis just given was simply translated into German, the appearance of the word "English" would have to be translated "Englisch" and not "Deutsch", and the sentence in quotes would have to remain "Man is a rational animal"; so even the German translation of that analysis could not be understood by someone who has no knowledge of English. The only way to circumvent this seemingly trivial problem, and make the analysis accessible to a German speaker would be to replace the word "Englisch" with the word "Deutsch" and then to translate the quoted sentence as "Der Mensch ist ein vernunftiges Tier". But to do that would be to give a different analysis of the report than was given in English. The German analysis, once

modified so that a German speaker could understand it, is not a translation of the English analysis, since the word "English" is not translated by the German word "Deutsch". This objection shows, I think, that Carnap's analysis flounders on a technical, but fatal, point. The same argument is easily converted to cover Carnap's analysis of belief, by replacing "said" in (1) with "believed", and "wrote" in the analysis of (1) with "believed".

Church's argument does not, however, have any force against the accounts of "says that" and "believes that" which I am offering, for the simple reason that those locutions are here taken to be used to give reports in the fact sense of what is said or believed, and not in the utterance sense. Carnap's analysis of reports of assertions and beliefs suffers from an indifference to that distinction. Although, from the fact that Seneca said that man is a rational animal, it follows (trivially) that Seneca stands in some relation with some sentence, that is not what is being reported when we say that Seneca said something or other in the fact sense. We are, as Church notes, conveying "the content of what Seneca said without revealing his actual words".¹⁷

The constructive suggestions of Church do not, however, seem to be adequate for answering the sort of problem we are posing. If belief reports are taken to be statements about propositions rather than sentences, and some person's relationship to them, we are left with no understanding of

Quine's "invidious attitude". There is no clear reason why one and the same proposition cannot be reported by various different expressions in the same language, and so no clear reason why substitution should fail if the substituted term is a synonym of the expression it replaces. Why should reports which include reference to an abstract entity fail to permit free substitution, when reports which include reference to concrete entities do not? For reports of what is said, what Church takes to be a statement about a proposition we take to be a statement about a person (Seneca), some activities that he performed (namely the acts of uttering or writing words), and some achievements he managed (namely the communicative achievement of making an assertion or statement and the subordinate achievements of individuation and predication).

There is one final attempt to break down any essential connection between having a belief and being a language user which we will consider. In the article already mentioned, David Kaplan argued that to have a belief about X one must have some name of X. Kaplan means "name" to have a very technical sense, and in the penultimate section of his article, he suggests that to account for the full range of beliefs that people have we must include visual images among names. "If we cannot even say it with words, but have to paint it or sing it, we certainly cannot believe it with words". It is doubtful (but unclear) whether Kaplan would think that an image or collection of images is suffi-

cient for constituting a belief, since one would suppose that I could have images of things about which I have neither knowledge nor opinions. But if having an image can play a necessary but not sufficient part in having a belief, then his suggestion would not conflict with our general thesis about the priority of language to belief providing that one of the other necessary ingredients for belief is a language. There is, at any rate, a significant problem with giving images the status of beliefs - images, or the having of any particular image, can be reported without creating an intensional context, but beliefs cannot. My believing that there is an oak in the backyard cannot be analyzed in terms of my simply having an image of an oak standing in the backyard, because to have an image of an oak in the backyard, whether the image is in my mind or in my wallet, is one and the same thing as my having an image of a Lithocarpus. If beliefs are to be intensional, then some other ingredient must be added to the having of an image to account for the intensionality. We have already seen that there is good reason to think that the extra ingredient must be an expression for a language. Having an image of a word, or an image which is labelled "oak", might play some role in having a belief about an oak, but having an unlabelled image is no more like having a belief than having a picture of a man is like having an opinion as to what his name is.

I stand, then, by my original claim that to have an

opinion presupposes having used a language. In no other way do we account for the intensionality of belief reports and attributions. Our attitude towards using a particular expression, and barring the substitution of coextensional terms or synonyms, ceases to be an invidious attitude only if we acknowledge that what is being reported when we report a belief is, in part, a use of an expression. Keeping sight of the distinction between the fact and utterance senses of what a person says, and using the fact-sense of what is said as our model for what a person believes, allows us to reject a host of criticisms against the thesis that believing something presupposes using a language. Having already given a somewhat detailed account of the workings of indirect quotation we have done most of what need be done to explain the intensionality of belief reports. The point at which indirect quotation breaks down as a model for belief reports is simply in the fact that to believe something, or to come to believe something, does not involve uttering any expressions. In the next section of this chapter I will try to show how forming an opinion can involve using a language, even though it need not involve uttering any expressions.

E. The Role of Language in the Acquisition of Beliefs and Knowledge

Thus far I have argued that it is more reasonable than not to hold that believing something presupposes language.

Nothing has been said yet to indicate what the role of language in belief is, only that it has a role, and that the role does not consist in identifying belief with using a language subvocally. I will now provide just enough of an account of the role of language in belief to make good my claim to provide an understanding of the phenomenon of the failure of substitutivity in belief reports, and to show that believing something is entirely on a par with saying something in the fact sense.

The assumption of the conventionality of language led us to say that a creature which has and uses a language must have and exercise a number of nonlinguistic skills. To follow the conventions for speech of our linguistic community we must employ these nonlinguistic skills, and in using a language we must subordinate the performance of nonlinguistic acts and the management of nonlinguistic achievements to the performance of acts of language and the management of linguistic achievements. To say that a particular tree is an oak, we must utter the word "oak" and a necessary (but not sufficient) condition for what we say to be true is that that particular tree be one of those kinds of things of which there is a convention in our language of uttering the word "oak" in order to say what it is. If we aim to succeed in saying what something is, or what color it is, we must make an effort, however slight, to maintain continuity with the practices of our linguistic community. To utter the correct word we must utter that word which our

community has trained and educated us in using for that job. The failure to be able to utter the correct word on some occasion, even if we have already been taught the proper word, shows us that one of the prelinguistic faculties we must have to maintain continuity with our community is the faculty of memory. My failure to recall the right word on a given occasion cannot be the failure of a skill, the very exercise of which already presupposed using a language, simply because that prior use of language would itself presuppose that I had recalled the very words which I used. There are various activities that I may perform in order to recall something - I might recall where I hid something by closing my eyes and picturing the room in which I hid it, or picturing myself going about searching for a place to hide it, or I might recall someone's name only after running through the letters of the alphabet. But remembering or recalling a word is an achievement which is managed by doing these other things. I do not remember John's name if, after considerable pondering, I say "Joe". In order to say "Joe" I must recall a name, but not the right one for saying who this man is. Any use of a language involves recalling expressions, and I will give memory an important role in the use of language which is essential for coming to a belief.

I can recall a word, whether the right word or the wrong word for a particular job, without uttering the word. Recalling is a mental achievement which is essential for

any use of language since it is a mental achievement which allows us to maintain continuity with the conventions of our community. I suggest then that a necessary condition for thinking that this tree is an oak is that one recall the word "oak" - which is not to say that one has or ever will utter the word. It is this necessary condition which is responsible for failure of substitution in reports of a belief. To think that this is an oak is not the same as to think that this is a Lithocarpus, because to think the former one must have recalled the word "oak", but to think the latter one must have recalled the word "Lithocarpus". Since we have already seen that believing that something is the case is not to perform an activity, it follows that whatever activity one might have performed in subordination to managing the achievement of recalling the right word is not identical with the belief. Nor is the achievement of recalling the word to be identified with having the belief, since that achievement is durationless, but one might have a belief (or hold it) for any length of time. My only claim is that recalling an expression is a necessary condition for having the belief. The role of the achievement of recalling an expression can only be that doing that is part of what is involved in acquiring a belief. The expression recalled plays the role, I suggest, in what is believed, that the expression uttered plays in what is said.

But clearly to recall an expression is not sufficient for acquiring a belief. Unless some sense can be made of

recalling the correct or the incorrect expression, no sense can be made of having a true or a false belief. Just as a necessary condition for uttering the right or (the wrong word is that there is some objective aimed at in uttering the word, so a necessary condition for recalling the right or the wrong word is that there be some objective aimed at in recalling that word. This is only to say that the achievement of recalling an expression must, if a belief is to be acquired, be subordinated to some further end. Since the words that we recall are public words - words instituted and taught for public speech - recalling any word can be subordinated to a communicative achievement. But there is no reason to suppose that every word recalled is recalled with the motive of performing some public act of speech. The ends to which one recalls an expression, in such a way that recalling that expression is part of acquiring a belief, can be almost as various as the ends for which one utters an expression as part of making a claim, statement or assertion. I may acquire beliefs or knowledge as to what is in a box in order to tell someone about it, or I may acquire the belief or knowledge in order to decide what to do with the box. Then again, I may recall an expression for saying what this is in order to know what it is (is this a lily or a dandyion?) and I may have any of various reasons for wanting to know what it is; that is, any of various objectives to which knowing or thinking something can be a means. Before recalling an expression can be part

of acquiring a belief, it must be subordinated in some structure of mental acts and achievements, just as uttering a word must be subordinated in some structure of linguistic acts and achievements before it will count as part of saying something, asserting, stating or claiming something.

As Zeno Vendler writes,

In thinking about something one goes through a series of mental acts, often involving some changes of mental states: one may guess, assume, realize, or conclude that something is the case; regard, consider, or view a certain thing in many ways; contemplate, plan, and decide to do one thing or another; and wonder about consequence. One might also adapt beliefs, give up suspicions, modify intentions - that is, "make up" or "change one's mind" in these and many other ways. The idea that one might be thinking about something without performing any of these or similar acts is as incomprehensible as the idea of talking about something without saying anything at all.¹⁸

Vendler is talking about thinking about something, and not about thinking that something or other is the case, but he provides us with a perspective in which the mental achievement of acquiring a belief fits into a nexus of mental activities. What should be kept in mind is that acquiring a belief can also fit into a structure of public activities, as when we form an opinion in order to make a claim, cast a ballot, or find a missing object. The thesis which I am defending is only the thesis that recalling an expression only plays a part in forming or acquiring a belief when that achievement is situated within some structure of subordination.

Our discussion in this chapter has focussed on the

notion of belief. It should be clear, however, that what has been said about belief has, with some modification, a much wider application. The aspect of belief which we have centered the discussion on is the intensionality of belief reports and attributions. But reports and attributions of knowledge share with reports of belief the feature of intensionality. The conclusions we have drawn about the relationship between acquiring a belief and exercising our natural faculty of memory in recalling a word apply across the board to the acquisition of knowledge. As with belief, neither the activity of pondering nor the achievement of recalling is identical with, nor a constituent part of, knowing something; they are only essential constituents of the achievement of acquiring knowledge. In coming to know or believe that something is an X, one makes a judgement as to what it is, and to do that one must perform mental acts and manage mental achievements which include, at one level of subordination, recalling the expression "X". In doing this one maintains continuity with the practices of a linguistic community, in the resources used for thought as well as speech.

We have already seen that a necessary condition for the truth of what one says is that one correctly follow the conventions of a community. The same condition is necessary for reaching a true opinion or acquiring knowledge. To judge this to be one of those things of which there is a convention of saying "X" for saying what it is, and so for

reaching the opinion or acquiring the knowledge that it is an X, requires that this thing actually be one of those things for which there is a convention of saying "X" for saying what it is. If it is not one of those things, which we identify for our purposes only as "one of those of which there is a convention of saying ...", then we can at best acquire a false belief. Important epistemological problems arise when we ask how acquiring a true belief differs from acquiring justified belief as to what something is. In the following chapters we will examine one of the ways in which one acquires justified belief and not merely opinion.

In this chapter we have provided, through considerations of the intensionality of indirect discourse as a model for understanding the intensionality of belief, support for a very general principle which we can now state. In general, if reports or attributions of any particular mental act, achievement or state have contexts which fail to allow for the free substitution of coextensional terms or synonyms, then performing that act, managing that achievement or being in that state, presupposes that the agent of the attribution used a language in performing the act, managing the achievement or reaching the state. The intensionality of an attribution is only explained by supposing that the act, achievement or state attributed is distinguished from some other act, achievement or state by virtue of involving in an essential way a particular expression.

There is, in general, much leeway for debate as to

whether a particular attribution has the feature of intentionality; but if it is agreed that some attribution has that feature, then it should be agreed that the act, achievement or state attributed involves in an essential way the use of a particular expression. Having beliefs and having knowledge are central cases of propositional states, that is, states the acquisition of which presupposes a use of language. Perhaps having a sensation falls as clearly as any state on the nonpropositional side. I have argued that intensionality is a sufficient test for propositional states, but not that it is necessary. We are left with no clear test which allows us to decide of a large range of mental acts and states whether or not they presuppose linguistic resources. There is some reason to think that a wide range of emotions (including jealousy and envy) presuppose beliefs and, if so, having such emotions presupposes having been initiated into the linguistic practices of a community. States like "wishing that S", "fearing that S", "hoping that S" and "expecting that S" seem to have the feature of intensionality, and so to be post-linguistic. But there is no immediate reason to think that "being frightened by an X", "expecting X", "anticipating X", "recalling X" or "recognizing X" are intensional, and so no immediate reason to think that they are not pre-linguistic. In the end, however, each of these cases must be decided on its own. If the general approach to language learning presented in the last chapter has any merit, then

there is reason to think that there will be a large number of mental acts, achievements and states which fall on the prelinguistic side, and so which form part of the natural support we have when we come to acquire a conventional language.

What I have tried to show in this chapter is that there are good reasons for thinking that the acquisition of belief and knowledge requires the support of linguistic resources. The considerations which lead me to say this are of two different sorts.

First is that on its own, the hypothesis that a nonlinguistic creature has a belief or a particular piece of knowledge is unfounded behaviorally - the attribution of a belief or a piece of knowledge is logically underdetermined by any behavior which that creature is capable of exhibiting. The question as to what a nonlinguistic creature believes is undecidable. On the other hand, the question as to what a nonlinguistic creature is frightened by is, on any given occasion, decidable. On methodological grounds, if a particular piece of behavior is explained either by saying that the agent was frightened, or by saying that it has some belief, and the agent is nonlinguistic, then we ought to prefer the attribution of fright to the attribution of belief.

The second consideration carries more weight, and at an absolute minimum, shifts the onus of the argument to those who wish to attribute knowledge and belief to nonlinguistic

creatures to provide over-riding reasons to do so, and to provide us with an account of intensionality of which does not entail that a creature with intensional states has a language. The phenomenon of intensionality is a genuine anomaly in a theory of nonlinguistic beliefs and knowledge, and our refusal to allow free substitution into belief and knowledge attributions is genuinely invidious. But on the hypothesis that belief and knowledge acquisition involves a use of expressions, the anomaly disappears since intensionality is then an expected consequence; and the refusal to permit free substitution ceases to be invidious since that refusal would then become the simple demand for a correct attribution. It is more reasonable than not to accept, among competing hypotheses, the one which rids us of anomalies and invidious attitudes.

Our problem is not to read the minds of other creatures, but to decide what is most reasonable for us to say about them, and I have tried to show these methodological considerations are unequivocal on that matter.

To judge that something is a such and such, or is such and so in color, is to exercise certain of our native skills, but to do so in subordination to following conventions established in a linguistic community for speech. To acquire a belief or to acquire knowledge is to manage one species of achievement, the management of which presupposes that we exercise the skills that our community trains us in. Just as the use of language for speech requires the

support of our native skills and faculties, from being able to make a sound or mark to being able to recall which sound or mark to make, so the acquisition of knowledge and belief requires the exercise of some of the same skills, plus the linguistic resources whose use our community trains us in when we learn to speak. Like a fighter who, through training acquires certain skills, acquires the title, and then holds it for a given length of time, we acquire certain skills, exercise them in order to acquire a belief, and then hold it for a given length of time. Holding a belief is no more to be identified with the acts we perform, or the achievements we manage, than being the champion is to be identified with fighting a bout or scoring a punch. Holding a belief is to acquiring a belief as being the champion is to winning the bout. In the most general, noncommittal senses, to have a belief or to be a champion is to be in some state. What I have tried to argue is that to reach the state of having a belief one must have and exercise certain linguistic skills, and manage certain linguistic achievements. In the remaining chapters, I will argue that perceiving something does not involve the use of linguistic resources, and so that to perceive is not to either know or believe something; then I will argue that sufficient skill in using our linguistic resources allows us to reach, in one kind of perceptual circumstance, the propositional state of knowledge, and not merely belief.

The oldest and longest philosophical tradition takes knowledge and belief to be prior to language, and language to be prior to the formation of society. On the view presented thus far, the order is reversed - a community, however small or poorly organized, sets precedents and maintains conventions for speech, and these conventions provide us with the resources for acquiring knowledge and beliefs.

F. Language and the World: The Case of Proper Names

Earlier I drew out and emphasized a difference I have with Kripke concerning the so-called causal relationships which our speech has to the things about which we speak. Now that the theory of judgement which I will put to use for epistemological purposes is complete, it will perhaps be thought that the difference I have with Kripke comes back to undermine that theory of judgement. The charge against the account of judgement which I have in mind is that my disagreement with Kripke makes it look as though I have, in effect, cut myself off from the possibility of using the 'causal theory of names' for one of the most important purposes for which it was devised - namely, to promote an understanding of the relationship which obtains between a name, or a speaker's use of a name, and the thing or person of which it is a name. Dennis Stampe has pointed out that we have causal theories of one sort or another for knowledge, memory, belief, evidence, proper names, common

names and reference. What these phenomena have in common which leads us to construct causal theories of them is, in Stampe's words,

the fact that they involve a relation to an "object", of one kind or another. It is, after all, this relation, this "involvement of an object", that we are tempted to understand as some kind of causal relation. Further, the entities that share this property of involving an object, may be conceived as sharing a common nature, for in each case, that which stands in that relation may be conceived as being a representation, or a psychological state of representational character, their various "objects" being what they variously represent. If this is so, the thought suggests itself that the truth underlying our attraction to causal theories is that representation is an essentially causal phenomena.¹⁹

It seems that I have cut myself off from the possibility of putting a causal theory of language to use in explaining the representational character of speech, and so of thought. The relationship which ties our use of words in a language to other speakers - in particular, to those speakers who introduced the words into our language with some quasi-baptismal act - is the relationship of following conventions. But that relationship goes back only as far as the first speaker in a chain of speakers - it does not reach back beyond the first speaker to the thing itself, and so is not a relationship between our use of a word and the object. By emphasizing this fact, as I have done earlier, the account of language, and so of judgement, which I am employing, appears not to be "anchored" to the world. What certainly does follow from the preceding account of lan-

guage is that the relationship between our speech and the world, as exemplified by the relation "X is the name of Y", cannot be explicated as the relation of "following". Since part of what I announced as the anti-Cartesian intent of this project is to break down the picture of thought and language as human artifacts, constructed by us to picture or represent the world to us, it will be important before proceeding any further to show how the account of judgement does this, and yet how the account also succeeds in explicating the sense in which our talk is about the world. The general idea all along has been to replace talk about language as a product of our activities with talk about language as a complex set of activities. We have, I claim, all the pieces in place for doing this by providing an understanding of the intentionality of language - the so-called relationship between language and the world. I will provide the account for the case of proper names. The extension of the account beyond proper names to other "referring expressions" is by no means trivial, but how it would proceed should be clear from the case of names.

First, the problem of what the 'naming relationship' is must be separated from the problem of what a name is, since a name may be used but not be in the purported naming relation. The name "Aristotle", in the sentence "Aristotle is a name" is not used as a name of Aristotle, even though it is his name. The sentence "Aristotle was a student of Plato" illustrates the use we want to understand - the name

"Aristotle" may be said, in that sentence, to name the man. But the use to which the name is put in that sentence has already been discussed - there the name is uttered (or written) in subordination to individuating Aristotle for the listener (or for the reader). That is to say, it is used in subordination to bringing the listener to know whom the rest of the sentence then goes on to say something about; its use is not subordinated to bringing the listener to know or believe anything whatsoever about Aristotle. The latter objective, if it is one, is an objective, the reaching of which depends upon the use of the entire sentence, not on the use of the name.

Secondly, it is important to separate concerns about how it is possible for us to use an expression as a name from concerns about what kind of use we make of an expression when we use it as a name. In insisting on a distinction here, what I am insisting on is only that we be prepared to provide answers to two very different questions in subordination to completing a theory of names. They are:

1. What is the use to which an expression may be put such that to put an expression to that use is to use it as a name?
2. What condition must be met before it is possible for us to put an expression to use as a name (i.e., to put it to the use cited in answer to (1))?

I will refer to that part of a theory of names which is given in answer to (2) as a statement of the material conditions for the use of a name. This is, in short, the distinction between saying what job names do in our language (giving the formal conditions) and saying what facts about language learning and use make it possible for an expression to do that job (giving the material conditions). No one would deny that there is a distinction between saying what role singing plays in our lives and saying what it is that enables us to sing; but that is the same distinction which I am employing here for our discussion of using expressions as names.

This second distinction is important to keep before us since, for example, the things that enable Aristotle's acquaintances to use an expression as his name (which might include their having seen him) and which make it possible for us to use the same expression as his name (which would not include our having seen him) might be very different. But these differences, however great they turn out to be, are differences only in the material conditions for the use of an expression as his name, they are not differences in the formal conditions. The formal conditions for the use of his name must be the same for us and his acquaintances, since we all use the expression as his name.

In light of these distinctions, and in light of the account of judgement which has preceded, I suggest that for an expression "X" to be a name of X in a language is for

the use of that expression to be sufficient, by virtue of the conventions of the language, for individuating X for speakers of the language. This is not meant to be a statement of the material conditions for the use of an expression as a name, but only of the formal conditions. To say that an expression is a name is to say only that its use is sufficient for individuating someone, and so to use an expression as a name is to use it for individuating someone. To use it in this way is to use it independently of the use of other linguistic or nonlinguistic resources for individuating a person, since the use of it alone is sufficient for individuation.

On this account of the use of an expression as a proper name, demonstratives like "this" and "that" are not conventionally proper names because their use is not sufficient for individuating something - individuating something with the use of demonstratives depends, in addition, on coordinated use of ostension and on perceptual achievements. Likewise, other sorts of referring expressions which are commonly subordinated to individuating something for a listener - expressions like definite descriptions - are not, on this account, proper names, since their use for individuating something depends on the listener having certain beliefs about who those descriptions fit (or are meant to fit by the speaker). Demonstratives, definite descriptions, and other noun phrases are not, then, sufficient for individuating something.

This thesis about what it is for an expression to be a name in a language, and about what it is to use an expression as a name, is a thesis about the intentionality of names - and so, the intentionality of an important part of language. How it is a thesis about intentionality can perhaps best be seen by contrast to the view which it commits me to rejecting. I reject the view which would say that by uttering a name, say "George", we are able to individuate George because there is some prior relationship (the naming relation) which obtains between George and his name. My claim is simply that the 'prior' relationship which obtains between "George" and George which enables us to succeed in individuating the man by uttering the name is not the naming relation; rather the prior relationship is one which constitutes the material conditions for the use of his name as a name of him. What the naming relation is - the relation between the man and his name which is expressed by saying that "George" is the name of George - is given by citing the formal conditions for the use of the expression as his name. So, we do not succeed in individuating George by saying "George" because "George" is the name of George; rather "George" is the name of George because we can succeed in individuating George by saying "George". (In the previous sentence, the two occurrences of the word "because" must be understood as follows: In its first occurrence, it indicates the order of causation; in its second occurrence it indicates the order of analysis.) The only prior rela-

tionship which obtains between George and his name - the causal relationship - which enables us to use the expression as his name, is not the naming relation.

Saying what the use is which is made of an expression such that its having that use is what is meant by saying that it is a name in a language exhausts saying what the naming relation is. But, on the account which I have given, an expression is the name of something in a language if the use of that expression is sufficient, by the conventions of the language, for individuating that thing. But that relation is not, in the end, a relationship between "George" and George; it is a relationship between the act of uttering (or writing) the name and the achievement of individuating the man. The relationship between the act and the achievement is that, by the conventions of the language, performing the act is subordinated to reaching the achievement.

We can, in many but not all cases, individuate something without uttering any expressions - simply, for example, by holding up and pointing to something. If we couldn't do this sort of thing, we could never get a language off the ground or teach it to others. When we have a language which is only moderately rich, we are able to use a word to aid or assist us in individuating, as when we say "this", and so supplement the act of pointing. With a somewhat richer language we can use uttering words to replace ostensive acts where ostension could have served as well - as when we say "The tall man in the corner ...". Or finally,

we can use uttering words to do the same job in situations where ostension is no longer possible - George is not in sight, he is out of town, he is dead, or he has been dead for centuries. But in all these cases, what is done is, in one respect, the same; we have subordinated acts to managing the achievement of individuating George. In the more 'primitive' cases, the achievement of individuating George consists in getting one's listener to see George.

The speaker's success in doing that depends upon the acts and resources of his listener (i.e., the listener's following the line of the pointing finger, turning his head, focussing, and so on). In the more 'sophisticated' cases, the achievement of individuating George, who is perhaps centuries dead, also depends upon the acts and achievements of the listener - in those cases his listening and attending and, most importantly, his recognizing the word uttered and recalling the conventional use of the word. In short, in these more 'sophisticated' cases, the speaker's success in getting the listener to know whom he is speaking about depends on the listener's linguistic competence. These cases, where we are far removed from the baptismal font, may be thought to be especially problematic for my account. I have said that an expression is a name of someone in a language when the use of that expression is subordinated, by the conventions of the language, to individuating that person. The use of the expression is sufficient by itself for individuating that person. But

when the person who is referred to is long dead, and not a possible object of ostension or perception, then the speaker's success in individuating his referent may seem impossible to understand. It may appear especially difficult to see how this could be done on my account since I have characterized individuation in such a way that it must be possible to individuate a referent for a listener even though nothing whatsoever is known about the referent. If it is not possible to individuate a long dead person, then my account of the formal conditions for the use of a name must be acknowledged to be defective, since it would have the consequence that we could not use any expression as a proper name for a long dead person. When we are in front of a man we can understand more readily how a speaker's acts, whether ostensive or linguistic, can individuate that man for us. But how the achievement of individuation is possible in other circumstances, and what that achievement could be like, is not clear. If I say that a listener's linguistic competence enables us to do that, then it seems as if all that is important to a theory of names gets buried in that remark. How, after all, could we ever come to have competence in the use of names for long dead persons?

By answering this question it may be thought we will come to an understanding of how names, or our use of them, are related to the people themselves. Then and only then will we understand the relation which obtains between a person and his name - the relation expressed by saying "X

is the name of 'Y'. But this is a thought which arises out of a confusion of the formal and material conditions for the use of a name. The problem of saying what the conditions are which enable us to use an expression as a name for someone long dead is a serious problem; but it is not a problem for an account of what it is for an expression to be a name of a person long dead.

The difficulties involved in saying what the material or 'enabling' conditions are for using a name are only problems for a theory of what the formal conditions are for the use of a name, if the latter account is such that it has as a consequence that no actual material conditions could make it possible to learn the use of a name (under that account of what a name is). In short, there must be something wrong with a theory of what names are which has the consequence that we cannot learn and use names. This is close to the kind of difficulty which Bertrand Russell's theory of names confronts in an unsatisfying way. On his account of what a proper name is, it would not be possible to learn and use a proper name for a person. The conclusion which Russell felt compelled to draw is that what we ordinarily call a proper name is not, in fact, one, but only an abbreviated description. He was led to this conclusion by what are effectively the problems that arise around names for long dead persons, but his conclusion covered all persons, including the ones we are in the presence of, since Russell's theory of perception has the result that

our friends and acquaintances are as far from sight and ostension as Aristotle is for us. Some such avenue as this - simply denying that a familiar class of expressions, pre-philosophically taken to be names, are in fact names - is open to any theory of what the use of a name is, but I will not take one. Instead, I will offer some remarks about the material conditions for the use of names of long dead persons which, I think, should remove any doubts about the possibility of learning and using such names on the account of the use of an expression as a name given above.

I have said that to use an expression as a name is to utter the expression in subordination to individuating something or someone, where the conventions of the language are such that the use of that expression is sufficient for managing that achievement. As with any convention for the use of a word, two conditions must be met for a speaker to learn the conventional use for the word. First, it must be possible to individuate for the listener the person whom the convention will enable the speaker to speak about; second, the speaker must then provide the linguistic novice with the expression which it is conventional to use for speaking about that person. The training which is provided consists in teaching the novice to subordinate uttering the expression to, in the case of names, individuating the person. On the account of names suggested, the novice will have acquired linguistic competence in the use of a name when he has acquired the skill to utter the name in subor-

dination to such an achievement. But the problem arises here - the student's acquiring that skill depends upon our first individuating the person for him. When the person whose name it is stands before us, we can do that by pointing and saying "that man is called Aristotle"; but our situation with regards to Aristotle is not so fortunate. Most (but not all) common nouns cannot present the same kind of difficulty since, although we may be far removed from the instances of a kind of thing with respect to which the nouns were instituted into the language, we can often re-enact the baptismal ceremonies with other instances of the same kind. A faithful re-enactment will maintain continuity with the history of the language and with the further reaches of the present community, and an unfaithful one will, at worst, result in an idiolect. But in teaching our students the name "Aristotle" we have no such recourse to further instances. In general, a convention for the use of a name, for using "X" to refer to X rather than to some other individual can arise only if X, rather than some other individual, is the item individuated by the speakers who perform the baptismal, stipulative acts, or who re-enacted those acts in using and teaching the convention. Without recourse to ostension and perception, how do we individuate Aristotle, and how do we do so in such a way as to prevent us from all speaking in differing idiolects when we use his name?

It is worth noting that on one view about the nature of names, the view which is held by Frege and John Searle, in which names have a meaning and that meaning is given by a description or belief (or a cluster of such), the possibility of using the same name with very different meanings in the same linguistic community is a very real possibility. Also, under any theory of names in which the referent of a name is fixed by means of descriptions or beliefs, and which allows the plausible assumption that speakers can hold widely divergent beliefs or provide widely divergent descriptions, a consequence which can ensue is that different speakers may be using the same name to refer to very different people. In either case, whether the meaning or the referent of a name is thought to be fixed by speakers' beliefs, or by descriptions they might provide, divergence of meaning or referents threatens. But, for names of long dead persons, there seems to be no way to individuate for another speaker the person for whom we want to teach a name without running up against the threat; we tell our stories about the person and introduce the name for speaking about the person. I want to accept the premise from which this difficulty is thought to follow, but suggest that if the distinction between material and formal conditions for using a name is kept in sight, the difficulties do not in fact arise.

The premise which I accept is, then, that the only way we have of individuating some persons for introducing or

teaching their name is by providing descriptions or expressing our beliefs about them. As a first step, notice that it is possible to have beliefs about some person without knowing or believing that that is who your beliefs are about. Thus, I may acquire many beliefs about Aristotle without believing that they are beliefs about Aristotle rather than some other person. I may meet a man at a party and after a conversation with him come to hold many beliefs about him but not realize that the man is Aristotle. I think all along that I am speaking with Plato. It is also possible to acquire many beliefs about Aristotle, long after he is dead, without knowing or believing that it is Aristotle about whom I have these beliefs. For example, due to a printer's error, my copy of his works contain no author's name; on reading his works I come to have many beliefs about the author - about his views and about his character - but have no idea who the author of these books is (I may even hold the false belief that the author is not Aristotle, since I have heard the name mentioned, and believe Aristotle to be a great philosopher who wrote dialogues).

Having such beliefs about Aristotle is sufficient for being in a position to now learn the use of the name "Aristotle" - I may be told that the author was called "Aristotle", I may be taught the name "Aristotle" for saying whom I hold certain beliefs about. In general, it is possible to individuate Aristotle for a person who holds

beliefs about Aristotle by saying something like "That person about whom you believe that ...", without that person knowing or believing that the person you individuate for him is Aristotle.

Furthermore, it is possible to individuate Aristotle for a listener who holds all false beliefs about Aristotle, so long as it is still Aristotle about whom he holds those beliefs. I may acquire only false beliefs about the man I spoke with at the party (I may even be wrong in thinking that the occasion was a party, and wrong in thinking that I spoke with the man - perhaps it was at a demonstration and I only overheard a conversation in which he was involved). Likewise, in the case where we are further removed in time from Aristotle, I may so badly misunderstand the difficult texts of Aristotle that I form all false opinions about his theories and his character. I think that the author of these books is constructing an elaborate defense of the doctrine of transcendent forms and, judging from his style, that he must have been a disciple of Plotinus. Circumstances may conspire so that the entire set of beliefs which a person holds about Aristotle may be false, and that person not know or believe that his beliefs are about Aristotle. Nonetheless, it would still be Aristotle about whom he holds these false beliefs. It remains possible to individuate Aristotle for such a person by saying something like "That man about whom you believe that ...".

Finally, it would be possible for everyone alive to hold only false beliefs about Aristotle and at the same time for no one to know that it is Aristotle about whom they hold those beliefs. In such a situation it would be possible for one speaker to individuate Aristotle for his listeners and say "Let's call by the name 'Aristotle' the person who defends reincarnation and the doctrine of transcendent forms, etc. . . .". So long as the person about whom they hold these beliefs is Aristotle, then Aristotle has been individuated and a name taught for saying who he is. Elaborate examples like this could be constructed which are consistent and coherent simply because who our beliefs are about is indifferent to whether those beliefs are true or false. Nor does it matter if some other person happens to fit our description perfectly; it would not be a description of that other person. We cannot be led to the conclusion that, in the very badly confused community I have described, the newly introduced name "Aristotle" is the name of someone else who happened to fit the description which the community gives. Drawing that conclusion requires the very implausible assumption that our descriptions are always either true of someone, or are not descriptions of anyone at all. Avoiding that conclusion, which threatens divergence of referents or meanings, requires only the more plausible assumption that it is possible to mis-describe something - if something is misdescribed, then it does not 'fit' the description but it remains, nonetheless,

what has been described. It cannot be a requirement of an act of speech counting as an act of describing something that the description given be true; describing is, in short, an act or activity, not an achievement.

I have defended by examples, saying that one can have beliefs about someone, say Aristotle, without knowing that the person about whom one has those beliefs is Aristotle. This provides no obstacle to saying that Aristotle is the person about whom one has those beliefs. This also remains true independently of any assignment of truth-values to those beliefs. Because of this, I have been led to say that even if an entire community has only false beliefs about Aristotle, and no one knows that the person about whom the community has those beliefs is Aristotle, it is still possible to individuate Aristotle in that community. He can then be given a name or, if he has one, a name for him can be taught. In this, the most pessimistic of scenarios, we may individuate him by saying something like "That person about whom everyone believes that ...". I concede that in such a case no one will know or believe that the person about whom we are speaking is Aristotle, but maintain that knowing or believing that Aristotle is the person whom we are talking about cannot be necessary for individuating him, or for talking about him. Requiring that would make all talk about Aristotle impossible. The achievement of individuation is, we saw earlier, only one of the components of the complex achievement of bringing a

listener to know what we have said. The subordinate achievement of individuation is, logically, prior to saying something (in the fact sense).

The worry that, if we have enough false beliefs about Aristotle, we will end up talking about someone else depends upon confusing the material and formal conditions for the use of his name as a name of him. If the relation expressed by "X is the name of Y" was reducible to, or consisted of, a chain of true beliefs leading back to, for example, Aristotle, then the possibility of using "Aristotle" as a name of Aristotle would be threatened by anything which threatened the maintenance of that chain. But in saying that that chain could only be a material condition for the use of "Aristotle" as a name of Aristotle, I have denied that the relationship between the name and the man, expressed by saying "'Aristotle' is the name of Aristotle", can be reduced to, or consist of, any chain leading back to Aristotle. The causal chain which does lead back to the man provides the material conditions which make it possible for us to individuate the man by uttering his name. But that chain does not consist of, or reduce to, a set of statements, descriptions or beliefs which are held by us, were held by his contemporaries, and were held by intervening generations of speakers. Truth and falsity are features of our descriptions and our beliefs, but they are not features of the subordinate acts of referring or of the subordinate achievements of individuation. But, since

maintaining a chain of achievements of individuation which extends from us back to Aristotle, is sufficient for us to be able to individuate the man by uttering his name, then our success in that enterprise remains unaffected by the falsity of the things we go on to say about the man after we have successfully individuated him. The ancient speaker at the baptismal fount who introduced the name "Aristotle" and taught others to utter it in subordination to saying who this person is, had to first succeed in individuating Aristotle. The first link in the causal chain which makes it possible for us to speak about Aristotle was put in place then. That chain was then continued by each speaker who also succeeded in individuating Aristotle, regardless of what those speakers then went on to say about him. We speak in such a way as to add an additional link to that same chain each time we individuate Aristotle, regardless of what we go on to say about him. That subordinate achievement of individuation is a very minimal achievement which depends upon preserving no information at all about Aristotle. If we have succeeded in preserving some information about Aristotle through the centuries, that is an achievement independent of our having maintained the ability to use his name to talk about him.

Maintaining the ability to use Aristotle's name to talk about him requires, as I have characterized it, maintaining two very different kinds of links between the generations of speakers. One of the links at each step along the way

consists in acts performed - namely, the act of uttering his name. There is little dispute about how this chain is maintained - what ties each generation to the preceding generation of speakers is the act of 'following', which requires only the ability to imitate one's linguistic elders in uttering the same name which they utter. But the convention of uttering that name in subordination to individuating Aristotle requires, as well, that each generation succeed in individuating for the next man himself, thus making it possible for one generation to teach the next to subordinate uttering that name to managing the achievement of individuating that man. It is this second link in the chain, the achievement of individuation, which is a matter of philosophical contention. I have argued that maintaining this part of the chain does not depend upon maintaining beliefs or telling stories which are true or even mostly true. It is sufficient that the beliefs we have and the stories we tell actually be about the man we mean to individuate for our listeners. But how, then, do we manage to actually keep our stories about someone who is long removed from the possibility of ostension? I will illustrate this by an example.

Suppose that you overhear a conversation about a man in which the man, who is not known to you, is referred to a number of times as "George". You may enter the conversation yourself by saying "George sounds like a terrible person to me". In fact, the gossip which you have over-

heard about George is malicious, and he did none of the things attributed to him. Who were you speaking about when you uttered the sentence? You were speaking about that person, whoever he is, whom they were gossiping about. If asked by another late arrival who you were talking about, you will only be able to reply "The person they were talking about - the person who did all those terrible things". The person they were talking about was George, but he did not do all those terrible things. But if they were talking about George, and you were talking about the person who they were talking about, then you were talking about George. Your success in individuating George was parasitic on the original speaker's having individuated George. The second late arrival may then go on to tell others about George, and his individuating George will depend upon your having individuated George. When asked who George is he will only be able to reply, as you have replied to his inquiry, by saying "The person about whom so-and-so was speaking when he said ...". His use of the name "George" to individuate George is parasitic on your story about George, but not, again, on the truth of your story. It does not matter whether the original conversants, overheard by you, were acquainted with George, or if their gossip was itself parasitic on an earlier conversation. So they may reply to your inquiries about who George really is by saying "George is the person about whom so-and-so were speaking earlier". But if the earlier speaker was talking about George, and

they were talking about the person that speaker was talking about, then they were talking about George. And if they were, then so were you; and if you were, then so was the late arrival who was talking about the person you were talking about, and so on. Later embellishments which further sully George's reputation will not alter the fact that it is still George about whom they are gossiping. Although this series of links will carry with it no interesting information about George, the series of links having the form "'George' is the name of that person about whom so-and-so said that ..." will lead back to the stories which George's own acquaintances tell about him. So long as the name is also passed from link to link, then any speaker along the way may use the name "George" as the name of George himself. At one end of this chain of transmission there will be speakers who individuate George and pass on his name by pointing and saying things like "That man in the corner is called 'George'", and at the further end there will be people who can reach agreement with each other, about George, by saying things like "George is, after all, a terrible person". The later speakers, perhaps generations removed from George himself, will be able to utter the name 'George' in subordination to individuating George, and succeed in individuating him for their listeners, because when their listeners hear the name "George" uttered, they will understand that the speaker is talking about that person about whom someone else has said some-

thing, and that that someone else was speaking about that person about whom yet another person has said something, and so on. The expression "that person" in the previous sentence is open to substitution; they are all speaking about George, and if we only had him before us, we could coordinate the utterance of "that person" with an act of pointing at George. When we utter the name 'Aristotle' what we do is, to use a metaphor, point to another speaker's act of pointing to another speaker's act of pointing, and so on, until someone is pointing at someone who is pointing at Aristotle himself. What these speakers say about what they are pointing to is not important to determining what we will find after following the line of pointing arms back through the generations. The metaphor's most misleading aspect, however, is that the generations are not linked by a series of acts, but by a series of achievements - it is the achievement of individuation that has been managed again with each generation, where the success of each new generation is parasitic on the success of each preceding generation.

This chain of achievements of individuation runs from speakers of Aristotle's generation to speakers of our generation. But it does not run back to Aristotle - he is individuated at the beginning, but he is also individuated at every step along the way. And, as I have said earlier, the chain of acts which runs parallel to this chain of achievements if his name is passed along the generations -

namely, that chain which I have earlier called the relation of 'following' - does not reach back to Aristotle either. If both these 'chains' are well maintained, then later speakers may follow the convention of earlier speakers and subordinate uttering the name "Aristotle" to managing the achievement of individuating Aristotle. That we can do this is what is expressed by saying that "Aristotle" is the name of Aristotle.

The long chains of transmission enable us to use the name to individuate the man; that is, maintaining this relationship between us and earlier generations of speakers enables us to subordinate uttering the man's name to individuate the man himself. The historical chains thus make possible the relation expressed by "X is the name of Y", but that chain cannot be cited in providing an explication of the latter relation. It would be a mistake to suppose that the relationship between our utterance of the name and our achievement of individuating the man consists in, is reducible to, or analyzable in terms of the chains of transmission. They merely make possible reaching that achievement by performing the act of uttering. But this is to say, as I stated earlier in this section, that the relationship expressed by saying "X is the name of Y", the so-called 'naming relation', is not a relation which extends from a name back to the person or thing named. The material conditions for the use of a name (the chains of transmission) make it possible for the formal conditions to be

satisfied (for the act of utterance to reach the achievement of individuation). For an expression "X" to be a name of X in a language is for the utterance of the expression to be, in virtue of the conventions of the language, sufficient for individuating X for speakers of the language.

What is it for one's statements or thoughts to be about a particular in such a way that they are distinguished from other statements or thoughts about other particulars, but not distinguished from other statements or thoughts about the same particular? My answer has been that a statement or thought is about a particular just in case the statement (or thought) has, as a subordinate achievement, the achievement of individuating that particular. The achievement of individuation is managed by speakers (or thinkers) in a variety of ways - as we have seen, a reporter is under no obligation to use the same subordinate acts as the original speaker (or thinker) to individuate the same particular for his listeners. The achievement of individuation is one which can be reached whether or not a speaker, or his listener, has any beliefs, true or false, about the particular individuated. Successful individuation of the particular is, in fact, a presupposition of having any such beliefs. In certain central cases - in the language learning situation, for example - the acts subordinated to individuation will be ostensive and the achievement itself will be perceptual. In other cases, remote from either learning or baptismal contexts, the achievement may be managed by a

more complex activity such as telling a story or giving a description. Success, in such cases, may rely heavily on the background beliefs which the listener already has. But in other cases, also remote from the possibility of ostensive acts and perceptual achievements, we may manage the achievement of individuation without relying on a background of beliefs which the listener has or shares with the speaker - we may manage it in times of massive disagreement about the particular which the speaker individuates for his listeners. The resources which our language provides us with for this achievement are names.

The account of saying something and the related account of believing something, provided in this chapter, requires that a speaker individuate an item for speaking about or for acquiring a belief about. Although the discussion of this section has focussed on proper names, it can easily be seen how it would be extended to common names and to other expressions which might be uttered or recalled in subordination to individuating a 'topic' for 'commenting on'. My dominant concern is to better understand the role of perception in coming to have beliefs or to make claims which are justified, and so our attention will not, in the following chapters, be focussed on beliefs acquired or claims made so very far from the baptismal font. What should be clear, however, is that the account of beliefs provided in this chapter does not suffer from any special inability to accommodate beliefs acquired far from perceptual circum-

stances. More importantly, I have tried to make a case for saying that, in spite of the emphasis the account places on relationships which are internal to the linguistic community, our talk and our thought is not thereby 'severed' from the real world. The account of language learning and use which I have defended does not have the anti-realist or skeptical implications which might be expressed by saying that, on this account, our speech is not related to the world, but only to other speech. I have not denied that there are relationships between language and the world - some words are names of things, statements are about things, and conversations are about things. But such relations are not, on the account, primitive; they are themselves to be understood in terms of relations which are not between language and the world, but between a speaker's acts of utterance and his own achievements, and between the acts and achievements of one speaker and those of other speakers. In general, the intentional character of thought can be understood through the intentional character of language, and the intentional character of language can be explicated in terms of the more generic notions of 'following', and 'subordination' (of acts to achievements).

Although I have emphasized, as a paradigmatic case of the phenomena of intentionality of language, the relation 'X is the name of Y', this is not meant to suggest that speaking or thinking about some extralinguistic entity depends upon the agent's use of a proper name. Doing that

depends, rather, upon subordinate achievements of individuation. Demonstratives are also uttered in subordination to individuation, but they depend for reaching that achievement on ostensive acts and perceptual achievements as well. Definite descriptions may also be subordinated to the achievement of individuation; but their success in managing that achievement depends upon the listener's having or acquiring beliefs about the item individuated. But all this is just to say that names, demonstratives and definite descriptions may all be used in subordination to individuating an item - to use any one of them may be to perform an act of referring. Attributing a belief or statement about something to someone depends upon managing, with any one of a wide variety of linguistic acts, the achievement of individuating the same item which they have individuated. Central to the consideration of the problems which follow is the central nonlinguistic achievement which makes it possible for someone to have a belief about something - the individuating achievement of seeing something. This is the topic of the next chapter.

By attending more closely to the activities and objectives of an agent speaking a language, it is hoped that the shift of attention is a shift away from the view of language and thought as products - artifacts constructed to picture or 'represent' the world.

Notes for Chapter Three

1. Descartes, Objections and Replies, from Haldane and Rodd (eds.), The Philosophical Works of Descartes, Cambridge University Press, Cambridge, 1911, Vol. II, p. 66.
2. Malcolm, Norman. "Thoughtless Brutes", Presidential Address before the American Philosophical Association, 1972.
3. A verb "V" is intensional if and only if the use of that verb to report or attribute some act or state of (or to) an agent, P, provides us with an utterance (the report or attribution) of the form "P V-ed ...", in which the "..." is not open to the free substitution of coreferential or coextensional expressions, or has parts (an imbedded subject or predicate) which are not open to the free substitution of coreferential or coextensional expressions. Such contexts are not open to free substitution if and only if substituting a coreferential or coextensional expression for some expression in that context does not logically guarantee the preservation, in the resulting report, of the truth-value of the original. I will sometimes speak of the mental act or state as if it were itself intensional, but strictly speaking, it can only be reports or attributions of states and acts which are intensional. The actual act or state is only derivatively so-called.

4. Davidson, Donald, "Thought and Talk", in Samuel Guttenplan, Mind and Language, Clarendon Press, Oxford Press, Oxford, 1975.
5. This is not, in the long run, the most perspicuous way of putting the distinction; it has all the unclarity of the subject/predicate distinction itself. The examples I will use throughout are all such that a that-cause will follow the intensional verb, and a simple subject/predicate sentence will follow the that-clause, so no problems which arise from the subject/predicate distinction should infect these cases. There is a more basic distinction discussed by some linguists (e.g., Hockett, C.F., A Course in Modern Linguistics, MacMillan, New York, 1958) between the topic of the utterance and the comment of the utterance; this distinction should be seen to underwrite my use of "subject" and "predicate". The topic (subject) is simply what the utterance is about (e.g., a ball); the comment (predicate) is what is said about the ball (e.g., that it is red). Connecting the subject/predicate distinction with the topic/comment distinction pushes the former distinction away from grammar and towards logic, since which grammatical part of a simple declarative sentence indicates the topic and which indicates the comment is largely a function of the conversational context in which the sentence is uttered. Thus, if we are discussing what has happened to Mary

today, and I say "John kicked Mary", the topic is Mary, and that John kicked her is what I said about her. In the examples we will use, it can be assumed that the topic of the imbedded sentence coincides with its grammatical subject, and the comment coincides with the grammatical predicate. Hence, the distinction that I am making here between two sorts of failures of substitution should not be viewed as a grammatical distinction.

6. I suggest that the management of both achievements is sufficient for getting the listener to know what I said, where "what I said" is taken in the sense in which the same thing can be said in different languages. But it is to no advantage to attempt to argue that here.

7. I put the point this way as a concession to traditional discussion of substitution which views the principle that substitution preserves truth as a variant of the principle of the indiscernibility of identicals. Such discussion is carried on without any view to the possibility of one's report ever being understood by one's listeners. Since our concern is with the subordinate communicative achievement of individuating for our listener some individual, the requirement of coreferentiality of subject expressions for substitution need not be made. We may, as we have already seen, substitute "Y" for "X", even though "Y=X" is false, and still

manage that achievement. The requirement of coreferentiality is made with an eye towards the semantical consideration of truth, or correctness of reports, and without keeping in mind the pragmatic conditions that govern successfully reporting to another person.

Strictly speaking, coreferentiality ought to yield to a pragmatic requirement which allows substitution of "Y" for "X" only if "Y" picks out for the listener the same individual that was originally picked out by "X".

8. Donnellan, Keith, "Substitution and Reference", Journal of Philosophy, Vol. LXIII, No. 21, November 10, 1966, p. 687.

See also Donnellan's "Reference and Definite Descriptions", The Philosophical Review, 75, 1966, pp. 281-304; and his "Proper Names and Identifying Descriptions", Semantics of Natural Languages, Davidson and Harman (eds.), pp. 356-379.

9. Quine, W.V.O., "Reference and Modality", in From a Logical Point of View, Harvard University Press, 1953, p. 155.

10. Quine, W.V.O., "Reply to Professor Marcus", in The Ways of Paradox, Random House, New York, 1966, p. 182.

11. The two most notable examples of this are Quine, in "Quantifiers and Propositional Attitudes", from The Ways of Paradox, and Carnap in Meaning and Necessity.

12. The position explicated on the following pages is from Frege's "On Sense and Reference", in Translations From the Philosophical Writings of Gottlob Frege, P. Geach and M. Black (eds), Oxford, 1960.
13. Frege, G., "The Thought: A Logical Inquiry", from Philosophical Logic, P.F. Strawson (ed.), Oxford University Press, 1967.
14. See Leonard Linsky, "Reference, Essentialism, and Modality", in Reference and Modality, L. Linsky (ed.), Oxford University Press, 1971, pp. 91-92.
15. Kaplan, "Quantifying In", in Linsky, ibid, pp. 112-144.
16. Church, A., "On Carnap's Analysis of Statements of Assertion and Belief", in Linsky, ibid, pp. 168-170.
17. ibid, p. 169.
18. Vendler, Z. Res Cogitans: An Essay in Rational Psychology, Cornell University Press, 1972, p. 41.
19. "Toward a Causal Theory of Linguistic Representation", in Contemporary Perspectives on the Philosophy of Language, Peter A. French, et. al. (Eds.), University of Minnesota Press, Minneapolis, 1979, p. 81.

CHAPTER FOUR

PERCEPTION

A. Perception and Propositional States

The verb "to see" has a plethora of uses, most of which are of no interest to us in our inquiry. Our investigation centres on the perceptual use of the verb, the use which indicates a specific achievement the management of which presupposes the use of the eyes. Even though I might only see your point of view after having read your letter, I could have seen the same point, in exactly the same sense of "see", if you had explained your position over the telephone; but I could not have seen my brother as he disembarked from the airplane unless I had used my eyes, had been facing in the right direction, and the fog had dispersed earlier in the day.

The results of the preceding chapter have direct consequences upon an analysis of perception. The assumption of the conventionality of language led us to say that a necessary condition for setting a precedent for the use of a proper or common name, or for instructing an initiate to the language in the use of such expressions, is that it be possible to make clear to the initiate with respect to which individual (whether the subject of a proper name or a paradigm for a class of individuals) a precedent is being

set or a convention is being taught. The resources for managing this achievement were held to be, initially at least, ostensive acts performed by the word-smith or the language teacher. A second condition for establishing and maintaining conventions for speech is that an expression be uttered (or written). It should be clear at once that in supposing that such procedures are possible, we must also presuppose that an initiate into a language have and make use of certain perceptual faculties - he must, for example, see the object indicated by the language teacher, and he must hear the word (or sound, not to beg certain issues to be dealt with later) uttered if he is to learn the language in anything like the normal way.

The results of the last chapter, together with the presupposition just uncovered, entail that there must be some forms of perception which do not presuppose or involve the having or coming to have of propositional states. Having beliefs and knowledge presupposes having acquired a language and having exercised skills that one acquires in learning a language, and acquiring a language presupposes having managed earlier achievements of perception. In short, it must be possible for me to see things, and to hear things, before it is possible for me to come to know or believe anything at all. It cannot be a necessary condition for perceiving X that I have some belief or knowledge as to what X is, nor can it be true that upon

perceiving X I necessarily come to have some belief or knowledge as to what X is.

This argument seems to me persuasive, but because it is so laden with the account already constructed, it is worthwhile to cite some independent considerations which support the same conclusion. First, however, we should be clear as to the nature of the minimal claim that is being rejected. The claim that I am denying is that "P sees X" entails "P knows (or believes) that X is F", for any F. It may well be true that a person with a relatively sophisticated conceptual framework rarely sees something and fails to come to have some belief about that which he sees. But this fact, if it is one, does not make it a necessary condition for seeing something that one come to have or already have some belief. Even if it were true that every time we see something we come to have some belief or another, that would offer no consolation to the view I mean to reject - it would not help us to know whether part of what it means to see X is that we come to have some belief about X, or whether it just turns out that every time we see something we also come to have a belief about it. But even this weaker claim, that there is a contingent connection between every perceptual achievement and some belief or another, seems highly implausible. Watching a friend walk, lost in thought, along a sidewalk, we may notice that he ducks his head slightly, or steps a little to the side, every time that he approaches an overhanging branch. If we

later ask him what belief he has about any one of those branches, a perfectly sincere response might be "What branch?". He may sincerely insist that he neither saw, nor remembers seeing any branches at all. And a perfectly plausible rejoinder for us to make would be "But you must have seen it, since you obviously walked in such a way as to avoid hitting it", and this rejoinder need not depend for its meaningfulness on the assumption that though he acquired some belief, he has since lost it.

The further implausibility of the general claim I am taking issue with, as well as the weaker contingent version of the claim, is brought out by examples like the following. Suppose that I work in a ball-bearing factory, and although my job is not to examine or inspect the ball-bearings, I am nonetheless stationed in such a position that I see, as they pass before me, literally thousands of different ball-bearings every day. It seems wildly implausible to suggest that I acquire literally thousands of new beliefs every day. Part of what is so dissatisfactory about the thesis that seeing entails believing is that no particular belief seems a likely candidate for the belief that I must come to have when I perceive any particular object. Nor is it clear whether I must come to have a true belief, or if any old false one will do as well. The requirement that I come to have some true belief seems too strong, even if the belief is so generally stated as to appear almost empty.

Fred Dretske, in Seeing and Knowing, cites an experiment in

which "subjects were shown various colored shapes on the wall, but were led to believe that the shapes they were seeing were actually their own imaginative constructions".¹ If the subjects were actually led to believe that they were not seeing anything at all, when in fact they were seeing various things, then there is no reason to suppose that they would have acquired any true beliefs about the things that they saw - not even the minimal belief that they saw something. As Dretske writes, seeing something is consistent with "any false belief one may care to mention about the generic character of what one is seeing. And since it seems obvious that it is also consistent with any true belief,² it is logically independent of such beliefs". In the end, I would prefer to rest my defense of the claim that seeing does not entail either believing or knowing on the general argument with which I opened this section; but whether or not that argument will do (and the thesis I am combating is not as strong as to count as a reduction of the account of language and belief) will depend upon whether or not we are successful in giving an account of seeing which does not itself presuppose the having of beliefs or knowledge. Giving just such an account is my aim in the remainder of this chapter, and as a first step towards doing that it will be useful to say something about the general nature of the possible objects of perception, that is, about the sorts of things that it is possible for us to perceive.

B. The Possible Objects of Perception: Publicity

One assumption in my initial argument for a nonepistemic account of perception might appear objectionable to some philosophers. I have told the story of language acquisition in such a way as to assume that the objects of perception are public objects - that one and the same thing is seen by different persons, and one and the same utterance is heard by different persons. There are, however, many philosophers who would prefer to say that the objects of perception are private - that only I can see that which I see, and only I can hear that which I hear. We must distinguish a weaker and a stronger thesis. The strong thesis, that all the objects of perception are private, will be rejected on two grounds: First, that the claim conflicts with the assumption of the conventionality of language; second, that the arguments which lead philosophers to hold this strong position seem to be fallacious. In addition, I am persuaded that a weaker claim, that some objects of perception are private, is also false. Contemporary literature has, I believe, contained ample persuasive arguments against both the weaker and the stronger theses, so much of what I have to say will be little more than a brief rehearsal of those arguments. I will treat the weaker thesis in much less detail for two reasons. First, whether or not there are some unusual cases in which it is best to say that we perceive private objects is of much

less interest for the epistemological thesis being developed. Secondly, adequate treatment of the cases of hallucinations, afterimages, sensations and the like would require an extended excursion into the philosophy of mind, leading us too far astray from our present objectives.

To begin with, we should note the conflict between the claim that all objects of perception are private and the assumption that we share a conventional language.³ A conventional language is, by its very nature, a shared and public language. The conventionality of language requires the possibility (or better, the actuality) of setting precedents for speech, of initiating others into the practice of following those precedents, and the possibility of transmitting and maintaining the conventions in a linguistic community. Let us assume the supposition that everything we perceive is private to each perceiver, and then try to say what goes on in a typical situation in which one person instructs another in the use of an expression. The language teacher perceives something, S, which is such that on many past occasions when he has seen things of this sort, he has eaten and derived both pleasure and nutrition. Meaning to create the linguistic resources for instructing his subordinates to go out in search of more of those things he coins a word, W, and proceeds to institute the use of that word in his language. The things he has eaten and derived nutrition from are not identical with anything that he has ever seen, since the things that he sees are all private,

and so not nutritious. There is no reason to suppose that this person holds a theory of perception like the one we are criticizing, and so no reason to suppose that he does not make the mistake of thinking that what he sees and what he eats are one and the same thing. If he does not already hold such a theory of perception, then he means to introduce an expression into the language for saying of that thing which he sees 'before him' what it is. Without knowing or thinking that what he sees is a private, perhaps mental state of himself, he means to introduce a word into the language for talking about it - the private object or mental state. But he happens to identify (mistakenly, on the view under consideration) that mental state with whatever it is that gives him pleasure and nutrition. So perhaps he means to introduce a word into the language for talking about that which he eats. One thing that does not occur to him, however, is that he ought to introduce two new words into the language - one for what he sees, and one for what he eats. It is fortunate that this does not occur to him, since he could not succeed in introducing a word into the language for what he sees, since no one else could see it, and there would be no way for others to learn the use of that expression, separate from the expression for what it is that he eats. But neither he, nor anyone else, can see what he eats.

Consider what happens when our language teacher points to something and utters a word for saying what something

is. If we understand pointing at something as stretching out one's arm and index finger and so controlling their position that they lie roughly on a line which extends from the agent to the object, then it is clear that he cannot point to what he sees, but only at that other thing which he eats. The activity of pointing is presumably subordinated in this case to making clear to others what it is that he then goes on to name. But of course he cannot manage that achievement, whether he intends to name what he sees, or what he eats, since no one else can see either of those things. Nor can they see his arm and hand move into position for pointing at it. We may assume that there is some significant relationship between what each of the others see, what he sees, and what goes on in the world beyond their 'range' of sight. A common originator of their private experiences, whether a beneficent God or a beneficent causal order, might guarantee enough continuity within each person's experiences to make it look to each person as though he was speaking to every other person about the same things. But this could only be an illusion. Since, in point of fact, they have all made the mistake of thinking that what were being named, when they learned a language, were the things that they see, and those are the things that they go on talking about when they use that language. P establishes a precedent (unwittingly) for talking about his own private states, and P 'follows' that precedent in talking about his own private states. If he

were to follow the precedent correctly, he would do so by talking about some state of P¹'s when he uses the expression W. But since P¹ could never have shown P² that private state, P² could not have learned to follow the precedent correctly. Although this leaves them with the illusion of speaking a conventional language, it is only an illusion, and there would, in fact, be no conventional language in operation. No one can correctly follow the precedents set by others, and no one can understand the utterances of others. But we have been guided all along by the assumption that there is a conventional language, not merely that there just appears to be one. Anyone who objects to this assumption on the grounds that we could not tell for sure that there is one, and does not just appear to be a conventional language, because of the sorts of considerations raised by the privacy of all objects of perception, thereby concedes that he does not understand us when we speak (but only appears to) and that we do not understand him when he speaks (but only appear to). Such a discussion need not be continued.

Our next step will be to rehearse, briefly, two of the major arguments which have led philosophers to say that all objects of perception are private. One form of argument which has led philosophers to this conclusion rests on a confusion between what we see and what we should say that we see if we mean to say something which is least susceptible to error. The driving motive for this form of

argument is the search for the incorrigible. The mistake has been to assume that because there are, epistemologically, less risky things to say in order to say what we see, then what we really see is what fits the less risky description but not what fits the more corrigible description of what we see. Since we are less likely to be shown wrong if what we say we see is a red patch, than if we say that we see a tomato, the right answer to the question of what we see is more likely to be "a red patch" than "a tomato". But from

(1) we are more likely to be correct if we say that we see a red patch than if we say that we see a tomato,

it does not follow that

(2) it is more likely to be true that what we see is a red patch, and not a tomato.

The truth of (1) is perfectly consistent with it being true that the red patch that we see is in fact a tomato.

Bertrand Russell writes:

A table viewed from one place presents a different appearance from that which it presents from another place. This is the language of common sense, but this language already assumes that there is a real table of which we see the appearances. Let us try to state what is known in terms of sensible objects alone, without any element of hypothesis.⁵

The opening remark, that a table "presents a different appearance" when viewed from different places, is innocuous

enough, but the notion of "presenting an appearance" is quickly transformed when he says that we "see the appearances" of the table. I also get a very different glimpse of something when my position is changed, and even sometimes a glimpse of different things, but there is no reason to suppose that what I see are the glimpses. I catch a glimpse of the table only if I see, momentarily, the table, and the table presents an appearance to me only if I see the table. Although Russell is not even supposing, in this passage, that it could be the table which is seen, it is clear that his aversion to saying that we see the table itself derives from his supposition that to say that is to assume that there is a real table. But that is not (he supposes) something that we could know, and so contains a large "element of hypothesis". What we say does reveal certain assumptions we make, but it reveals only what we think we see. Stating what we think we see in such a way as to remove any element of hypothesis may result in our saying less about what we see, but doing that will not result in our saying that we see less. If I see, from a distance, what I take to be a grizzly bear against the background of trees, I run a great deal more risk in saying "I see a grizzly bear" than in saying "I see a brown splotch". But from my saying the latter, and getting it right, it does not follow that what I see is a brown splotch and not a grizzly. The brown splotch which I see in the distance may well turn out to be a grizzly, in which

case, whether I say it, know it or think it, what I see is a grizzly.

There is a second fallacy which is equally widespread and which has led many philosophers to say that what we perceive is a private state of ourselves, such as sense data or sensations. The general argument runs as follows: X looks F to me; but X is not F; therefore, there must be something else which is F and that must be what it is that I really see. The argument is invalid because, in general, if X looks F to me, that is because I see X. If I did not see X, it would not look any way at all to me.

A.J. Ayer, in The Foundations of Empirical Knowledge, provides us with one instance of this fallacy. He writes:

When I look at myself in the glass my body appears to be some distance behind the glass; but other observations indicate that it is in front of it. Since it is impossible for my body to be in both places at once, these perceptions cannot all be veridical. I believe, in fact, that the ones which are delusive are those in which my body appears to be behind the glass. But can it be denied that when one looks at oneself in the glass one is seeing something? And if, in this case, there really is no such material thing as my body in the place where it appears to be, what is it that I am seeing? Once again, the answer that we are invited to give is that it is a sense datum.⁷

Two things should be noted, and set aside, at the outset.

— First, this is not an argument which, if it were successful, would show that all we ever see are sense data, since it is restricted to the special case of mirror images. It is, however, part of a sequence of arguments which is meant

to lead to that conclusion. Secondly, the case of mirrors provides additional problems which should be set aside. Ayer seems to take it for granted that the locution "I see a reflection of myself" has some primacy over the locution "I see myself reflected in the mirror". The first locution suggests that what I see, when I look in the mirror, is not myself, but something else - a reflection. The second locution suggests that I see myself, but adds something about the special means by which I see myself. Ayer seems to shift from preferring one locution to preferring the other. When he says "when one looks at oneself in the glass one is seeing something", he allows that by looking in the mirror, one can be looking at oneself. Why not, then, allow that what one sees is just what one is looking at - oneself?

The fallacy we are interested in here is this one:

- (1) My body appears to be some distance behind the glass.
- (2) My body is actually in front of the glass.
- (3) Therefore, there exists something else which I see and which is behind the glass.

The absurdity of the argument, and of the fallacy in its most general form, is unwittingly displayed by Ayer in his choice of examples. What seems seductively plausible when the argument is worked through with color or shapes is implausible with the property of spatial location. Ayer should be led to say, if he rides the argument to its

conclusion, that the sense datum which he sees is behind the glass. But it is simply bad metaphysics to let the argument go all the way to its conclusion. If it follows from the fact that I see what appears to be a dog but is not, that I see something which is called the appearance of a dog, then it ought also to follow that if I kick what appears to be a dog, but is not, then I kick the appearance of a dog. Besides being simply fallacious, the argument lends no support to saying that whatever it is that I do see is private. It can, as in Ayer's case, be three feet behind a mirror, or as in the last case it can be something which I might kick.

I have argued that there are grounds for rejecting the strong thesis that all objects of perception are private. This still leaves open the possibility that some of the things which we see are private. I would now like to offer some reasons for rejecting the weaker thesis as well. The best cases for private objects of perception seem to be afterimages, hallucinations, and sensations. I will discuss each of these cases briefly in order to show that there are significant differences between them and central cases of objects of perception - differences which do not obviously result from the fact that these are private objects or states.

First, it should be noted that an important part of our talk about these private objects is the role that the expression "has" plays in that talk. I can have a sensa-

tion, an afterimage, or an hallucination. Of course I may also have a cat, a cup, or a mountain, but not in the same sense of "have". Having a pain or a tingle entails feeling it, and vice versa, but having a coin in my pocket or a cup on my desk does not entail feeling them or perceiving them in any way, or vice versa. Second, it should be noticed that, in the case of sensations, the apparent perceptual locutions (I feel a pain in my foot) are logically equivalent to certain locutions with no suggestion of a perceptual report. "I feel a pain in my foot" says no more, and no less, than "My foot hurts". In the second locution one of the items mentioned in the first locution disappears. Instead of talking about an agent, a sensation, and a part of the agent's body, we talk only about the agent and a part of his body. No such rewrite of "I see the cup" or "I feel the coin in my pocket" is possible which eliminates mention of the cup or the coin. The implication of these ways of talking is not merely that to feel a pain is not to perceive an object, but that the pain is not in any sense an object.

If the use of "feel" in "P feels a pain in his foot" is the use of a perceptual verb, then there is no reason that we ought not to be able to have a pain without feeling it. But since having a pain in my foot entails that my foot hurts, and that entails that I feel the pain, the implication is that "feel" does not operate in this case as a perceptual verb. We have, in fact, a common nonperceptual

use of the word "feel" which has the same logical characteristics as the use of "feel" in "I feel a pain in my foot". If I feel sad it follows that I am sad, and vice versa. But to feel sad is not to perceive sadness, since I might do that (if the expression makes any sense at all) without being sad. To say that I feel sad is to say that there is some way that I feel. To say that I feel sad is to say that I am in a certain state - that is, it is to say how I feel, not what I feel. But exactly the same thing is true of feeling a pain in my foot. To say that my foot hurts, or that I feel a pain in my foot, is to say how I feel. (Or how my foot feels; sensations, unlike emotions, are not states of me. I am sad, but my foot hurts.)

The most important difference, however, between sensations and possible objects of perception, or between feeling a sensation and perceiving something, is brought out by John W. Cook in the following passage:

How, for instance, could one make out the difference between not feeling for pain and feeling for a pain but finding none? Here all talk about a kind of observation appropriate to sensations become obvious nonsense.⁸

The reason for this is clear; "feeling for a pain" is an activity, and feeling it or "finding it" is an achievement. Perceptual verbs like "see", "feel" and "hear" are, we shall soon see, achievement verbs. But if they are achievement verbs then there ought to be corresponding act verbs which culminate in those sorts of achievements -

verbs like "looking for", "touching", "listening", etc. I feel the quarter in my pocket only if, in my fumbings, I touch it. Though I may touch the wound on my foot I cannot touch the pain, and so I cannot manage the achievement of feeling it. That I feel a pain can only mean that I, or some part of my body, am in a certain state, not that I have managed some perceptual achievement. In short, pains (and other sensations) are feelings, they are not things felt on successfully feeling about for something.

It is this last feature which also points out the crucial difference between cases of perception and the having of afterimages or hallucinations. We may put the point by saying that there is no such thing as improving one's position with respect to afterimages, sensations or hallucinations. If I see a cow in a distant pasture, I may improve my position for viewing the cow by walking closer, or by waiting for the sun to rise and the mist to clear. Once my position is somewhat improved, it becomes possible for me to view the cow more clearly - I may see parts of the cow which I could not see before, or I may be able to make out the spots on its face. By performing additional acts, such as walking closer to the cow, I put myself in a position for managing perceptual achievements which at an earlier time were beyond my means. But if to have an afterimage is to see an afterimage, and if to have an hallucination of a dagger is to see an hallucinatory dagger, then I ought to be able to do something similar to improving

my position for viewing the cow. Consider the case where, after viewing briefly a yellow spot, I turn to the white wall and 'see' a blue afterimage. The afterimage is somewhat faint, and it is difficult for me to say just what color it is. Do I not improve my position, or my view of it, if I turn back to the yellow spot and look at it longer, and then look again at the wall? My afterimage is, as a result of my performing these acts, much more vivid. But doing this does not count as improving my view of the original afterimage; what I have done is produce a new and more vivid afterimage. Even if we individuate afterimages in such a way as to lead us to say that we are now having the same afterimage, I have still failed to improve my view of the afterimage, since I have not improved my position, but I have only changed the quality of the afterimage. It would be as though the cow suddenly grew much larger, so I could see it more clearly, and not as though I has walked closer to the cow, and so performed actions the result of which is that I manage new perceptual achievements. The same difference exists in the sensation and hallucination cases. I may attend very closely to the cow on the distant hillside, and not be able to make out its spots. So I may attend more or less closely to the ache in my foot, or I may even distract myself from it as well as I am able. But I cannot improve my position and so feel it better, in the way in which I can remove my gloves in order to feel the fine detail on the quarter. If, after all my attention has

been given to the pain in my foot, I can still just barely feel it, the only way that I can get the pain to feel more vivid or sharp is to do something like driving the needle further into my foot. But doing that simply alters the pain, and does not improve my perception of it. Likewise with hallucinations. If Macbeth approaches the dagger which he thinks he sees, and begins to 'make out' the engraving on the handle, or the blood on the blade, the character of his hallucination (or of the 'hallucinatory dagger') simply alters; he does not get a closer look at something.

There are further reasons to say that it is a mistake to suppose that someone who has an hallucination of a dagger sees a dagger, even though an hallucinatory one. To say that I had a near accident is only to say that I nearly had an accident, not that I had an accident, but a special kind of accident - namely, a 'near accident'. I suggest that the difference between a dagger which is real and a dagger which is hallucinatory is like the difference between real accidents and near accidents. There are not two kinds of accidents, real ones and near ones, and there are not two kinds of daggers, real ones and hallucinatory ones. The key to seeing this as a confusion is to keep sight of an ambiguity in expressions like "He falsely thinks that he sees a dagger". If I see a Bowie knife, but mistakenly take it to be a dagger, then I think that I see a dagger, but do not. The mistake, if I make it, is one of seeing

something and taking it to be something other than it is. If, on the other hand, I have an hallucination of a dagger, and mistakenly think that I really see a dagger, my mistake is not one of misidentifying what I see, even though I also hold, in this case as in the last, the mistaken belief that I see a dagger. My mistake here is to think that I see a dagger, not that I see a real dagger rather than an hallucinatory one. If we say that Macbeth actually sees an hallucinatory dagger, then we lose the distinction between mistakenly thinking that one sees a dagger, and mistakenly thinking that one sees a dagger. The former mistake is the mistake of misidentifying something, but the latter mistake is much more radical. On the view being criticized, having an hallucination is in no interesting way different from actually seeing something; it is exactly like the difference between seeing a dagger and seeing a Bowie knife.⁹

A final point about hallucinations and afterimages is that they necessarily look the way that they are. This is in effect the reason that there is no such thing as improving our position for 'viewing' them. What this should not be taken to suggest, however, is that we cannot have, say, an afterimage without thereby knowing just what color it is. There are lots of ways in which we can go wrong in saying or thinking what color something is without our having mistaken the color it looks for the color it really is. I may judge that the color of my afterimage is crim-

son, when in fact it is magenta, just as easily as I might make the same mistake about the color of the plum setting on my plate. Or, because it is such an odd color, and I have no confidence in my use of color expressions, I may abstain from forming any opinion at all as to the color of the plum or the color of the afterimage. Or finally, I may not know or have any opinions as to the color of my afterimage, or the monster in my dreams, simply because I have no color vocabulary, and hence lack the resources for making any judgements about its color. The fact there is no is/looks distinction for afterimages and hallucinations gives us no more guarantee of success than the fact that in some conditions the plum happens to look just the color it is; nor does it guarantee that we even make an attempt to judge what to say in order to say what color they are. If afterimages and hallucinations are epistemologically special, it is only in the sense that we are barred from making one kind of mistake in our judgements about them - we cannot mistake the color that they look for the color that they actually are, since they can only look the color that they are.

C. The Possible Objects of Perception: Particularity

In this section we will narrow further the class of possible objects by arguing that universals are not possible objects of perception. This conclusion will be of

much use since it is often said that the way in which we come to know, by perception, the color of some object is by seeing its color. The conclusion would be of no interest if it was held that there are no universals, but since many philosophers of perception do hold that there are universals, and that they are possible objects of perception, it will be necessary to argue the point. The general point which must be made is that we do not see the properties, qualities or features of objects, but only the objects themselves. I put the thesis by saying that we do not see universals because the features of things are generally held to be universals - that is to say, two different things can have one and the same feature - the same color, the same shape, or the same length. Since colors provide the clearest case of perceptible features of objects, I will give the arguments for the case of colors, but the same arguments ought to be easily applicable to shapes, sizes, lengths, and so on.

It should be clear at the beginning that in denying that we see colors I am not denying that we have color vision, or claiming that we are all color blind. It will be argued in a later part of this chapter that seeing something bears an important relationship to discriminatory behavior. To see that there is an important difference between seeing colors and having color vision we will need a notion of differential response or reaction, which contrasts with the notion of discrimination.

I will say that an object, X, has or undergoes a differential response, or reaction, to a variation in some range of stimulus if a variation in stimuli from within that range results in a different specific response in X. Examples of variation of stimuli from within a range would be a change in color or a change in temperature. A change in color or hue effects a change in the pattern of stimulation in the set of cones in our retina which counts as a differential response to the change in color, since there is one specific pattern of stimulation corresponding to each specific change in hue. The use of the term "differential response" should be kept sharply distinguished from most notions of discrimination. On the definition given, an ice cube has a differential response to variations in temperature, but we will not say that the ice cube discriminates different temperatures.

A key feature of the notion of a differential response is that if a creature has a differential response to something there is no suggestion that it knows, believes, or is aware (in any sense) that it has or is having such a response. It should also be kept in mind that it remains an empirical problem to say exactly what range of stimuli is responsible for a given range of differential responses.

The notion of a differential response is tied only to reactions and not to full-fledged actions. There is much disagreement as to where and how the line is to be drawn between action and reaction. The excitation of the cones

in the eye in response to incident light is a clear case of reaction and the starting of a race in response to the starting gun is a clear case of action, but there are many unclear cases. If a man jumps when a firecracker explodes behind him, has he acted or reacted? For our purposes we need not take a side on where to draw the line. The only requirement is that, whatever the correct way of marking the distinction, nothing will count as a differential reaction unless it is a reaction and not an action.

Finally, there is an obvious relationship between differential reactions and the actions which we are capable of performing. If a dog is to be trained to play dead when a certain chord is struck, to beg when another chord is played, and conditioned to salivate upon hearing yet another chord, then the dog's ears must be sensitive to the relevant differences in the sounds that he hears. That is, he must have a different reaction in his auditory hardware to different chords that he hears. If he had a uniform reaction response to different chords, then he could not learn, be trained, or conditioned to perform different actions upon hearing different chords. (That he hears different chords is perfectly compatible with them all sounding the same to him - that is, to his not having any differential response to that variation in a range to stimuli. In the same way, I can see ten very different horses, but because of my distance from them, they may all look the same to me;

no one of them has a different effect on my visual apparatus than any of the others.)

Having color vision is not having the ability to see colors, but only having visual apparatus sufficiently sensitive to differences in the colors of the things that we do see so that we have a differential response to two objects which differ in no other way but their colors. To have color vision is not to see colors, in addition to the things that are colored; it is simply for the different things that we do see to look different to us in a respect that we have come to call a difference in hue or color.

A consequence of this is that the question as to whether we have color vision is neutral with regards to the question as to whether or not we see colors. A person with no color vision at all ought not, for that reason, fail to see colors - if anyone at all can see them. Of course, neither the colors that he sees nor the objects which have the colors that he sees, would look any color to him. This immediately offers the following question. If we do see colors and we have color vision, then what among the things that we see looks colored to us? If the cup that I see before me does not look white, then what does? If the color of the cup looks white, but the cup does not, then under nonstandard viewing conditions (e.g., under a red light), it must be the color of the cup (i.e., white or whiteness) which looks pink. But if it is the cup itself which looks white, then of what use is the theory that we

see, in addition to the cup, its color? To say that we see the cup, which looks white, and we see the color of the cup seems a waste of words. Either way of speaking seems pointless or nonsensical.

There is an additional argument which convinces me that we do not see colors, and which can be used to show that we do not perceive, with any of the senses, the properties or features of things, if those features are taken to be universals. ~~The~~ argument, in its simplest form, runs as follows:

1. P sees the color of A.
2. The color of A is the same as the color of B.
3. P sees the color of B.

If P sees the color of A, and seeing the color of A is not the same as seeing A and A's looking some color to him, then in exactly the same sense in which he sees the color of A, he sees the color of B. He sees the color of B, by this argument, even though he does not see B, he need not be looking in B's direction, and even if B is buried on the far side of the moon and P is standing in his home. In fact, in exactly the same sense in which he sees the color of A, P also sees the color of everything in the universe which has the same color as A. This will be true even if P sees none of the objects which are the same in color as A. The absurdity of this conclusion is manifest.

Many philosophers would deny the second premise and hold with Locke that "all things that exist are merely particulars".¹¹ It will be argued that the color of A cannot be the same as the color of B, simple because one is the color of A and the other is the color of B; they have the important difference of 'belonging' to two distinct individuals, and in so far as they belong to distinct individuals, they are distinct colors. The consequence of this view, however, is that no individual has a color just so long as that individual is divisible into colored parts. If A is red, and A is divided into two separate parts A₁ and A₂, then A₁ and A₂ necessarily have distinct colors. But that means that A could not have been red after all, but must have been red₁ in part and red₂ in part. The process in division could go on, producing parts no two of which could be the same in color, until the original object is so finely divided that none of the parts has any color at all. The only things which could be properly said to have a color are those parts of any divisible individual which remain before the final division which produces parts which have no color. If the maximal number of colored parts into which A can be divided is N, then A must have been of N different colors before the process of division began. The oddity of this view is a result of the fact that places are used to individuate colors. It might be thought that this view could be maintained in spite of this odd consequence if we lumped together colors which are

similar, or 'exactly similar' and called them all by the same name. But such an attempt to reduce talk of universals to talk of particulars, by replacing the notion of sameness with the notion of similarity, has serious difficulties of its own. First, it is not clear that any notion of similarity can replace the notion of sameness in all the needed ways, since the notion of similarity seems to be built upon the notion of sameness - two things are similar only if they are the same in some respect. Secondly, the theory will not do for the reason that names or expressions are themselves universals - we cannot call similar things the same name, but only by similar names. But this consequence is unacceptable, since if we cannot utter the same word twice, we cannot follow precedents or conventions correctly, and so the entire structure of a conventional language could not be built. Since colors are universals, the argument with which we began this discussion seems to leave us with no choice but to say that we cannot see colors. Although this conclusion can be avoided by denying the second premise (a move which, I believe, has unacceptable metaphysical consequences) the denial of that premise forces us to treat colors as particulars, thus conceding the general conclusion I want - that universals are not possible objects of perception.

Although I have stated the argument for the case of colors, it should be noted that the same argument applies for any feature which is a universal - lengths, shapes,

sizes, and so on. Lastly, I would like to point out that although this is an argument about possible objects of perception, it also works to show that universals are not possible objects of ostensive acts. If I hold up or point to the color of A, and the color of A is the same as the color of B, then I hold up or point to the color of B, even though B itself is located in the opposite direction from the person pointing at A.

If relations are universals, so that A can stand in the same relation to B that C stands in to D, the argument already given shows that relations are not possible objects of perception or of ostension. The following dialogue illustrates an additional reason for saying that relations are not possible objects of perception.

P: Can you taste the difference between these two brands of beer?

Q: Of course.

P: Very well, taste this one first.

Q: (tasting) It tastes very sweet.

P: Now taste this one.

Q: (tasting) It doesn't taste as sweet as the first one.

P: Now taste the difference.

Obviously, there is nothing that Q can do in response to P's last instruction. He has already done all that he can do to taste the difference between the two beers, and his remarks indicate that he has tasted the difference. But

all that he has tasted is the beer. To taste the difference just is to taste the two beers and to exercise one's faculty of judgement in order to compare them. The same sort of account suggests itself for explaining the fact that we do say things like "I can see the color of the man's hat" or "I see the shape of the steeple against the sunset". All such utterances must be understood as reporting the perception of a particular and as reporting that some judgement has been made about the particular. To see the color of a star is to see the star and to judge that its color is such and such, or that it has a color. Talk about seeing colors, tasting differences, and so on must be understood as meaning that we can see what color something is, or see that something is a certain color, or see that something has a color. Seeing colors, like tasting differences, involves acts and achievements of judgement, not of perception, or not only of perception.

With this I will rest the case for saying that possible objects of perception are public and particular. It should not be thought that this is a very narrow restriction on the class of objects of perception, since it so far includes things as diverse as cups, mountains, the sky, a hole in the wall, shadows, rainbows, and mirages. Whether or not there are other considerations which will narrow the class further is something which can best be decided by pushing on to give an analysis of perception - in particular, an analysis of visual perception or seeing.

D. Causal Theories of Perception

It seems perfectly natural to begin an analysis of perception by saying that statements like "P sees X" express some relationship between P and X. To many recent philosophers of perception it seems equally natural to say that the relevant relationship between P and X is a causal relationship. What I would like to argue is that even if it is true that in most cases in which P sees X, there is some causal relationship between X and P, it cannot be a necessary condition for seeing something that one be in some causal relationship with what is seen. There are two reasons for rejecting a causal account. First, there are serious difficulties in viewing the causal theory as a possible analysis of the concept of perception, if we take the subject to be analyzed to be that concept of perception which we do, in fact, have at this time. Secondly, there are counterexamples to the claim that every case of seeing something involves some kind of causal relationship between the perceiver and what is seen.

Let us take the minimal claim that a theory must make, if it is to count as a causal theory of visual perception, to be:

P sees X only if some event involving X, or occurring at the surface of X, some property of X, or some fact about X, is causally efficacious in bringing about some event involving P, or causes P to be in some state, or causes P's visual apparatus to be in some state; and the relevant effect is brought about by the relevant cause via the specific causal mechanism M.

The second conjunct is necessary since the first conjunct alone would allow us to count getting poked in the eye by a stick on a dark night as a case of seeing the stick, providing that the poke in the eye caused us to go into the appropriate neural state. It is, however, just the last conjunct which causes difficulties for the causal theory. It is necessary to give an exact specification of the causal mechanism, M, required for a case of seeing, before we have anything which counts as an analysis of perception. There may, however, be means of filling out the necessary causal mechanism. H.P. Grice writes:

I suggest that the best procedure for the Causal Theorist is to indicate the mode of causal connection by examples; to say that, for an object to be perceived by X, it is sufficient that it should be causally involved in the generation of some sense impression by X in the kind of way in which, for example, when I look at my hand in a good light, my hand is causally responsible for its looking to me as if there were a hand before me, or in which ... (and so on), whatever that kind of way may be; and to be enlightened on that question one must have recourse to a specialist.¹³

While it is almost certainly true that in most cases there is some event involving the object that P sees which is causally efficacious in bringing about some event involving P, and if the latter event had not have occurred, P would not have seen what he did, there seems no reason to suppose that the causal mechanism can be usefully mentioned in giving an analysis of seeing. The fact that there is a specific causal mechanism in operation in all cases of

seeing (if it is a fact) may not be one of the formal or logical conditions for seeing but only one of the material or causal conditions. But it would be difficult, if not impossible, to decide which sort of condition the causal condition is, even if we had a clear statement of the specific causal mechanism.

It is clear that, with the exception of certain odd or borderline cases, we all know how to apply the word "see" with a high degree of success. This facility of ours with the expression must serve as a large part of our means for testing any proposal for what counts as the specific causal mechanism for seeing. If a specific causal mechanism is said, by some theory of perception, to be a necessary condition for seeing, and the account specifies a mechanism which does not happen to be present in a case which we are willing to call, without any reservations, "seeing", then that causal theory must be rejected. A consequence of this is that just so long as there are difficult or borderline cases, cases for which we feel ill at ease using the expression "see" and reluctant to prohibit its use for any good reason that we are prepared to give, we will have for those cases no test for the acceptability of any specification of the causal mechanism involved in seeing. If a theory claims that a certain causal mechanism is involved in seeing, and that mechanism does, in fact, operate in all the uncontentious cases, but in only some of the contentious cases, we could not have adequate grounds for accepting or rejecting

the theory. If we decide, on whatever grounds, to accept some specification of the causal mechanism, we may, upon doing so, refuse to use the expression "seeing" for any contentious case which does not fit with the theory. But we should then recognize that we are in no sense discovering the causal mechanism which happens, as a matter of fact, to underlie seeing; we are simply stipulating that this theory will be used, and so these cases will count as seeing and those cases will not. Even if our theory has a strong descriptive basis, it becomes at this point revisionary. This is what is so odd about Grice's suggestion. There is no good reason that the revision, clarification or precization of our language, if it is wanted, need be left to physiologists of perception. They are not "specialists" in such matters. Nor is there any reason not to do such revisions of our linguistic usage apart from or prior to any investigation into the physiological mechanisms; we already know where and when to use the word "see", and when there are difficult cases. If all we mean to do is stipulate that certain of the contentious cases will henceforth count as seeing, and certain others will not, we can do that with no understanding of the underlying physiological or optical mechanisms. In fact, if we are determined to make certain revisions in our use of the expression "seeing", then we would be better off to make those revisions by citing criteria for the use of the expression which can be applied by nonspecialists - by those who, as it turns out, use and need

to use that expression everyday. If a specialist needs a technical term, he would more usefully coin and institute a new expression, rather than create confusions by usurping an expression which already has a different meaning. If Grice's program is to make sense, and it is to be possible for specialists to discover the mechanism underlying all cases of seeing, then they must approach that enterprise with an account of perception already in hand; they must be guided by an account which can tell them which cases are cases of perception and which are not, even for the contentious cases, if they mean to uncover the actual mechanism, and not just stipulate what it will be. In short, they already need a philosophical theory of perception. Grice offers nothing towards such an account.

Just how specific any adequate specification of the appropriate causal mechanism for seeing must be can be seen by comparing the following cases. In the first case, on a cloudless day I glance into the sky and see the sun. The causal mechanism is, to a specialist, relatively clear. Photons ejected from the sun impinge on my retina, the photochemical response of the cones and rods of my retina sets off a transmission of ions along the neural pathways of the optic nerve culminating in an event in some part of my brain. I describe what has happened, if asked, by saying that I caught a glimpse of the sun. Contrast this case with the following case. On a cloudless day I glance into a mirror and see the sun over my shoulder. The mech-

anism is exactly the same as before, except that in this case the photons emitted by the sun take a less direct route to my retina - they are reflected to my retina by the mirror. A problem immediately arises for the second case. Do we still see the sun or do we only see a reflection of the sun when we look into a mirror? Whatever the correct answer turns out to be, understanding better the physics of photon reflection will not help us to answer that question. If we answer it at all, we must answer it on other grounds - we may refuse to say that we see the sun on the grounds that our metaphysics of perception allows us to view the reflection of the sun as a distinct public particular, identifying it with neither the mirror nor the sun itself. In that case, we will take the reflection of light by the mirror to count as a significant difference from the causal mechanism in the first case. Suppose, however, that we have no room in our metaphysics of perception for reflections as possible objects of perception. In that case we will not take the interference of the mirror to count against saying that we see the sun reflected; we will broaden our account of the mechanism of perception to allow for such cases. But having done that we encounter the third case. In the third case, I look at the moon on a clear night. The moon emits no radiation of its own, but is lit by the light reflected from the sun. How does the third case differ significantly from the second case? If we allow this to count as a case of seeing the sun reflected

by the moon, then we have no reason for saying that we see any of the objects in a room lit by an incandescent bulb - we can only say that we see the lightbulb itself. If, as it seems we must, we decide to say that it is the moon we see on a clear night, and not the sun reflected in the moon, then we must fill out the causal mechanism which counts as part of an analysis of seeing in such a way as to include the second case but exclude the third case. The relevant difference in the mechanism is, I suppose, that in the last case the reflected light is too diffuse or scattered for us to be able to see the sun reflected in the moon. As in our comparison of the first two cases, it is not our knowledge of the causal mechanism involved in seeing that allows us to decide that we do not have a case of perception; rather, it is our knowledge of seeing that allows us to say what the relevant causal mechanism for seeing must be. As the cases get more and more difficult, we will have to consider the effects of saying that we see in any given case on what we would then be committed to saying in other cases, or the effect of saying that we see in a given case on other commitments that we have already made in our metaphysics, our epistemology, or whatever. But in asking these considerations, we in effect end up with an analysis of perception. Not only Grice's version of the causal theory of perception, but the general project of giving an analysis of seeing by discovering the underlying causal mechanism, has the entire project backwards.

These sorts of difficulties are what lead me to say that there are serious difficulties in viewing the causal theory of perception as any kind of analysis of the concept of perception. But there is another objection to the causal theory, in the form of a counterexample.

Our statement of the causal theory of perception required some relatively exact specification of the causal processes involved in seeing an object. The counterexample that I have in mind counts as a counterexample to the weaker claim that for P to see X, P must be in some (interesting) causal relationship with X. It, a fortiori, counts against the statement of the causal theory given earlier, since that statement of the theory entails the weaker claim just mentioned. The counterexample counts against saying that it is even a necessary condition for P's seeing X that P be in some causal relationship with X. We have all heard, and have no good reason to doubt that physical theory is correct on the matter, that a flat black object reflects very little light to our eyes. It's black appearance is, in fact, a result of the object's absorbing most of the light incident upon its surface. We need only follow physicists in speculating on the possibility of a perfect black body, a body which absorbs all the radiation which would otherwise have some effect on our optical receptors, to find our counterexample to the causal theory of perception. The perfect black body would reflect no radiation to our eyes, but there is otherwise every good

reason to say that we can see such an object - if it moves, we can follow it with our eyes, we would know right where to reach for it, and we could even say how it looks to us (it would look a very dark black). If the case seems artificial, we need only compare it with the case of a dark shadow seen on a bright day. A shadow reflects no light, and so a shadow cannot be responsible for any events occurring on the retina of the eye; but there is no reason to say that we do not see them - in fact, we say all the time that we do. The only reason that could possibly be given for refusing to say that we do not see shadows, or could not see black bodies, is that they do not reflect light to our eyes, and so could not stand at the beginning of any causal chain which culminates in our seeing something.

It might be objected to the counterexamples given above that they do not, in fact, show what they claim to show simply because they are really cases of causation after all. It cannot be denied that if I see a perfect black body against a white background that the presence of the black body in my visual field makes a difference to the state that my retina is in at the time that I see the black body. And what is this notion of "making a difference", if not a causal notion? While I gladly concede that my retina would not be in the same states of excitation if the black body were not in my visual field, I deny that this means that the black body, or some event involving the black body, causes my retina to be in the state that it is. What

the black body clearly does is prevent light from reaching a certain part of my retina; it makes a difference to the state that my retina is in, in the sense that if the black body were not in my visual field then reflected light would cause that area of my retina to be stimulated in the same way that the rest of the retina is stimulated by the light reflected from the white background. And I am perfectly willing to concede that if the black body didn't make some such difference to the state of my retina, I could not be said to see it. But it makes a difference to the state of my retina without itself having any causal efficacy - it makes the difference that it does simply by virtue of not having any effect on me.

We began this section by saying that it seems perfectly natural to say that "P sees X" expresses some relationship between P and X, but we have now seen that the relationship cannot be a causal one. What else is left? An answer to this question will begin to emerge if we pose one more objection to the causal theory of perception. We can, I believe, make the following claim about seeing. There is no case of seeing an X where some part of our retina is stimulated in a certain way, and where that state of stimulation of our retina has some specific causal antecedent, which is not such that we can construct a case in which the same part of our retina is stimulated in exactly the same way with the same causal antecedents, but which is not a case of seeing an X. I have in mind the sorts of cases

discussed by Dretske in Seeing and Knowing, which include the following example:

Suppose that we attach a piece of beige paper to a beige wall and dim the lights until the paper appears (from where we are standing) as an undistinguished portion of the wall.¹⁴

Other examples might include a white sheet, seen at first on the dark ground, until the snow begins to fall and leaves the ground covered in snow of exactly the same color and brilliance as the sheet, but leaves the sheet itself uncovered. The important feature of such examples is this - the same object, the sheet, has continued to stimulate in the same way just the same portion of the eye of a stationary observer. If the causal relationship between the sheet and the man's eyes was originally sufficient for his seeing the sheet, and the same causal relationship now obtains, then he ought to be able to still see the sheet. But he cannot. We should not mislead ourselves into thinking that he must see it, since it reflects light to his eyes in just the same way that it did before the snow fell. We need only compare this case with the case of a perfect black body, in a perfectly dark room. As the lights gradually are brightened, there comes a point where the illumination is sufficient for us to see the black body. But notice that as the light increases, the other objects in the room also become visible, and if they did not we would have no background against which it would be possible to see the

black body. Now if we suppose that we must see the white sheet even after the snow falls around it, because it effects our eyes in the same way that it did before the snow fell, then we ought also to say that we see the black body before the lights come on (though we do not see any of the other objects in the room) because it also effects our eyes in the same way that it does when the lights are on and we see it - that is, it effects our eyes not at all in both circumstances. But just as we do not see the black body in the dark, we do not see the white sheet in the snow.

I hope that the examples discussed make clear what distinguished the cases in which we do see something from the cases in which we do not. The relevant difference seems to be, in Dretske's formulation, that we see what is "visually differentiated from its immediate environment".¹⁵ I will try to argue shortly that there are inadequacies in Dretske's own account of seeing as visual differentiation, but I shall devote the remainder of this chapter to developing and defending an analysis of seeing which adopts a notion of differentiation or discrimination, rather than a notion of causation, as the key to seeing.

E. Seeing As An Achievement

It will be useful to argue first that seeing is an achievement, and that all the other perceptual verbs are

achievement verbs as well. We can best do this by showing that the verb "see" fits the five criteria given for achievements in Chapter One.

First, it should be noted that "see" takes no present continuous tense if the verb is used in its perceptual sense. "I am seeing X" does not have a perceptual sense; I can be seeing the farmer's daughter, even though I am blind. If I am, over an extended period of time, viewing some object, the only things I can be said to be doing are things like viewing, scrutinizing, watching or observing. But those are all activities not achievements. To scrutinize, for example, the man in the lineup, I must see a number of things - his eyes, his mouth, his nose, and I may notice that he is trembling, watch the way he moves his eyes, and so on. Doing all of these things constitutes scrutinizing the man before me. If I remark, while turning the dial on the telescope and bringing the planet into clearer focus, "I am seeing it more clearly now", I either mean that I am getting the planet in better focus, or that I do see it more clearly now.

Second, I cannot start to see something, be half finished, and then stop before I have finished seeing it. I may start looking for something and stop before I find it, or I may start doing something which will lead to my seeing a certain thing, and then stop before I see it. But in none of these cases do I actually start to see something and not finish that.

Third, as soon as a simple present tense sentence, in which the main verb is "to see", become true, a present perfect tense form of the same statement is necessarily true. If it is true that I see the moon, then it is true that I have seen the moon.

Fourth, although I can only see something by performing certain acts, my seeing it does not consist in performing those acts. There may be no more I have to do to win the game than move one piece, or utter a single word. But winning the game is not the same as moving the piece or uttering the word; winning does not consist in doing any of those things. In the same way there may be no more that I have to do in order to see the star in the dusk sky than shift my focus slightly, or no more I must do in order to see the chair besides turning on the light, but my seeing the star does not consist in my shifting my focus, and my seeing the chair does not consist in turning on the light. Nonetheless, I do see the star by shifting my focus, and I see the chair by turning on the light.

Fifthly, when we count up the actions that we perform, seeing need never be one of them. Given that I am standing in the right place, facing the right direction, and my eyes are focused properly, I may do two things - reach up and flick on the light and turn my head to the left - and I will then see the chair. There is nothing else that I need do in order to see the chair. If, after doing these two things, I see the chair, it does not follow that I have

done three things - turned on the light, turned my head, and seen the chair. Although my seeing the chair does not consist in turning on the light and turning my head, seeing the chair is not something that I can or must do in addition to turning on the light and turning my head.

There is a very large range of activities which we can perform, none of which is identical with seeing, but any of which could be such that if we did not do it on a particular occasion, we would not see what we do see on that occasion. In order to see the Empire State Building, I may have to buy an airplane ticket. Buying the ticket is just one of many things that I must do, and it is one of the actions which I can subordinate to seeing the Empire State Building. There seems to be, however, a class of acts which play a much more central role in seeing, perhaps which play an essential role. I have in mind those acts which involve the use of the eyes - looking at, focusing, and so on. We may, however, view all these activities - buying the ticket, turning on the light, focusing and turning one's head - as culminating in an achievement of perception. The centrally located acts, those involving the use of the eyes, are necessary not for perception but for seeing. Hearing may also, on a given occasion, require that we turn our heads or, if it is the Berlin Philharmonic that we want to hear, we may also have to buy an airplane ticket. But if the series of subordinated acts do not end with acts which involve the use of the ears, we will not

call the perceptual achievement managed one of hearing. Whether a given perceptual achievement is one of seeing, hearing, smelling, tasting or feeling will depend upon whether the acts leading up to the achievement were acts of looking, listening, sniffing, eating or touching, and so whether the organs that we made use of in managing the achievement were the eyes, ears, nose, etc.

It is hardly enough to say that to perceive is to get a result or manage an achievement. To win a race or to checkmate an opponent is to do as much. We must also say what achievement, or what kind of achievement is managed when one perceives something. We can say what sort of achievement one manages when one wins a race. To win a race is to manage the achievement of crossing the finish line ahead of one's competition, after having run within the boundaries of a given track, when all of this is done in conformity to a certain set of rules. By going on in this way we can distinguish that species of achievement we call "winning a race" from those we call "checkmating" and "seeing". The problem before us is to characterize the sorts of achievements that we call perceiving.

The first distinguishing characteristic of the achievement of seeing is, as we have already seen, that we perform certain acts which result in our getting our visual organs into a position so that the presence of the object that we are said to see makes some difference to the state of those organs. I use the notion of "making some difference" in

the very broad sense discussed in the previous section, so that the presence of X, or X's being in the appropriate place and the appropriate environment with respect to P's eyes, can make a difference to the state of P's eyes even though X is not itself causally responsible for the state of P's eyes. The example of the black body shows that the appropriate positioning of X with respect to P's eyes can make a difference to the state of P's eyes simply because X is the thing, or one of the things, in P's visual field which does not cause a change in the state of P's eyes.

Everything we have seen so far strongly suggests that we should characterize the achievement reached by putting ourselves, and our eyes, in the appropriate position, as related in an important way to "visually differentiating" what we are said to see "from its immediate environment". Dretske not only takes visual differentiation to be importantly related to seeing, but makes it, in effect, his definition of seeing. He writes that "S seeⁿ D = D is visually differentiated from its immediate environment by S"¹⁶ and that "S's differentiation of D is constituted by D's looking some way to S, and moreover, looking different than it's immediate environment"¹⁷. In view of the specific nature of the failure of the causal theories of perception, Dretske's analysis looks to be a promising alternative. In the next section, however, I will argue that Dretske's analysis is inadequate, not because he analyses seeing in terms of visual differentiation, but because of

the account of visual differentiation that he gives. It will then be open to us to provide an alternative analysis of seeing.

F. Visual Differentiation

The first difficulty with Dretske's analysis that I would like to mention is that the definition allows that we see afterimages and hallucinatory objects. I have already given my reasons for thinking that such a consequence is undesirable. Of course it is always possible for Dretske's account to stand as is, and for the class of possible objects of perception to be narrowed down on other considerations - considerations external to the definition of seeing. I suggest, however, that it would be preferable to having an account of seeing which, on considerations internal to the definition, would limit the class of possible objects of sight - that is, it would be better if we could understand why certain things are properly said to be seen and others are not, simply by knowing what sort of thing seeing is. The criticism is very weak, since it depends upon it already being agreed that afterimages and hallucinatory daggers are not possible objects of perception.

The second point is, however, more serious. It does not seem to me to be possible to reconcile Dretske's account of seeing with the claim that seeing is an achievement. To say that something, D, is visually differentiated by S

makes it sound as though visually differentiating D is something that S does, an act he performs or an achievement that he manages. But S is said to visually differentiate D when D looks to S different than its immediate environment, and D's looking some way to S hardly seems to be an achievement of S's. If there is a relation between the visual differentiation of D by S, and D's looking some way to S different than its immediate environment, it seems more likely that it is not a relation of identity, but rather that when one does see D (when one manages that achievement) then D will, as a contingent matter of fact, look some way different than its immediate environment to S. If seeing is an achievement, and it is the achievement of visually differentiating what is seen, then it seems unlikely that visually differentiating something could be for that thing to look different than the rest of its environment.

Dretske's definition is problematic on other grounds. Although he notes the problematic cases himself, it does not seem to me that Dretske deals with those cases adequately. We will look at two such cases, using Dretske's own examples. In the first case he mentions what he calls a "minor complication". He supposes that after S "runs his eye over the entire surface of a wall" someone asks him if he saw a certain triangular section of the wall which the interlocutor indicates by tracing out a certain section of the wall with his finger. Dretske agrees that the natural answer to the question would be yes; "even

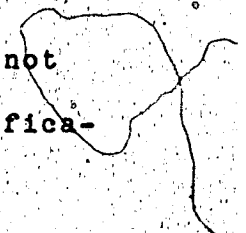
though, at the time S saw the wall, he did not visually differentiate that particular section of the wall from the remainder of the wall's surface".¹⁸ The reason that he failed to differentiate that part of the wall is simply that that part of the wall did not look any different to him than did the rest of the wall - it all looked the same. Although Dretske agrees that in cases¹⁹ like the white sheet surrounded by snow, we ought not to say that we see the sheet because it does not look any different than its surroundings, he makes the case just cited an exception, saying that

in such cases as this I think we can say that S saw that (triangular) section of the wall simply on the condition that that part of the wall looked some way to him at the time he saw the wall.²⁰

Thus far there seems to be no relevant difference between this case and the case of the white sheet in the snow, and one would expect Dretske to say the same thing about both cases. Either S does not see that part of the wall or, if he does, then he also sees the sheet. But he treats the two sorts of cases differently, and tries to justify his doing so by saying that

Differentiation is unnecessary when the individuality of what is seen rests solely on such arbitrary tracing operations or exclusively on its relative spatial location and configuration.²¹

But the example under consideration clearly does not fit the original definition of seeing, and this justifica-



tion looks like an ad hoc adjustment made to accommodate what must otherwise be seen as a counterexample to the definition. To bring out how I think Dretske's account goes wrong in dealing with this case, I mean to modify his story slightly to produce a similar case. Instead of 'running his eye' over the entire surface of the wall, suppose that S stood at a distance of, say, twenty feet from the wall and opened his eyes only for a moment - long enough to take in the entire wall at a glance - and then closed them again. He saw the entire wall, and from the distance at which he saw it, if there had been a dark triangular spot near the centre, he could easily have seen it. But since there was no dark spot near the centre, and the wall was of a uniform light color, he did not see one. The question then is whether he saw that part of the surface of the wall which, had a certain triangular spot been there, would have been covered. In this case, though, he took in the entire surface of the wall at a glance, he did not run his eye over each part of the surface. Although he would, if asked, naturally reply that he saw the entire wall, and though he may in fact have just happened to have looked at and taken note of just the spot of wall we are concerned with, it is more likely that he would not want to say that he saw just that part of the wall - unless of course he just happened to have given, briefly, that part of the wall some of his attention. But if this latter sort of thing did not happen, he would correctly reply that he

did not see that particular section of wall, that he did not, in fact, pay attention to any particular part of the wall, even though he would agree that he did see the entire facing surface of the wall. The difference between Dretske's original case and this second case is simply that in the original case he ran his eyes over the entire surface of the wall; he gave, as it were, each part of the wall some of his attention. He not only looked at the wall in its entirety, but he also looked at each part of the wall separately when he ran his eyes over the entire surface of the wall. He looked at, focussed his eyes on, and gave his attention briefly to the part of the wall which was later indicated, since he did as much to every part of the facing surface of the wall when he ran his eyes over the wall. Had he been searching for a small mark on the wall (though there was none) he would have to run his eyes over the wall in just the same way. To say, after conducting such a search, that he did not see that part of the wall would be to concede that he had not conducted a thorough search. But since he looked at, and even examined that part of the wall, we say he saw that part, even though it did not look any different than any other part of the wall - it was not visually differentiated in that sense. In the second case, however, he did not look at, focus his eyes on, or give any particular attention to that part of the wall. We should, I suggest, compare the first case with the case in which, staring into one part of the dusk

sky, we catch sight of a star only after varying our depth of focus. Given that certain conditions are already met, there is no other act we need perform other than, or in addition to, varying our depth of focus in order to manage the achievement of seeing the star. Had we not focused our eyes at the appropriate distance, we could not have seen the star. Likewise, had S not performed some of the particular acts that he did, he would not have seen the part of the wall that was indicated. There is, of course, another difference between the two cases - the star looks different than the surrounding sky, but the part of the wall does not look different than the surrounding wall. The fact that the person searching for the mark did see the part of the wall indicated, however, shows that the relevant difference between the first case (in which the part of the wall is seen) and the second case (in which the part of the wall is not seen) has nothing to do with how the part of the wall looks, but with the additional acts performed in the first case but not in the second. What this example shows is that the achievement of seeing something is managed only after the performance of sufficient actions and, although it may be the case and is almost certainly generally the case, that when the achievement is managed what is seen will look different than its surroundings, it cannot be a necessary condition that what is seen looks different than its surroundings.

There is one last case that must be discussed in considering the inadequacies of Dretske's original definition. The last cases discussed have not given us reason to reject the definition of seeing as visually differentiating, but only reason to reject the account of visual differentiation given. By viewing the visual differentiation of X as an achievement managed by performing actions, we can yet say that a person who performs sufficient actions of the right sort has done enough to have visually differentiated what he is said to see. In managing the achievement of differentiating some part of the wall, by performing acts from a certain range, essentially involving the positioning and use of the eyes, we perform sufficient acts to have managed the achievement of seeing that part of the wall, whether or not that part of the wall looks any different from the rest of the wall. The last case we will look at points, I believe, to a similar conclusion. This case Dretske calls a "limiting case" on this definition of seeing. Once again, however, it is not clear why this should be viewed as a limiting case and not as a counterexample:

Touch your nose to a large smooth wall and stare fixedly at the area of the wall in front of you. There is not much doubt that you see the wall, or at least a portion of it. It is also fairly clear that you do not differentiate it from its immediate surroundings.²²

Dretske suggests that, in this case and others like it, "the differentiation clause becomes imperative when nothing

appears to S that is not part of D". Once again the adjustment in the original definition seems ad hoc; the case simply does not fit the definition of seeing. But once again the acts performed seem to be sufficient for saying that the wall is seen - our eyes are open, we are, if we are not so close that we cannot, focusing on the wall, and are attending to it - we "stare fixedly" at it.

We saw, in our discussion of causal theories of perception, that there seem to be very good reasons to prefer an account of seeing in terms of differentiation to an account in terms of causation. But the "minor complication" and the "limiting case" mentioned by Dretske seem to indicate that the account of visual differentiation he offers is inadequate. At the same time, I have been developing the achievement side of seeing, and trying to show how the performance of certain subtle acts like focusing, or shifting one's attention, can make the difference between reaching or failing to reach that achievement. We have, I believe, prepared enough ground to propose an alternative definition of seeing.

G. Perception and Discriminatory Action

Thus far I have been urging that, as we keep sight of the achievement aspect of perception, we should give those actions, the result of which is the perceptual achievement, their proper due in the analysis of perception. There is an extremely wide range of actions which can be subordi-

nated to managing a given perceptual achievement, some of which might involve the use of the legs (walking close enough to see), some of which might involve the use of the arms or hands (turning on the light, cupping one's hand around one's ear in order to hear the noise over the wind), but the total collection of which necessarily include some actions involving the positioning or the use of the sensory organs. As looking is related to seeing, touching is related to feeling, and tasting (the action) is related to tasting (achievement). In the case of seeing, there is a relatively wide range of actions which involve the use and positioning of the eyes - turning one's head, scanning the horizon, squinting against the glare, focusing, looking around, and so on. In so far as the diverse perceptual achievements of seeing, hearing, tasting, smelling and feeling can be distinguished by the separate sensory organs involved, they can also be distinguished by the diverse acts involving the separate organs. Sniffing is subordinated to smelling, and not to hearing. Some of the acts can be on a large scale, like walking; others, like concentrating one's attention or focusing one's eyes, can be more subtle. In any given case it may turn out that a large number of such acts will be necessary in order to perceive a given thing, or it may turn out that only one very subtle act is needed, depending upon such things as the agent's proximity to the object, the viewing conditions, the size

of the object; in general, it will depend upon the agent's present perceptual position and circumstances.

For some achievements it is possible to state in a relatively brief way what acts are necessary and sufficient for managing that achievement, or to state in a useful, noncircular way what results must be got before one can be said to have managed the achievement. There are many ways to run a race, but when a runner crosses the finish line ahead of his competition, after engaging in the activity specified by the rules and in the way specified by the rules, he wins the race. Likewise, if a chessplayer makes moves, in turn, which are allowed by the rules of the game of chess, and puts his opponent's king in check while leaving no avenue of escape or protection, he manages the achievement of checkmating his opponent. But it is unlikely that any winner in a game of chess performed all the same subordinate actions as any other winner. What we need, in order to have an analysis of perception is some account which tells us when the acts performed by an agent, or the results got, are adequate for him to be correctly said to perceive something. When are an agent's actions sufficient, however much they may differ from the actions of another agent, for what he has done to guarantee the result that he perceives? By answering this question we will provide an account of perceptual achievements. I will proceed by offering an account of the achievement of seeing, which can be generalized in obvious ways to give an

account of perception in general, and then explicating the crucial semitechnical notions involved in the account. I propose to define seeing as follows:

To see X is to manage, as a result of performing acts involving in an essential way the use and positioning of the eyes, the achievement of reaching a position good enough for performing discriminatory action which has X as its object.

There are two parts of the definition which will need some explication, the notions of "reaching a position" and of "performing discriminatory actions".

The notion of a position, as I am using it, retains all of its nontechnical meaning as, roughly, spatial location, or the placing and posture of various parts of the body. If by crawling through the grass I am performing the activity of positioning myself to shoot, then when I reach the top of the grassy knoll I am, after that moment, in a position to shoot. Given that my eyes are properly focused, that I am facing in the right direction, and so on, the act of stepping one pace to the left may be sufficient for my then seeing something; there need be nothing more that I must do at that point besides stepping to the left. My motion to the left counts as a change in my position, since I move from one position to another. In this case, reaching the new position is sufficient for seeing something which was, at an earlier time, hidden from my sight. Included in the notion of position used in the definition is the positioning or posture of my body - in particular,

the head and the direction in which I am facing. The total position that I am in at any given moment will include, as one constituent part, the direction in which I am facing; thus, by turning my head I can reach a new perceptual position. In some cases there need be nothing more that I must do in order to see something. Once again it must be emphasized that if the view that perception is an achievement is to be taken seriously, then, in the last case mentioned, we cannot suppose that after turning my head, I must do something else in addition; that is, that I must turn my head and then also see X. Seeing X is, in the example, the achievement got as a result of my turning my head. Finally, ~~the~~ notion of position with which we are operating must be extended to include the changes effected by the subtler acts involving the use of my eyes. How my eyes are focused, whether or not I am squinting against the glare, and so on, are all relevant to the total position that I am in. This extension of the ordinary notion of position is made in order to allow us to count any of the acts we perform in subordination to seeing as changes in position. In an example discussed before, if I am already staring in the right direction into a dusk sky and then shift my focus ever so slightly, I may only then see the star which has, all along, been right in front of my eyes. This slight shift of focus counts as a change in position. Once again, besides shifting my focus slightly, there is nothing else that I need do in order to see the star,

providing that the other constituent parts of my total position are such that I need not, say, turn my head, open my eyes, or brush aside an overhanging branch. If there was something else called seeing the star which I must do, then seeing could not be an achievement.

Finally, it should be added that I will include external factors generally put under the rubric of "perceptual circumstances", in the notion of position. It may be that, because it is simply too dark, none of the actions I can perform will put me in the position of seeing something. But if I wait for the sun to rise, there will come a point where walking closer or squinting into the dark will improve my position sufficiently so that I will then see the house on the hill. Or I may shine a flashlight in that direction. I may, by turning off a flashlight in my hand, so impair my perceptual position that I can no longer see something which my earlier position was good enough for seeing. If we were, as God is said to be, all-seeing and omnipresent, there would be no such thing as our improving our position; but if God was literally all-seeing, it must be a contingent feature of Him, since even if His omnipresence gives Him the best of all positions from which to use His eyes, it at least makes sense to suppose that He might have them closed.

Thus far, the use of 'perceptual position' is limited in such a way that two perceptual positions will count as different perceptual positions when the list of objects of

perception differs for those two positions. Since my primary focus is not going to be what is seen, but rather a judgement as to the color of what is seen, I will extend the notion of a perceptual position in one additional way - to include certain states of a person which may vary or change even though there is no ordinary sense in which the person has changed or shifted position. In particular, I will count as differences of position those differences in the states of a person which can make a difference in how things which they see appear to them. So, for example, if two persons are in perceptual positions which are otherwise indistinguishable, but one of them is under the influence of hallucinogenic drugs, I will say that their perceptual positions differ in just that respect. Exactly which states of a person are a part of that person's perceptual position and which are not becomes, for the next chapter, an important concern. I will argue there that no change in a person's propositional states can count as a change in that person's perceptual position - that is, that what we see and how it looks is independent of what we know or believe. It should be clear, however, that changes in the state of a person's sensory organs can and do make significant differences in what that person sees or how what they see looks to them. In general, I will say that if two persons see different items, then they are in different perceptual positions. An examination of any particular case in which two persons see different things or in which

the things seen look different to the two, ought to reveal a difference of the kind I have been describing - a difference in the position from which each of them sees. I take it to be a matter for empirical research to determine if a difference in, for example, what two people desire or how strong their desires are, can make a difference in what the two people see or how what they see can appear to them. I will argue, however, that we have a priori grounds for saying that differences in propositional attitudes cannot make for such differences in perceptual position. So the difference between the perceptual position of a person who sees the first star at dusk and a person who does not may be only a difference in the slightest shift of focus, but it may also be a difference in how much they have had to drink. It cannot be a difference in what they think.

By making differences in how things look (as well as differences in what is seen) count as differences in one's perceptual position, I have not modified (but only filled out) the substance of the central notion of a 'position'. The reason for this is that the sort of position I am talking about is tied to a notion of discriminatory action. I will proceed next to a discussion of discriminatory action, but it may be noted in advance that, under the account of discriminatory action to follow, what sorts of discriminatory acts we are in a position for is a function of both what is seen and how things look to us. If, for example, we see berries, some of which are red and some of

0

which are purple, then we are in a position to engage in discriminatory behavior - to treat berries differently from other things in our environment. But if the red berries and the purple berries look the same to us in certain circumstances, then we are not, in those circumstances, in a position to treat red berries any differently from purple berries. (A racist to whom blacks and whites suddenly become indistinguishable could not continue his racist practices.) Establishing and maintaining any discriminatory practices whatsoever depends on there being differences in how things look to us.

As I have characterized the notion of position to be employed, our total position at a given time might be such that either we do or do not see a given object. From my present position, I will see an ashtray if I lift my head slightly and focus at a distance of about six feet. If I do not focus my eyes properly, I will not be able to see it, even though I lift my head. But lifting my head will improve my position (for seeing that ashtray) since there will then be that much less that I must do in order to see the ashtray. What is now needed is a way of characterizing one's position when it has been so improved that one does see; a way, that is, of saying how good one's position must be before it is necessarily good enough so that it no longer makes sense to say that, in that position, one does not see. In the definition given above, I suggest that one's perceptual position is good enough if it is a posi-

tion from which one can perform discriminatory actions which have as their object what it is that one is said to see. I will next explicate the relevant notion of discriminatory action with an object, and then will show in what sense a certain position will count as one from which one can perform such actions, while from a different or worse position one cannot.

I will say that an agent, P, performs a discriminatory action or engages in discriminatory behavior or activities if P performs an action which has some individual, X, as its object. The action or activity will have X as its object if the action is directed towards, at or away from X. So pointing at something, chasing it, fleeing from it, stepping over it, or turning one's head to follow it would be actions which have something as their object. Activities like walking, writing and playing are not discriminatory actions, though following someone, writing to someone, or playing with something are. Falling on something, tripping over it, or bumping into it are not discriminatory actions, simply because they are not actions.

An action's having an object is not the same as it's having an objective, though the two are importantly related. Without mentioning an agent's objective in performing a particular action, it is difficult, if not impossible, to make clear just what that agent is doing. Thus, to say that I am running is to give a very incomplete characterization of what I am doing, since it will be true that I am

running if I am running a race, running after the bus, or running from a rabid dog. To say that I am running in order to improve my health is to give a more specific characterization of what I am doing than is given by simply saying that I am running. It is part of the consequence of working within a teleological model that, since everything which counts as an action does so by virtue of aiming at some objective or goal, there is no action which cannot be so described as to include mention of some objective towards which it is directed. Saying what an act has as its objective sometimes involves mentioning an object, and sometimes it does not. Discriminatory action is that action which is such that a complete statement of the objective of the action requires mentioning some object - having, avoiding, following, etc., that object is the objective of the action. That is the difference between running down a path which happens to lead away from the bear and running from the bear. Running in order to improve my health and running from the bear are both actions which have objectives, but only the latter action takes some object - the bear - as an object. "He is chasing", "He is following", and "He is pointing" are incomplete characterizations of a person's action which are not completed by saying "He is chasing for his health", "He is following out of curiosity", or "He is pointing in order to show you". On the other hand, "He is chasing rabbits for his health", "He is following you out of curiosity" and "He is pointing at the star in order to

show you ..." do complete the characterizations of his actions. Unless there were such things as actions directed towards, at, or away from objects, there could be no difference between, say, stretching and reaching, chasing and running races, and running after and running away from something. Actions such as these, which involve in an important way in their complete description mention of an object, are the ones I mean to label discriminatory actions. But a few more points will need to be made to sufficiently delimit the relevant class of actions. First, if P is A-ing X, there is no commitment made in saying that, or in its truth, that P knows or thinks that what he is doing is A-ing, nor that what he is A-ing is an X. P need have no knowledge or opinions as to what he is doing, or as to what the object of his action is. I may see something on the horizon and move closer to it to investigate without having any idea as to what it is; or a child may reach for a shining object passed before its eyes without knowing that that is what he is doing. It may well turn out that there is always some mental act or state of P which has X as its intentional object - if there were not, then we may in the end fail to be able to make a distinction between P's just happening to be moving in a direct line away from X and his avoiding it, or running from it. My point is only that P need not have any propositional attitude towards X. I have already argued in earlier chapters that the behavior of nonlinguistic creatures must be understood and explained

without presupposing that they have either knowledge or beliefs. A fortiori their discriminatory behavior does not need that presupposition, even if it does turn out to be true that they must be in some mental states or be performing some mental acts. But if any sense can be made of saying that mental acts or states can have, or must have, intentional objects, then the class of acts I wish to call discriminatory are just those public acts which are motivated or explained as being a result of such mental states and acts. Of course the intentional object of some mental state need not be actual, though the objects of discriminatory acts must be. I cannot run from a ghost, if there are no such things as ghosts, though I can run from what I believe to be a ghost. Mental acts which have nonexistent intentional objects do so largely via the intermediary of a false belief, though not necessarily. I may wish for a million dollars, but my belief that there is such a thing as a million dollars is not a false belief, though even if it were false that there is a million dollars left anywhere, I could still wish for a million dollars. But what I cannot imagine that there could be is a case of a mental act or state which has as its intentional object something which does not exist and which is not accompanied by some belief or other about that object. It is this feature of mental acts and states which have nonexistent objects as their intentional objects, together with the fact that discriminatory acts seem to require agents who are in such

states or performing such mental acts but who need not have any beliefs about the object of their discriminatory acts, which lead me to exclude such things as 'running from' ghosts from the class of discriminatory acts. If there are no such things, we cannot run from them, run after them, or point to them, though we may think that that is what we are doing. A consequence of this for the definition of seeing is that there could be no such thing as seeing something which does not exist - unless there are such things as hallucinatory daggers, we could not see them. As we have said already, hallucinatory daggers are not one among the many kinds of daggers that there are - they are not real, but neither are they one kind of artificial dagger; they simply do not exist.

A final point about discriminatory acts is that no discriminatory act has as its object a universal. We have already seen reasons for denying that we can point to a color, and not just to a colored thing. The same sort of considerations show that it is senseless to say that we can run from the color of an object, or chase it, or follow it, with our eyes.

The role that I have suggested discriminatory action be given in an analysis of perception ought not to be confused with the role that the acts subordinated to the perceptual achievement have been given. The two 'levels' of action are separate, but are connected via the notion of reaching a position. The connection is this: A certain position is

reached as a result of my having performed certain acts - the reaching of the position is the achievement to which the acts are subordinate; at the same time, it is by virtue of my having reached a certain position that certain discriminatory acts are made possible for me. We were looking earlier for a way to characterize the achievement reached when the actions one performs are sufficient for seeing something. My suggestion is, then, that one reaches the achievement of seeing X just when the position one reaches, as a result of performing the relevant acts, is good enough to make possible further acts of a discriminatory sort which have X as their object. When one reaches a position which is such that by virtue of being in that position it becomes possible for one to perform discriminatory acts, one necessarily perceives X. Prior to performing the acts which put one in just that position, one's (earlier) position was not good enough for performing discriminatory acts which had that thing, X, as their object. There are three points to emphasize: First, one's being in the position one is in must be the result of having performed the relevant acts if the position is to count as one of seeing X. That is, the acts must include some use and positioning of the eyes. Secondly, if the position one is in is to count as seeing X, then it must be by virtue of being in that position that certain discriminatory acts are made possible. Thirdly, being in that position only makes possible, in a

sense to be explained, certain discriminatory acts - none need ever ensue upon reaching that position.

If I am standing looking into the dusk sky and shift my focus slightly, I will catch sight of a certain star. The last act, of shifting my focus slightly, so altered my position that the new position I reach makes it possible for me to tell you what color the star is. But I may not say anything. My total position prior to that final act was not such that it was possible for me to tell you that. If I spot the rabbit on the hillside; it thereby becomes possible for me to point at it, or to point it out to you, even though my position is not yet good enough to tell you what color it is. Prior to spotting the rabbit, I may have pointed in the same direction, and even though the rabbit was in the same place and you managed to spot it as a result of my having pointed and directed your attention to that spot, it was not by virtue of my being in the position that I was in that it became possible for me to point at the rabbit. Had I perhaps looked longer, or more carefully, or squinted, I might have succeeded in so altering my position that then, precisely as a result of my having performed that extra act and so of my having reached that new position, it would have been possible for me to show you where the rabbit was standing. Although I might have accidentally got the same result without having seen the rabbit myself, it would not then have been the position I was in that enabled me to do it.

This account of seeing is, I believe, easily generalized to become an account of perceiving. The difference between seeing, hearing, smelling, tasting and feeling is a result of the different subordinate perceptual acts and the organs of sense which are used in reaching a position adequate for discriminatory behavior. Focusing may put me in a position which makes possible further discriminatory behavior on my part, but so will listening carefully, sniffing and touching.

I suggested that when one perceives something one reaches a position good enough for performing discriminatory acts. The difference in my total position before and after shifting my focus slightly is that before performing that act I was not (but after, I am) in a position good enough to go on to perform discriminatory acts which have as their object a certain star. The range of discriminatory acts which can have the star as their object can be very wide - I can point at it, shoot at it, hide from its glare or, if I have the necessary linguistic resources, I can describe it. But it should be noted, as the point will be of greater importance later, that the definition only requires that some discriminatory act be made possible as a result of one's being in a position of a certain sort - of the sort we call "seeing" - not that any or every discriminatory act is made possible by reaching such a position. Certain discriminatory acts require, if they are to be made possible by virtue of one's reaching a certain position, a

much better position or a much different position than certain other discriminatory acts. For example, though I may reach a position adequate for the discriminatory action of following a certain object as it moves, the same position may not be good enough for me to perform the discriminatory act of saying what color it is. So even though I may be in a position good enough for seeing the object, I may not yet be in a position good enough for saying what it is, or for saying what color it is. Reaching a very different kind of position may be necessary for performing different kinds of discriminatory acts, but so long as I reach a position which is such that by virtue of being in that position some discriminatory acts are made possible for me, then I perceive the object of such discriminatory acts. In the next chapter we will see just what sort of position is such that it makes possible certain epistemologically basic acts as judging what it is that is seen, or judging what color it is. For now we need only note two points: First, that a position good enough to make possible some discriminatory acts need not be good enough to make possible every discriminatory act or every kind of discriminatory act; secondly, though I have spoken of certain acts of judgement as belonging to the range of discriminatory acts made possible by reaching certain perceptual positions, not every act of judgement is made possible by virtue of reaching some perceptual position, nor is any act of judgement possible simply by virtue of one's being in

some perceptual position - other conditions must be met as well.

It is, of course, manifestly false that I must perform some discriminatory act or other before it is true to say that I see the star, or the tree across the street. But I have not suggested so much. My claim has simply been that, say, after a final shift of focus there are discriminatory acts which I can perform, by virtue of my having done that, whereas before shifting my focus I could not have done those things. In some cases, this is more obviously a logical truth than in other cases - seeing that there is a star in the sky involves, as we shall see before long, certain acts of judgement, but I could not see that there is a star in the sky, though I might come to know it in many other's ways, unless I see the star, and in the case we have just mentioned, I would not have seen the star unless I had shifted my focus. Performing that final perceptual act, though it did not guarantee that I would come to know that there was a star in the sky, made it possible for me to perform the acts which would have that result. My general claim has been only that to perceive something is to manage the achievement of reaching a position from which I can do certain things. The account of perception will not be complete until we see in what sense of the word "can" seeing has the result that I can do certain things - in what sense reaching a certain position makes it possible for me to do certain things. Let us distinguish four

senses of "possible" or "can" - it will be relatively easy to see that only one of them serves in the analysis of perception.

There is, first, a sense of logical possibility - I can draw a triangle on the blackboard, in that sense, even though I do not, as it happens, have a piece of chalk handy; there is no contradiction involved in saying that I draw such a triangle. But I cannot draw a round square on the blackboard, because there is a contradiction involved in saying that I do draw one. My thesis about perception is not, however, so weak as to say no more than that there is no contradiction involved in saying that when in a certain position I point to the tree across the street. Nor does being in such a position make it logically possible for me to point to the tree - that was always logically possible.

There is, secondly, a sense of "possible", often called epistemic possibility, which has, roughly, the sense of "might" or "might be" - it is epistemically possible that Joe is in town if there is some, however slight, reason to think that he is in town, and no conclusive reason to think that he is not in town. But my being in a certain perceptual position, and so my seeing something, cannot mean that there is some reason to think that I might perform some action directed towards or away from the tree. If you know that I am afraid of trees, and that I see it, that might give you reason to think that I might run, but my seeing

the tree cannot entail that there is some reason for anyone to think that I will do anything about the tree or to it. If the analysis of perception I have offered is right, then there will be some sense of "possible" or "can" which is such that "I see the tree" entails "I can do A", or "It is possible for me to do A". Epistemic possibility always, as it happens, takes the form "It is possible that S", where the "S" following the that-clause is filled out by a complete statement. The sense in which seeing something entails that I can do something, which prior to seeing I could not, always takes the form "It is possible for me to ...", as perusal of the preceding discussion will show. This sense of "possible" is not followed by a that-clause, but needs only an act or an achievement verb to fill it out. There are two major uses of "possible" or "can" which take this form, the "can" or "possible" of ability, strength or power, and the "can" or "possible" of opportunity.

It is possible for me to lift a small table. I can lift it because I have the strength. But it seems obvious that focusing my eyes, and so reaching the appropriate position to be said to see a star, does not give me any additional strength or powers for performing discriminatory acts which have the star as their object. If I do not have the ability to lift my arm and point to the star, focusing my eyes will hardly help me to acquire that ability. And if I have the necessary abilities to chase and catch the

rabbit, losing sight of him will hardly take away those abilities.

The final sense of "can" or "possible", and the one which seems to me to be needed in the analysis of perception, is the "can" or "possible" of opportunity. I may have the strength to lift a small table, but still cannot lift one, simply because there are none around. If I climb the ladder I will be able to pluck an apple from the high branches, even though climbing the ladder gives me no new powers or strength, but simply because from the top of the ladder I will be in a position to reach the apples - I will have an opportunity to do something which those standing on the ground do not have the opportunity to do. On the other hand, I may have an opportunity, which no one else has, to pick the apple, but not be able to do it because my arms are too tired from climbing the ladder - I lack the strength to do it. Opportunities to perform certain actions are just the kind of thing that one can acquire solely as a result of a change of position; but strength and powers are not altered by a change of position - only one's opportunity for exercising them. I may have all the skills and powers of a great hunter, but never succeed in bagging a trophy, since there is no game around, or I may be a poor shot and an inept tracker but have a wall full of trophies, simply because the game is so plentiful. In the first case, I would have sufficient skill, but no opportunity to use those skills to a certain end, and in the second case, the

special opportunity I have eliminates the need for much skill in order to reach the same ends. One can be able to do things that others fail at either because one has some special skill or powers, or because one has an opportunity which no one else has.

Not every opportunity that one has is a result of the place one happens to be - I may have the opportunity to campaign for office, if I wish, only because my name was drawn in the lottery. But sometimes to have a certain opportunity to do a certain thing is nothing more than to be in a certain place - the man atop the ladder has the opportunity, simply by virtue of being there rather than at the bottom of the ladder, to pick the best apples, and the guard posted at the door has the opportunity, simply by virtue of being where he is, to catch a glimpse of the Queen. My suggestion has been, however, that being in a certain position, in the semitechnical sense of "position" we have been employing, is also to have an opportunity to do certain things. Simply turning my head (and so catching sight of the charging bull) puts me in a position, if my eyes are open and functioning, etc., to jump out of the way of the bull. The person who was like me in every other respect, but didn't turn his head, didn't manage the achievement of reaching a position adequate for discriminatory behavior of the sort I showed, not because he lacked the ability to jump out of the way, nor because he felt suicidal, but only because he lacked one kind of opportunity

for doing that. His position was not good enough so that by virtue of being in that position it was possible for him to avoid the oncoming bull - that is, he didn't see it coming. If, as a result of turning my head, I reach a position good enough to make possible avoiding the bull, that is only because by turning my head I catch sight of it.

We now have, I suggest, given a complete account of seeing. For P to see X is for P to reach a position, as a result of actions involving the use and positioning of the eyes, being in which position is identical with having one species of opportunity - the opportunity to perform discriminatory acts, and which opportunity one acquires solely by virtue of reaching that position.

Earlier we saw that there are reasons, not necessarily tied to any particular analysis of perception, for saying that only public particulars are possible objects of perception. It should be noted that the account of perception just given has the same conclusion as its consequence, and so the account of perception provides us with, I think, a better understanding of why the possible objects of perception should be so limited than an account which does not have such a consequence. Considerations limiting the possible objects of perception ought to, if possible, be internal to the analysis of perception itself. Universals are not possible objects of discriminatory acts of the sort which are made possible by virtue of one's reaching a

certain position. I can run from the bear, but not its color, even if the bear is such that I would not have run from it unless it had been of a particular color. There are no such things as perceptual acts that one can perform in coordination to perceiving logically private states, and so no such thing as improvements or changes in one's perceptual position which would make possible the achievement of perceiving such things. I cannot have a pain and yet fail to feel it until I grope for it and touch it.

The relationship between perception and discriminatory behavior can profitably be compared with the relationship between sensations and some forms of overt behavior as that relationship is understood on certain theories of sensations. One of the concerns which motivates this account of perception is the concern that motivated Wittgenstein's arguments against the possibility of a completely private language. Our fixing the use of the word "see" in a public language depends upon there being public criteria for seeing, just as the word pain is fixed and kept in circulation with the support of public pain behavior. As it stands, the account of perception offered keeps the relationship between a given case of seeing something and a given piece of public discriminatory behavior a contingent relationship. It remains, in fact, a contingent matter that there should be any piece of public behavior associated with a given instance of perception. Likewise, a Wittgensteinian analysis of pain gives an important role to

certain kinds of public behavior, though there is not thought to be any strict logical relationship between the truth of "P is in pain" and "P shows pain behavior". He can either be in pain and not show pain behavior or show pain behavior and not be in pain. Nonetheless, the pain behavior is thought to be a natural expression of pain - a dog normally limps when his foot hurts, and the man normally winces when pricked by a needle. Although Wittgenstein did not do so, one could make the connection between the sensation of pain and the public pain behavior into a logical connection by introducing the semitechnical notion of a disposition into the account of sensation. Thus, on such an account, "P is in pain and P exhibits no pain behavior" remain perfectly intelligible, but "P is in pain and P has no disposition to give public expression to his pain" is inconsistent. To have a pain is, on such a view, to be disposed to do such things as cry, wince, limp or scream. It is only by introducing some such notion as a disposition, which is always a disposition to do such and such, that the behavior can be imported into the analysis of pain. The notion of position plays just such a role in the account of perception that I have offered. By introducing the notion of a position, which is always a position for doing such and such, we can import the discriminatory behavior, which may or may not ensue upon perceiving something, into the analysis of perception. "P sees X and P performs no discriminatory acts having X as their object"

could be true, but "P sees X and P is not in a position for discriminatory acts having X as their object" is inconsistent, since to perceive something simply is to be in such a position. The need for the notions of disposition and position is roughly the same in the two theories. Both terms serve as intermediaries between a logically private analysandum and the public behavior with which it has, logically, an unhappy relationship.

If we return now to an example given much earlier we can see the usefulness of such a move. I mentioned a case in which we observe a friend walking, lost in thought, down the opposite sidewalk. Each time that he passes under an overhanging branch he ducks his head slightly. When asked later if he saw the branches, he replies in all sincerity that he did not. It would be perfectly proper for us to reply that he must have seen them since he ducked to avoid each and every one of them. But of course it does not follow from the fact that he ducked that he saw them - it may have been that he had a peculiar muscle spasm in his back just as he passed each of the overhanging branches. My point is only that we rightly take our friend's movements to provide us with a good reason to say that he saw them - even a good enough reason to override his own denial that he saw them. In the same way, a man's screaming and writhing behavior provides us with a very good reason for thinking that he is in pain, even though it does not follow from the fact that he behaves thus that he is in pain. We

have, with the account of perception given, an explanation of why his behavior should give us a good reason for saying that he saw the branches. More importantly, however, we should note the fact that in order to show that it did not follow from his behavior that he saw the branches, we redescribed his behavior - he didn't duck in order to avoid the branches, but instead he had muscle spasms. The fact that the muscle spasms count against saying that he must have seen them, or undermines our claim that his behavior indicates that he must have seen them, is precisely because the muscle spasms are not discriminatory behavior like avoiding the branches. On the other hand, we can construct cases in which the person does perform discriminatory behavior - such as ducking to avoid the branches - even though he did not perceive the branches. We may suppose that he is actually blind, and is being guided in his movements by a sighted friend who is watching his path. If the guide says "There is a low hanging branch immediately ahead", the person ducks to avoid hitting it. In this case he fails to see the branches not because the behavior was not really discriminatory behavior having the branches as its object, but because the action was not made possible for him by virtue of his reaching a certain position as a result of performing actions involving in an essential way the use and positioning of the eyes. Nor, of course, could he avoid the branches because he heard them - even though the use of his ears for listening to his guide's instructions

was part of what made it possible for him to avoid the branches. He had to exercise his judgement, and rely on his knowledge of the language, since he could have heard the instructions and not have understood a word of the language in which they were given - it was not, then, solely by virtue of the position he was in as a result of the use of his ears that he was able to avoid the trees. It was, however, by virtue of that position that he was able to come to know what his guide said, and so to follow his directions. He heard his guide and not the trees.

The dispositional account of pains does not emphasize the pain behavior to the extent of disregarding the 'subjective' experience of feeling the pain or being in pain, since that experience is captured in the talk about dispositions - to say that my foot hurts is not to say that I behave in a certain way, but to say that I am disposed to do so. It takes, in many cases, a great deal of effort not to behave in those ways. Likewise, the account of seeing does not emphasize discriminatory behavior to the point of denying that there is a subjective side to seeing something. The subjective aspect is meant to be captured in the notion of reaching a position sufficiently good for certain kinds of behavior. To say that the blind man simple does not have certain subjective experiences which those with sight have is, in part (but only in part), to say that the performance of actions involving the use of his eyes will not put him in a position adequate for discriminatory behavior. I say

"in part", because not only does the blind man not see things, but as a result of his not seeing things, his total experience lacks other features which ours have - for example, nothing looks any way at all to him. The total experience of a blind person must differ in a number of ways from the total experience of a sighted person, so in saying that the use of a person's eyes cannot, if he is blind, put him in a position for discriminatory behavior is meant to capture only one aspect of that total difference - the aspect of his not seeing things. Since we have already separated questions about seeing things from questions about the differential responses that one's eyes, if functioning, might have to the things which one sees, we need not suppose that the advantage of sight only makes a difference as to whether things are seen.

In the previous three chapters we developed an account of the achievements of acquiring or reaching propositional states. To acquire knowledge is to manage an achievement, and to manage such achievements we must exercise certain skills and employ our linguistic resources. The account of perception proposed in this chapter is nonepistemic or nonpropositional, since to see something neither involves nor presupposes the employment of a language, nor does it necessarily result in our coming to know or believe something.

It seems highly implausible to suggest that there is no epistemological relationship of any interest between seeing

something and coming to know something about it - that is, no cases in which we acquire a bit of knowledge which we could not have acquired in that case and in that way if we had failed to see just what we did. A major test for any theory of perception must be that it succeed in accounting for the role of perception in that acquisition of knowledge. If to see were, a fortiori, to reach some propositional state, then the transition between seeing and knowing would be relatively smooth. But the task appears, prima facie, more difficult for a nonepistemic theory of perception.

Our next and final task is to show that there is a significant relationship between perception and knowledge, and to say just what that relationship is. First, however, it will be necessary to offer some criticisms of a deeply entrenched alternative account of the relationship between perception and justified belief. The defects of that account leaves a lacuna which, I hope to show, our account fills adequately.

Notes For Chapter Four

1. Dretske, Fred I. Seeing and Knowing, The University of Chicago Press, Chicago, 1969, p. 7-8.

Dretske provides, in the second chapter of his book, the most extended defense I have found of the thesis that seeing does not essentially involve having or coming to have any particular belief - that is, that

seeing is nonepistemic. He claims for his conclusion, however, only that seeing something does not essentially involve having any particular belief, and leaves open the possibility that seeing does essentially involve having some belief or another (see footnote 2, p. 17). He appears to qualify his account in this manner only because he wishes to avoid confrontation with the view that for anything to count as a sentient being, and hence to count as a perceiver, it must have some belief or other. Since I have, in effect, already confronted and rejected that view in the previous chapter, I will make no such qualification of my claim that perception is nonepistemic or nonpropositional. I find nothing incoherent about the view that many sentient creatures, and hence many perceivers, have no beliefs at all.

2. Ibid., p. 9.

3. The private language argument of Wittgenstein, in his Philosophical Investigations, seems to me to be conclusive in showing that there is conflict between holding both that our language is public and conventional and that all the objects of perception are logically private. The argument is, however, disputed in all of its forms, so I will not venture to rehearse and defend that argument here.

4. Versions of this argument are used by Descartes, Berkeley, Russell, Moore and Ayer, to mention only a prominent few. C.W.K. Mundle, in his book Perception: Facts and Theories, Oxford University Press, London, 1971, takes this particular fallacy to be the major source of the sense datum theory.
5. Russell, Bertrand. Our Knowledge of the External World, George Allen & Unwin, Ltd., London, 1914, p. 84.
6. Winston, H.F. Barnes, in his essay "The Myth of Sense Data", from Swartz, Robert J., Perceiving, Sensing, and Knowing, Doubleday Anchor, 1965, calls this argument the "sense datum fallacy" and takes it to be the major source of the sense datum theory.
7. Ayer, A.J. The Foundations of Empirical Knowledge, Macmillan Press, London, 1940, p. 4-5.
8. Cook, John W., "Wittgenstein on Privacy", in George Pitcher (ed.), Wittgenstein: The Philosophical Investigations, Doubleday & Co., Garden City, 1966, p. 298.
9. This is not meant to be an account of hallucinations in general, only an account of the mistaken beliefs that may, but need not, accompany hallucinations. In the

example just given, the mistake is one of failing to distinguish seeing a dagger from merely thinking that one sees a dagger. I do not suppose that all hallucinations involve a similar mistake or, for that matter, any mistake at all. First, I may, while hallucinating, know full well that this is what is happening, and so not suppose that I actually see the pink elephant. Secondly, no account of hallucination could be acceptable without strong compensatory arguments which did not allow for the possibility that nonlinguistic creatures, and so creatures with no knowledge or beliefs, might have hallucinations, since there is good reason to believe that experimental animals under the influence of hallucinogenic drugs have hallucinations. This is not the place to work out an account of the concept of hallucination, but there is one point that I would like to make. It is well known how readily suggestions, prior fears or desires, and such produce specific hallucinations in individuals or even groups that are in the appropriate states - say under the influence of hallucinogens. An account of hallucinations is, I think, suggested by the ambiguity already noted in the expression "P thinks that he sees a dagger", which is compatible with the claim that having an hallucination does not require being in any propositional state, and which helps to make sense of the role of prior suggestions, fears or desires in the specific nature of the

hallucinations that a creature has. An hallucination is consequent, I suggest, upon an impairment or deterioration (whether due to drugs, extreme fatigue, fever, etc.) in one's ability or capacity for differentiating between radically different mental states with a given (intentional) object. It is not, that is, either a confusion of objects (hallucinatory daggers for real daggers), nor a confusion of any one thing for anything else, but a loss of one's natural capacity or ability for differentiating distinguishable mental states. Thus, overwhelming horror at acts I have just performed with a dagger might undermine my ability to distinguish thinking about a dagger from seeing a dagger; the deteriorating effects of an hallucinogenic drug might transform my fear of bears into what I might respond to as a glimpse of a bear in the doorway. But any creature which has fears, wishes or desires could have sufficient mental resources for hallucinating, even though the same creature has no knowledge or opinions.

10. The conclusion I am defending here is argued for in greater detail in Richard Bosley's On Truth: A Neo-Pragmatist Treatise in Logic, Metaphysics and Epistemology, especially Parts I and II. My formulation of these arguments was influenced by an earlier draft of that book.

11. Locke, John. An Essay Concerning Human Understanding,
Book III, Ch. III, p. 6.
12. This short dialogue was related by Herman Tennessen.
13. Grice, H.P., "The Causal Theory of Perception", in R.J.
Swartz, op. cit., p. 463.
14. Dretske, op. cit., p. 23.
15. Ibid, p. 20.
16. Ibid, p. 20.
17. Ibid, p. 20.
18. Ibid, p. 24.
19. Ibid, p. 25.
20. Ibid, p. 25.
21. Ibid, p. 25.
22. Ibid, p. 26.
23. Ibid, p. 27.

CHAPTER FIVE

PERCEPTION AND JUSTIFICATION

A. Perception and Judgement

What, given the accounts of judgement and perception argued for in previous chapters, is the relationship between judgement and perception? Put in just this way, the question is as problematic as any answer that might be attempted. In this section I will answer the question after distinguishing between a number of different relationships which might be taken to be relationships between judgement and perception, but I will do so in a way which steers clear of any of the interesting epistemological relationships which might be proposed by one or another philosopher. I will discuss those relationships in a way which is independent of whether or not the judgements arrived at are justified. The problem of the justificatory relations between perception and judgement will be discussed in the remainder of the chapter.

To be clear about the possible relations between perception and judgement we can begin by keeping before us the distinction between the acts and the achievements relevant to each. On the side of judgement we have a distinction between the (complex) achievements of saying or thinking of something what it is, what color it is, etc., and the acts

of uttering or thinking which we subordinate to those achievements. I will refer to the terms of this distinction as, respectively, achievements of judgement and acts of judging. On the side of perception we also have a distinction between the achievements of seeing something, hearing it, etc., and the acts of looking, focusing, attending, etc. I will refer to the terms of this distinction as, respectively, perceptual achievements and perceptual acts. These pairs of distinctions provide us with the relata for the relationships to be considered.

In addition to the distinction between the possible relata, we can usefully distinguish four possible relations which might obtain between them. We can ask, about any two of the four relata distinguished, whether one is logically necessary or logically sufficient for the other, and whether the one is causally necessary or causally sufficient for the other. In the accounts of perception and judgement worked out above, I discussed the relationships which obtain between acts of perception and achievements of perception. My present concern is with the relationships, if any, between perception and judgement. This narrows the possible pairs of relata to four:

- a) Perceptual acts and acts of judging.
- b) Perceptual acts and achievements of judgement.
- c) Perceptual achievements and acts of judging.
- d) Perceptual achievements and achievements of judgement.

With respect to each of these pairs of relata, we can ask if either term is logically or causally sufficient for the other, and if either is logically or causally necessary for the other. With four unordered pairs of relata, or eight ordered pairs of relata and four possible relations to be asked about for each of the ordered pairs, we have thirty-two possible relations to be asked about. I do not doubt but that we could find for each of the thirty-two relationships some philosopher or psychologist who has held that the relationship does, in fact, obtain. The accounts of perception and judgement in the previous chapters force us to say, however, that only four of the relationships obtain - and the 'weakest' four at that, with two of the four obtaining only trivially. We can see this by considering two cases, each of which lies at an extreme opposite to the other. These are the cases of the nonlinguistic creature with full perceptual faculties, and the case of the linguistic creature with no perceptual faculties - the 'brute' case, and the ideal case of sensory deprivation.

On the account of judgement worked out above, I am committed to saying that no creature without a language or without mastery of some part of a language, can have any knowledge or beliefs. But the account of perception allows that such creatures may nonetheless manage the full range of perceptual achievements. The joint commitment of these accounts is, then, to the conclusion that acts of judging (since these are linguistic acts) are neither logically nor

causally necessary for acts or achievements of perception; achievements of judgement are neither logically nor causally necessary for either acts or achievements of perception; achievements of perception (and, a fortiori, perceptual acts) are neither logically nor causally sufficient for either acts or achievements on the side of judgement. The upshot of these conclusions is that no change in one's body of beliefs or in one's knowledge at a time is necessary for a change in one's perceptual position at that time, and that no change in one's perceptual position at a time is sufficient for any change in one's knowledge or beliefs at that time. (This last point must be understood as making no commitment to the question of whether a change in one's perceptual position is sufficient for any change in the extent to which some of one's beliefs are justified; what does not change, solely as a result of changes in one's perceptual position, is what one believes or knows).

There are, in short, no inferences possible from the fact that someone perceives something to the fact that he either knows or believes something. I will, henceforth, refer to the fact that no such inference is possible by saying that perception is nonepistemic, and to the view that such an inference is possible by saying that perception is epistemic. That perception is not epistemic ought to be evident from the fact that propositions asserting that someone perceives something (having the form "P sees X", "P hears X", etc.) and propositions reporting that

someone knows or believes something have a radically different form. The context of a report of a perceptual achievement which says what is perceived, and the context of belief reports which say what is believed, differ in important ways. What is seen, for example, is reported by a noun or noun-phrase; but what is believed is reported by a that-clause, a factive nominal. The latter reports, in part, an act of predication. The resulting difference between the two sorts of reports is not, then, merely a grammatical difference. Reports of the things we believe can never be inferred from reports of the things we perceive because no act of predication can be inferred from any perceptual achievements. The difference shows itself conspicuously in the fact that belief reports will always have contexts which are opaque to substitution, but reports of perceptual achievements never will. When perceptual verbs are sometimes followed by factive nominals (e.g., "P sees that ...") there is no report given of what is perceived, except occasionally and then only accidentally. I may see that George is dead when I see George, but I may also see that he is dead when I see only the obituary column, or see only the remains of his car. Likewise, I can see that it is raining, but in saying that I make no commitment whatsoever to the referent (if it should be thought to have one) of "it". Finally, not only is there no inference possible from the fact that I perceive something to what I believe, but there is no inference possible

to the proposition that I believe something at all. The permissibility of substitution in perceptual contexts but not in belief contexts shows this; no perceptual achievement or set of perceptual achievements will yield an act of predication. So, linguistic acts and achievements are not necessary for perceptual acts and achievements, and perceptual achievements are possible without linguistic acts or achievements.

The second case, mentioned earlier, for seeing possible relationships between perception and judgement is the ideal case of sensory deprivation. This is a case in which a person with full linguistic capabilities, who has mastered and uses a language fluently, and who has acquired a complex array of beliefs, is placed in a perfect sensory deprivation tank. This tank is operated in part by incapacitating with a curare-type drug the sensory organs of the subject. Not only, then, is the subject deprived of sensory stimulation, but he cannot perform acts like turning his head, focusing his eyes, extending his arms or fingers, or cocking his head to the side to listen. In this situation the subject may, upon considering his plight, come to believe any number of things which he has never before believed - for example, he may come to believe that he is dead, or that he has no arms, no hands, etc. This is not unlike the situation Descartes thought he had reached by the end of the first meditation. Finally, this person will also be under the influence of an additional

drug which makes it impossible for him to attend to or concentrate on anything for more than a few brief seconds - he cannot perform any of the acts we would normally think of as reasoning - he cannot draw any inferences. Nonetheless, he still acquires beliefs. He may, for example, think that he has gone mad.

The accounts of perception and judgement in previous chapters allow us to say of this case that the subject has no perceptual experiences whatsoever, but that he nonetheless acquires beliefs. So, having a perceptual experience - performing perceptual acts and managing perceptual achievements at a time - cannot be necessary (either logically or causally) for acquiring a belief - that is, it cannot be necessary for judging or for achievements of judgement at that time. Perception, at a time, is not necessary for judgement at that time. Furthermore, neither judging nor managing achievements of judgement are either causally or logically sufficient for managing a perceptual achievement or for performing a perceptual act. To this extent we may say the use of the faculty of judgement is independent of the use of the faculty of perception.

I have qualified the independence of judgement from perception by saying that having a perceptual experience at a time is not necessary for judgement at that time. The qualification is necessary since, as has been argued in the preceding chapters, making judgements depends upon having a language, and having a language depends upon perceptual

achievements (those involved in successfully learning the use of some set of expressions). So there is dependence, albeit indirect, of judgement on perception. Managing perceptual achievements at some time or other is necessary for performing acts of judgement and, therefore, for managing achievements of judgement at some later time. Perception is necessary for acquiring the resources we employ when we make a judgement; but it is not necessary for later employing those resources. Just how weak this relationship between perception and judgement it can be seen by noting that in this sense in which perception is necessary for belief, what we perceive leaves what we later come to believe completely indeterminate. If I learn the use of the word "cup" in ways which involve perceiving a cup, I will then have acquired an essential resource for later coming to believe of something that it is a cup, and also an essential resource for later coming to believe of something that it is not a cup; furthermore, nothing is contributed to what I will later come to believe that of. I may come to believe of a book that it is a cup and I may come to believe of a cup that it is not a cup, and the past perceptual achievements, on which my ability to acquire either of these beliefs depends, could be the same ones. It would not in principle be possible to predict or infer, from all my past perceptual experiences, what my beliefs are.

I said above that of the thirty-two relations possible between perception and judgement, only four obtain. They are all consequences of understanding the sensory deprivation case in light of the theories of perception and of judgement in previous chapters. The four relations are:

1. Perceptual acts are necessary for acts of judging.
2. Perceptual acts are necessary for achievements of judgement.
3. Perceptual achievements are necessary for acts of judging.
4. Perceptual achievements are necessary for achievements of judgement.

The sense of "necessary" in each of 1 - 4 is that the first relatum is causally necessary for the second relatum. There are four relations rather than only two because of the fact that perceptual achievements depend on perceptual acts, and achievements of judgement depend on acts of judging, so both acts and achievements of judgement depend on (remotely) acts as well as achievements of perception. We can summarize these relations by saying that there is only a remote, 'historical' relationship between perception and judgement.

Although I have illustrated the tenuousness of the relationship between perception and judgement using the extreme cases of the 'brutes' and the subject of perfect sensory deprivation, the contribution of perception to

judgement and of judgement to perception in our own cases is no different, since we are, as of yet, not concerned with whether or not what we believe is justified. For us the business of perception also proceeds without engaging the mechanisms of judgement, and the range of judgements that it is possible for us to make only remotely depend on our past perceptual experiences. But unlike the brutes, we do have the resources for acquiring beliefs, and unlike the subject of sensory deprivation, what we actually will come to believe and what we are justified in believing at the present time is in some interesting way a function of our perceptual position. Those relationships remain to be understood.

We are left with a rather sharp distinction, and only the most remote of connections, between perception and judgement. We may bring out the most general formulation of the distinction by introducing an expression which contrasts the total set of results of one's acts of judgement with the total set of results of one's acts of positioning and using the sensory organs. I will refer to the total set of what is believed (justifiably or unjustifiably) by an agent at a time, together with the total set of what is known by the agent at that time, as the agent's epistemic position at that time. Acquiring knowledge or beliefs, changing beliefs, and ceasing to know or believe something, will all count as changes in one's epistemic position.

Because a propositional state is an intensional state, and reports of such states fail to permit free substitution of coextensive terms, and because failure of substitution in reports of a state results from that state being such that reaching it requires a use of language, and because a linguistic act is a conventional act, it follows that no new propositional states can be reached without the performance of a conventional act of language. Because no new propositional state can be reached without the performance of a conventional act of language, it follows that one's epistemic position cannot be improved (i.e., no new knowledge or beliefs can be acquired) without acts of uttering or of recalling words.

Changes in our perceptual position, on the other hand, are brought about by natural acts and events - bodily movements, changes in the lighting, etc. - all of which are possible independently of the existence of any conventions in any community.

In short, the distinction between our epistemic position and our perceptual position is drawn along the lines of the ancient distinction between convention and nature.

Accretions to our epistemic position depend upon acts of judgement - acts of recalling words or expressions of a language. No change or set of changes in one's perceptual position is either necessary or sufficient for producing such a change in one's epistemic position. Likewise, no acts of judging and no achievements of judgement are either

necessary or sufficient for producing a change in our perceptual position. Our epistemic and perceptual positions are independent, each of the other. The remote, historical relationship between our epistemic and perceptual positions is only that learning the linguistic acts necessary for changing our epistemic position depends upon our being, at some time, in a perceptual position good enough for imitating the conventional practices of a community.

But now, it will be objected, this relationship between perception and judgement is far too weak to support our body of empirical knowledge. How, if judgement operates in the dark independently of perception, and if perception cannot move the gears of judgement, do we avoid the outcome that our body of judgements have nothing to do with the world? My response to this worry takes two parts. First, perception has already played an essential role in explaining how it is that our judgements are about the world; I will develop this point in the next section. Second, the extent to which perception plays a role in how it is that our judgements about the world are justified has been greatly exaggerated; I develop this point in the remainder of the chapter.

B. Perception and Judgements About the World

The weak relationship between perception and judgement which I have argued obtains is sufficient to explain how it

is that, in one sense, our judgements 'have something to do with the world'.

If an object, X, is individuated for an initiate into the language by acts of ostension, then the act of pointing is successfully subordinated to getting the initiate to see X. He is then in the relevant relationship with X to go on, in future cases, to individuate it for others by speaking, providing that he masters the linguistic skills which are then taught him for doing that. By virtue of having seen X he has discriminated it from other items in his perceptual environment and is in a position to take discriminatory action towards it. One form his discriminatory action might take is that he might say what it is or what color it is. This act of uttering the expression, say, "This is red", is a discriminatory act directed at X and not some other thing, because (and only because in this case) X is the object which he sees and at which he points. There are, in this case, two discriminatory acts - he points at X and he says something about X. Likewise, the initiate may learn the name "S" of X and say "X is red"; X is then the item about which he speaks (i.e., towards which he performs the discriminatory act of saying what color it is) because his use of "S" is subordinated, by the conventions of the language, to the achievement of individuating X. In the first case, his complex achievement of saying what color X is contains, as a subordinate achievement, the simpler achievement of seeing X or of bringing his listener

to see X. In the second case, where he utters the name but does not point, he individuates X. In both cases, when we report what it was that he said (fact sense) we must manage the subordinate achievement of individuating X. We may succeed in managing this achievement by saying "He said that X is red", thus getting another to know what he said is red by saying "X"; or we may succeed by saying "He said that this is red" (while pointing), thus getting another to see what he said is red. In either case we subordinate an act, either of language or of ostension, to individuating X whenever we give a report of what he said. There is, then, a clear sense in which what we report him as having said 'has something to do with the world' - namely, we attribute to him a judgement about X, one of the things in the world.

On the account of the role of ostension in setting and teaching precedents for the use of words in a language (Chapter Two, Section A), and on the account of the use of language in managing achievements of individuation (Chapter Three, Section F), this relationship between judgements and the world is one which is adequately accommodated by the weak, historical relationship between perception and judgement. That our talk, and so our beliefs, can be about the world is explained by that remote dependence of judgement on perception.

The relationship which is characterized by saying that our judgements are often about the world is explained by citing the remote dependence of judgement on perception.

My judgement that a cup is white, which I express by saying to you "This is white", is a judgement about the cup and not something else which I individuated for you. I individuate it for you, and not something else, by performing acts of language and ostension which brought you to see it, and not something else. My judgement that George is white, which I express by saying "George is white", is a judgement about George, and not someone else, because the conventions of the language in which George is the name of a particular man subordinate the utterance of "George" to individuating George, and not someone else. In the standard case, the utterance of "George" is conventionally subordinated to individuating him, and not someone else, and its use for that purpose is taught to others, by ostensively indicating George, and not someone else, to speakers of the language. The success of this operation consists in getting others to see George.

Just as the dependence of judgement on perception is a very weak and indirect dependence, so the relationship between our perceptual and epistemic positions which is explained by it is a very weak relationship. It ~~is~~ weak because it only contributed to an understanding of how it is possible for our statements and beliefs to be about the world; it is consistent with supposing that, even though we make statements and have beliefs which are about the world, all of what we say and believe about the world may well be false or unjustified.

The accounts offered of both perception and individuation are such that, as I have emphasized above, both what we see and what we succeed in individuating are indifferent to what we then go on to say about them. My success in individuating X by saying "X is red" does not depend on anyone having any true (or justified) beliefs about X, and seeing X does not depend on anyone's having any beliefs whatsoever. So, for all we are guaranteed by the role which perception plays in making it possible for our judgments to be about the world, we could speak endlessly about the world but say nothing which we are justified in taking to be true. We might fix our topic even though all of our comments on it go far astray.

It will, then, be objected that the relationship between judgement and perception which I allow is too weak to sustain any confidence that our beliefs about the world are justified; it is too weak to explain the role which, the objection would continue, perception must surely play in accounting for the fact that some of those beliefs are justified.

I will respond to this objection in the remainder of this chapter.

C. The Inquirer and the Apple-Picker

The epistemological problem of perception has been, traditionally, to say what the role of perception is in

coming to know or to justifiably believe something. The groundwork I have laid thus far separates rather drastically perceiving something and coming to know or believe (justifiably or otherwise) something. Acquiring a belief is a matter of employing a language, perceiving an object is a matter of employing the mute organs of sense, and nowhere, it would seem, is any room for interaction between seeing something and what we come to believe about it. Even to say, as I have suggested we ought, that judgements about objects of perception are discriminatory linguistic acts and that perception is a special kind of opportunity for discriminatory action, cannot lay a bridge between perception and justified belief. A natural supposition seems to be that seeing something, or seeing it in the relevant circumstances, makes some contribution to the status of the judgement one makes about the object. A person who does not see an object may not be justified in saying that it is red, whereas a person who differs only in that he does see the object may be justified in saying that it is red. The contribution of perception appears, in such cases, to be that seeing an object confers justification or warrant on the judgement, and that contribution is what marks the difference between a wild guess or idle speculation and a justified belief. But if what I have argued in the previous two sections is correct, then the role of perception is entirely explanatory - and it plays exactly the same role in the acquisition of an unjustified belief

as it plays in the acquisition of a justified belief. I have found for it no justificatory role.

To locate, in a general way, just what such a role for perception might be, it will be useful to begin by placing perception and justified belief in the context of the general activity which epistemology means to investigate - inquiry. First, the activity of inquiring may itself be subordinated to a wide variety of higher objectives - ranging from finding a lost article of clothing, to testing a scientific hypothesis, or to satisfying someone's curiosity. We may give up an inquiry simply because we tire of it or because we despair of ever succeeding, but we may also advance an inquiry without completing it, or we may bring an inquiry to some satisfactory conclusion. Although there are objectives dominant to the activity of inquiring which may or may not be reached, no inquiry can be counted a success unless it reaches, as an objective dominant to the activity of inquiring (though, perhaps, subordinate to finding the lost article of clothing or satisfying curiosity), the intermediate objective of acquiring knowledge or, as least, justified belief. The objective of inquiring is not to acquire beliefs, but to acquire beliefs which we are justified in holding to be true. Within the context provided by the previous chapter, the question about how perception contributes to our reaching this objective ought to be viewed as of a kind with the question about how

having an opportunity contributes to the success of any enterprise.

Suppose that, because of the height of apple trees in the region, no one has an opportunity to acquire an apple unless he is standing on a ladder. There are, of course, many rungs or positions on a ladder, but some of them provide, because of the positioning of the ladder, a better opportunity for acquiring apples than others. Suppose, furthermore, that one of those positions is such that no climbing up or down the ladder and no repositioning of the ladder itself could provide one with a better opportunity for acquiring an apple than one would have by virtue of standing in that particular position. As standing on the ladder provides an agent with an opportunity for acquiring an apple, so perception provides an opportunity for acquiring a justified belief. But in neither case would having a good position be sufficient for acquiring the apple or the justified belief. One must, in both cases, also perform certain acts - reaching for the apple, or judging what to say of the object of perception; and in neither case does having the opportunity guarantee success in reaching the objective when the acts are undertaken. Likewise, it is clear that in neither case would performing those acts guarantee success in getting the desired results when one is not in the relevant position (i.e., when one does not have the relevant opportunities for success). What is needed to sustain any hope of success is the convergence of

both the relevant skills and opportunities. Furthermore, in neither case would the skills and opportunities together guarantee the success of the enterprise undertaken, though their convergence would make failure highly unlikely and in need of some special explanation.

When an apple-picker is placed in the best possible position for acquiring an apple, and yet fails, we must account for his failure by discovering what went wrong with his reaching and grasping. His failure must be, assuming no problems with his opportunity, a failure to perform those acts skillfully enough. But what is the relevant skill which an inquirer might fail to have or exercise, the relevant acts which he must perform to acquire knowledge or belief? I have argued that an inquirer must be able to judge of, for example, the item seen, what color it is - that is, he must perform the act of recalling the word which it is a convention in his community to utter in subordination to saying what color something is. This act is a discriminatory act, in the sense that he must recall one word rather than another for saying of that object which he sees what color it is. It is, then, an act of the kind for which perception provides an opportunity. To maintain the comparison with the apple-picker we must say that in a case where perception provides a good enough opportunity for advancing or concluding his inquiry, his failure to acquire a justified belief cannot, a fortiori, be due to a defect on the side of his opportunity; his

failure to acquire a justified belief is a failure on the side of his actions - a failure to recall the word which it is conventional to use in his community for saying what color the object is. But a failure on the side of actions consists, in this case, in a failure to recall the word which it is conventional to utter in a given community for some objective; maintaining the comparison with the apple-picker, a person's failure to acquire a justified belief, in a situation in which his perceptual position is as good as can be for doing that, reduces to a failure to recall the conventional expression for the object (or for the color of the object) about which he makes the judgement. But that would imply that, in cases where the inquirer's perceptual position provides a good enough opportunity for success, all that he must do successfully in order to acquire a justified belief about the object of perception is succeed in recalling the conventional expression for saying what it is (or what color it is).

The apple-picker, standing in the best possible position on the ladder, may fail to acquire apples because he lacks the skill to pick them, or he may have the skill to pick them but still carelessly bungle the job. No better positioning of the picker on the ladder, and no better positioning of the ladder itself, will be sufficient to ensure success. But likewise, however good our perceptual position is, it will never be sufficient to ensure that we acquire a justified belief. We may always misjudge the

color of the object of perception, either because we lack the skill to follow the convention for the use of color words, or because we fail to exercise the skill we have. The upshot of this comparison, if it is a workable one, is that in the epistemic evaluation of an inquirer's belief, we must (for perceptual beliefs or statements) evaluate both the speaker's opportunity for success, and his exercise of skills - that is, we must evaluate both his perceptual position and his act of judging. But if we find that his perceptual position provides him with an adequate opportunity for success, then we can conclude that his belief or statement is either justified or unjustified on the basis of an investigation into his act of judging. If he has recalled the word or expression which it is convention to use in his linguistic community, then we can infer that his belief is justified.

The result of the comparison looks immediately implausible. It might readily be granted that the syntactic or phonetic adequacy of a report could consist in its conformity to the syntactic or phonetic conventions of a linguistic community. But to suggest that there could be, in conformity to the conventional practices of a linguistic community, the criterion for the epistemic adequacy of a speaker's report seems an open invitation to epistemological relativism and skepticism. Given, however, the account of perception (as an opportunity) and the account of judging (as the exercise of a skill in using words), the consequence

ensures that the relationship between the inquirer and the apple-picker is not one of analogy - they are instances of the same generic phenomenon. I am, then, committed to defending the view that perception provides an opportunity for performing acts of judgement in which our epistemic successes and failures are successes and failures in conforming, in the use of words, with conventional linguistic practices, and that the tests for success and failure in our inquiries is just the test for a report's conformity, or lack of it, to linguistic conventions.

I will approach the defense of this thesis through a detailed consideration of the epistemological relevance of commonly acknowledged failures to conform with linguistic practices - namely, verbal mistakes. By doing this, I hope to remove the appearance of implausibility which the thesis just stated no doubt has, and the spectres of relativism and skepticism which hang around it.

D. Perception and Verbal Error

We do tend to assign a higher epistemic rank to claims and reports which are about a speaker's perceptual position - reports about what is seen or how something looks - than to other classes of claims or reports. There are two very different ways in which our high epistemic regard for perceptual reports shows itself in our actual practices. Most conspicuously, it shows itself in our practice of

assenting to perceptual reports and dissenting from non-perceptual claims when the two come into conflict. But this respect for perception also manifests itself in our practice of occasionally excluding from the class of linguistic competent utterances verbal performances which are, phonetically and syntactically, perfectly good. In the most common cases, when a speaker has already proven himself competent in a language, and when he immediately corrects himself or is prepared to accept without hesitation another's corrections, we call the report or some part of the report a "verbal slip". But we even go so far as to exclude speakers from the class of linguistic competents when they consistently utter or intractably defend the most blatant falsehoods in the relevant perceptual circumstances. We charge them with misusing a word. This practice demonstrates our respect for the epistemological powers of perception in the following way: We would not be so quick to suppose that our friend makes a purely verbal error when, in clear sight of a lovely new house, he remarks "What a lovely new mouse!", if we didn't also suppose that his perceptual position makes it hard for him to be so very wrong. If his perceptual position were so bad that houses and mice looked the same to him from where he stood, the charge of verbal error would not be forthcoming. If we ignore the contribution made by perception to our confidence in reports and claims, we would be hard put to defend our distinction between knowledge or justified belief and even

the wildest of lucky guesses; but we would be equally hard put to defend our distinction between false opinions or unlucky guesses and verbal errors. The reverse side of our distinction between verbally incompetent utterances and the expression of highly unjustified assertions is the distinction between merely verbal competence and justified assertion. This distinction as well feeds on our respect for perception - an eloquent annunciation of "It is red" from a dark room is usually only eloquent; but in good perceptual circumstances the same sentence could falsify whole theories.

Standardly, philosophers have been content to assume, for the purposes of a theory of justification, that the perceptual claims up for investigation are made by linguistically competent speakers, and then have gone on to the business of trying to specify perceptual circumstances, or a special class of objects of perception, with respect to which reports are noninferentially justified or even incorrigible. My strategy will be to turn the tables and take for granted an adequate perceptual position and then to examine the interaction between the justification of a perceptual claim or report and the verbal competence of the utterance. Along this route there are gains to be made in understanding the role of perception in the acquisition of empirical knowledge.

The central notions of 'linguistic competence' and 'verbal error' will remain somewhat unclear until a later section of this chapter - they are part of what will be in

dispute, so to say too much now would be to beg the question against certain views to be considered. For now, we can say that the notions must be restricted in such a way that a speaker may be pronouncing all words correctly, uttering syntactically well-formed sentences which are meaningful in his own language, and still be making a verbal error or displaying linguistic incompetence. We have this when our friend says "What a lovely new mouse" in the earlier example, and we have a case of it when a speaker new to English says "I am electricity", when in fact he is an electrician.

The verbal errors and slips which have traditionally been used for epistemological mileage are all of this form - they involve the utterance of the wrong word, in a crucial position in a sentence, with the result that the sentence actually issued is perfectly good in all ways but one - taken at face value it is obviously false. I will refer to such utterances as displays of semantic linguistic incompetence. It is verbal error of this sort that Locke is writing about when he says:

If there ever happen any doubt about it [what our ideas are], it will always be found to be about the names, and not the ideas themselves.²

Or A.J. Ayer, when he says about reports of one's own sense data:

... all that one can properly mean in such a case by saying that one doubts whether this is green is that one is doubting whether 'green' is the correct word to use.³

And when he says later that, with respect to facts about our sense data:

To make an assertion that does not correspond to the fact that is either to tell a deliberate lie or else to make a verbal mistake.⁴

C.I. Lewis joins in indicating the same kind of mistake when he writes that:

There could be no doubt about the presented content of experience as such at the time when it is given, but it would be possible to tell lies about it.⁵

And then adds in a footnote:

Correlatively, it is always possible to make a mistake of formulation, even when there could be no possible error concerning what is formulated.⁶

W.V.O. Quine relies on the same notion of a verbal error when he constructs his behavioral account of observation sentences as a special subclass of sentences which would be elicited from any witness "conversant with the language"⁷ and makes it clear that when the relevant sentence is not elicited we are entitled to trim "our speech community to exclude our well-meaning but ill-spoken friend".⁸

The literature concerning perceptual reports, from which these samples are drawn, is laden with such caveats.

These remarks are all taken from philosophers who hold that, with respect to some class of objects of perception or some features of our perceptual position, we are infallible, or our reports are, when sincere, incorrigible.

Thus, when our reports are too obviously false, the error must fall on the side of giving the report, in the utterance. But it is important to notice that infallibilism and incorrigibilism are not necessary companions to the use of a special notion of linguistic ineptness. If we believe that we are fallible in formulating judgements about our perceptual position and that all our reports are corrigible, so that we must always be prepared to yield to conflicting but equally corrigible reports, we might nonetheless maintain that, in some perceptual positions, certain beliefs are so highly unwarranted that the patently unjustified statements of them must be written off as verbal errors. Even if our friend can't be certain, after a tour of the house, that it is a house, the belief that it is a mouse might be so strongly defeated or so unwarranted that it would take a slip of the tongue to praise it as a lovely mouse. In short, our use of notions of semantic linguistic incompetence doesn't depend on the view that we can't get certain things wrong; the view that we are very unlikely to get things wrong will do as well.

In the last chapter a notion of a perceptual position was introduced; the needs of the present chapter require shifting attention away from perception and focusing it on

the correct and incorrect use of language in various perceptual positions. The way to draw attention away from perception and focus it on verbal correctness of perceptual reports is to follow the example of the philosophers just quoted. The evaluation of a perceptual report is focused on its verbal aspect by taking the report to be given from a perceptual position which cannot itself be faulted for any errors that are made - presumably because it is so good that there is no way that it could be improved upon. I will now provide an account of a special class of perceptual positions which are not in need of improvement - which are, in light of the objective of following a convention for saying what color something is, adequate.

E. Adequate Perceptual Positions

It is an important feature of perceptual positions that with respect to specified objectives which a person or a community may have, a person's perceptual position may be better or worse, it may improve or deteriorate. We may be too far from an object to see it, the lighting may be too poor to come to know what color something is. By walking closer, turning our heads, squinting our eyes, or adjusting the lighting we may improve or worsen our perceptual position for reaching some objective we have set for ourselves. But only in light of some objective does it make sense to

speak of any change in one's perceptual position as either an improvement or a deterioration.

In what follows I will be talking about adequate perceptual positions - that is, perceptual positions which are good enough for reaching some objective. Many, but not all, alterations in such a position will count as deteriorations relative to the specified objective; some changes will leave the position adequate, and some will count as improvements. Turning out the lights will deteriorate most perceptual positions (there are exceptions - identifying the color of a neon light and observing the cloud belts of Jupiter require darkness) and brightening the lights will often improve a perceptual position - even if it is already adequate for some objective. But continuing to brighten the lights beyond a certain intensity does not count as an improvement and quickly begins to ruin the perceptual position. But because two different perceptual positions can both be good enough for some objective and yet one of them be better than the other, the 'weaker' of the two can be improved even though it is already adequate. There is, however, no reason to suppose that every perceptual position can be improved relative to a given objective. The activities of walking closer and brightening the lights will both cause a perceptual position to improve only up to some threshold; continuing, however, to perform these activities will cause the position to rapidly decay. To eliminate some otherwise irrelevant worries which might arise

later, I will narrow the class of adequate perceptual positions I am referring to to those perceptual positions which have already been improved, relative to a given objective, in such a way that no further modification of them will count as an improvement. These are, if you like, the best of all possible perceptual positions. My use of the expression "adequate perceptual position" in what follows will then refer to what is really only a subset of what are properly called adequate perceptual positions. Many have held that, with respect to the objective of making judgments about our sense data, our perceptual position is always adequate. We might, perhaps, be able to locate some objective for which every perceptual position is adequate in this sense.

It is when a speaker is in an adequate perceptual position that his false reports about what he sees or how it looks, barring insincerity and deliberate falsehoods, become serious contenders for verbal errors. When a speaker's perceptual position is beyond reproach, we see more clearly that his linguistic acts are not.

F. Adequate Perceptual Positions for Judging

Our habit of criticizing the competence of a linguistic performance given in certain perceptual positions suggests a way of specifying an important class of perceptual positions. Linguistic competence is tied to perceptual circum-

stances in the language building and learning situation. We acquire semantic linguistic competence by learning and mastering, to some extent, a vocabulary. Our perceptual position plays an uneliminable role in our initiation into a language, and in our acquiring the use of certain parts of the vocabulary (e.g., the color words). If our perceptual position is not good enough at some point to learn the vocabulary, then it is a truism to say that we will not, in that perceptual position, learn the vocabulary. But if a perceptual position is good enough to learn and master the use of a vocabulary, then that same position, or one relevantly similar⁹ to it in lighting, distance, etc., is good enough for us to correctly use that vocabulary again. This is not to say that we will use it, or that if we use it we will do so correctly; it is only to say that no modification of our perceptual position - of factors like lighting and distance - will improve our chances of correctly applying the vocabulary to the objects of perception. A perceptual position which is good enough for us to master a segment of a vocabulary is a fortiori good enough for us to be able to apply that vocabulary each time the position is regained. Any failure to use, for example, the correct word for saying what color something is, when we are in such circumstances, cannot be explained by citing a defect in the perceptual position.

It must not be supposed that we can ask how we can be sure that there is or ever has been a perceptual position

good enough for learning our vocabulary. We cannot always search for ways of improving our perceptual environment until it becomes adequate for teaching and learning a prior vocabulary. Our perceptual environment is given, in the sense that we first find ourselves in our various perceptual positions, and then we devise vocabularies to talk about what we perceive and how it appears; we do not begin with a vocabulary already in place and then seek out a perceptual environment in which we can find the distinctions which that vocabulary marks out. We simply devise a vocabulary which is adapted to the perceptual environment and its salient features as we find them. So, if there is any perceptual situation with enough constancy for a community to get a foothold in it, and devise a vocabulary to talk about features of interest and importance, then recapturing the same or relevantly similar situations will be to reach a perceptual position good enough for continuing to use the vocabulary. We need only consider the differences it would have made to our color vocabulary if we, and our color vocabulary, had evolved under a red star. Our color language would then have had no color term coextensive with our term "white", unless that color term had had a use given to it under 'peculiar' lighting conditions (i.e., under white light), much as expressions like "phosphorescent green" have come into our vocabulary through the use of ultraviolet light in mineralogy. It should, then, be evident that when the speaker is in a perceptual position

which includes seeing the object and the object's appearing some color to him, when the speaker's perceptual position is good enough for learning and mastering the use of the correct word for saying what color it is, then the utterance of the wrong word is not an error attributable to a defect in the speaker's perceptual position. The failure must be blamed on the verbal performance.

In evaluating a speaker's linguistic competence, we need do no more than determine whether his usage conforms with the conventional usage to which he is exposed in the language learning situations. The conventional usage to which he must conform is the use made of the expressions in the circumstances in which he learned the language. To conduct a fair test of the adequacy of a person's verbal skills, we must put him in the same (or relevantly similar) perceptual position as the one in which the usage is taught. It is, after all, the ability to use the words the same way as his teacher showed him which we are testing for. Only given the same or relevantly similar tasks and opportunities will disparities of performance indicate disparities in skill.

Perceptual circumstances are, then, adequate for uttering the correct (i.e., conventional) word for saying what color something is when they are adequate for learning the correct use of the color vocabulary. Perceptual circumstances which are not adequate for such ends are, a fortiori, not adequate for conducting a fair test of a speaker's

linguistic competence; and perceptual circumstances which are adequate for such ends are good enough for conducting a fair test of the speaker's linguistic competence. Consequently, a speaker's failure to utter the right word in such perceptual circumstances is correctly taken to indicate a defect in verbal skill or a failure in linguistic performance.

I conclude that the philosophers cited above, on verbal error, are correct to suppose that in certain perceptual circumstances a speaker's failure to say, for example, what color something is must be a verbal error. Where a speaker's perceptual position cannot be faulted, and so cannot be improved upon relative to the objective of saying what color something is, then the speaker's failure to utter the word which it is conventional to use for saying what color the object is, is nothing more than that - it is nothing more than a failure to utter the conventional word. It may be a slip which is quickly retracted, or it may be a display of a more serious incompetence in the use of that word, but in either case it is a mistake of a purely verbal sort.


We have found, then, the subset of adequate perceptual positions which are adequate for the particular objective of judging of something what it is or what color it is. They are those perceptual positions which are adequate for learning the use of the relevant expressions or class of expressions used in judging those things.

Since the important class of adequate perceptual positions for the purposes of this investigation are those which are adequate for judging correctly (i.e., those which are adequate for the use of an expression in accord with the conventions for its use), it will be helpful to summarize the account of those perceptual positions and to respond to two objections which might be raised to the account.

The account is developed as follows: If a speaker's perceptual position at a time is adequate for learning the use of some vocabulary, V, then that same position is adequate for using V again. Let us refer to this position as the "learning position". But a speaker's perceptual position at the time when the vocabulary is actually learned is, a fortiori, adequate for learning the use of V. It follows that that same position is adequate for using V when that position is reached again. The extension to other perceptual positions is then made in the following way: Any other perceptual position which is the same as, or relevantly similar to, the learning position is adequate for learning the use of V, since if it is the same as or relevantly similar to the learning position, then V could have been learned in that other perceptual position. But if any other position is such that V could have been learned in that position, then that position is such that the vocabulary actually acquired in the learning position can be used there with equal hope of success. So, it follows

that an adequate perceptual position for the use of an expression in accord with the conventions of the language is any position in which the use of the expression could have been learned, and that the use of the expression could have been learned in any perceptual position which is the same as or relevantly similar to the one in which it actually was learned.

It might be objected that this account avoids saying what an adequate position for using *V* is - the characterization of an adequate position for judging is entirely comparative. A position is said to be adequate because it is the same or relevantly similar to some other position, but nothing is said about what the features of that other position are such that it is adequate for learning, and so for using, *V*, and so nothing is said about what an adequate position for judging is except that it is one which is the same or relevantly similar to that earlier, but completely uncharacterized, position. My response to this is that the important characteristic of the learning position which makes it adequate has already been stated - namely, that the vocabulary was in fact learned in that position. Any other features which that position might have which are relevant to our learning the vocabulary in that position are, in a sense, a matter of accident, and must be expected to vary from one set of expressions to another. As I have argued above, we acquire a vocabulary which is adapted to the perceptual environment in which we happen to find



ourselves - whatever it may be like. It is a contingent matter that our vocabulary - for example, our color vocabulary - has as adequate conditions for its use those perceptual conditions which it does. If the perceptual conditions under which our color vocabulary evolved, and under which it is passed from one generation of speakers to another, had been very different, then we would have had a color vocabulary which marked out a very different set of distinctions than ours does, and so the perceptual conditions which would have been adequate for the correct use of our color vocabulary would have been very different than they are. So, any detailed characterization of the class of perceptual positions which are in fact adequate for the correct use of our color vocabulary, though of interest for other purposes, is indifferent to our present concerns in epistemology. All that need be said for present purposes is that they are those positions which are the same as, or relevantly similar to, the perceptual positions in which we actually do acquire our vocabulary - whatever those perceptual positions may be like.

A second objection which it may be tempting to raise is that the vague notion of 'relevant similarity' (which is essential for tying adequate perceptual positions for judging to learning positions) is being used to conceal the fact that a wide range of drastically dissimilar perceptual circumstances are good enough for learning a vocabulary, and those may have so little in common with each other that

the link I claim between learning positions and adequate positions may obtain only trivially. That would be a serious objection since the tie between learning positions and adequate perceptual positions is needed if citing the adequacy of a speaker's perceptual position is to help sustain the charge that his errors must be verbal. I will respond to this by specifying some similarities between two perceptual positions, which are relevant similarities, in the sense that if these similarities do not obtain, then learning a vocabulary in one of the positions will not allow us to say that the other position is adequate for employing that vocabulary.

Suppose that there are two conventional expressions in a color vocabulary which are subordinated by the conventions to saying of objects of two different colors what color they are. Then, I suggest, there are at least two necessary conditions which must be met by any perceptual positions which are adequate for learning the conventional use of those expressions:

- (1) In those perceptual positions things of the same color look the same in color; and
- (2) In those perceptual positions things of a different color look different in color.

I make no claim for these conditions being sufficient, but only for their being two important necessary conditions for learning this color vocabulary. There may be a very

wide range of perceptual circumstances which satisfy these conditions, but circumstances which do not satisfy either condition are so defective that the use of the color words cannot be taught and learned. We acquire, in circumstances which satisfy both conditions, the skill to apply the color words to new cases; if these conditions are not met, then we will fail to acquire that skill. Suppose, then, we proceed to a new perceptual position in which one of these conditions is not met. If (1) is not satisfied, then, in that new position, things of the same color will look different in color. Part of the skill we acquired in the learning position was to apply different words to things which look different in color. If we then proceed to do that, then we will proceed to apply different words to objects which, in the learning position, we learned the convention of applying the same word to. So, to use the skill learned in perceptual conditions which satisfied (1) and (2) in perceptual conditions which do not, is to fail to follow the convention for the use of those words. If (2) is not satisfied by our new perceptual position, then things of a different color will look the same in color. Part of the skill we acquire in the learning position is to apply the same words to things which look the same in color. If we then proceed to do that, then we will proceed to apply the same word to objects which, in the learning position, we learned the convention of applying different

words to. To do this is, again, to fail to follow the convention for the use of the word.

An explanation of why we fail to follow the conventions is given by citing the fact that the perceptual position we were in differed from the learning position in the ways specified. In the cases I have mentioned, the explanation of the failure to follow the convention does not have, as a part, any failure to have or exercise the skills we acquired when we learned the language. In fact, our having and exercising those skills provided part of the explanation of the failure to act in accord with the convention we learned to imitate in the language learning situation. The failure to accord with the convention cannot, therefore, be plausibly thought to be a sign of linguistic ineptness. The specification of the relevant similarities as given by (1) and (2) is, then, sufficient to sustain the claim that failure to follow communal linguistic practices in inadequate perceptual circumstances is not plausibly construed as 'merely verbal' failure.

It should be clear on the basis of this example that what kind of similarities will count as relevant similarities will vary from one set of words to another. The perceptual conditions which satisfy (1) and (2) will not be adequate for learning all of the vocabulary of a language, but only that part the learning of which depends on the color things look. What will count as relevant similarity between two perceptual positions will be a function of the

vocabulary whose application is meant to be carried over from one position to the other. There is no a priori reason to suppose that a completely general specification could be given of relevant similarities for the adequacy of perceptual positions for the use of an entire language.

Nonetheless, I would speculate that a generalization of (1) and (2), as follows, would cover a very extensive part of any vocabulary whose use is to be learned ostensively:

- (a) Things of the same kind appear the same; and
- (b) Things of a different kind appear different;

where what things are of the same or different kind is to be understood as being in part a function of the conventions of the community which we are to learn and apply in these perceptual conditions. If a prior vocabulary is learned with the support of ostensive acts, then if a perceptual position is adequate for learning that vocabulary then that position will meet conditions (a) and (b). Any other perceptual position which meets both conditions will be relevantly similar to the learning position, and so will be adequate for using that vocabulary. For different parts of the ostensively acquired vocabulary, these criteria for relevant similarity will have to be further restricted by saying in what respect things must appear the same or different. I think enough has been said at this stage to proceed with the use of the notion of relevant similarity without fear that the link it expresses between

learning conditions and adequate conditions is too flimsy to sustain the claim that charges of verbal error cannot be supported except when a speaker's failure to follow a convention occurs in adequate perceptual circumstances.

We have seen in this section that the relationship between a speaker's perceptual position at a time, and his act of judging at that time, which is expressed by saying that his perceptual position is adequate for judging at that time, is a relationship which obtains when and only when his perceptual position is the same or relevantly similar to the perceptual position in which he learned the language. So whether or not our present perceptual position is adequate for judging what to say in order to say what color something is depends only upon whether or not our position has that relationship to the 'remote' perceptual circumstances in which we learned the language we are using. In an earlier section I argued that the remote, historical, dependence of perception on judgment, which accounts for the ostensive learning of words, is sufficient to explain one sense in which our judgements 'have something to do with the world' - namely, the sense in which our judgements are about things in the world. That conclusion left open the possibility that our judgements, though about things in the world, may nonetheless all be false or unjustified. Although we may have things of the world as our topic, our comments on those topics may not be to the point. We have, now, one important piece in place to begin

to unfoot that worry. What we say to be the case about the world depends upon what expression we utter; what we think to be the case depends upon what expression we recall. Our achievements of judgement, in short, depend upon our acts of judging. But our acts of judging are acts learned in certain perceptual circumstances and, we can now add, our continuing to engage in that activity - in following those conventions - also depends upon maintaining continuity between the perceptual positions in which we judge and the perceptual positions in which we learned to judge. If we can only restore ourselves to perceptual circumstances of the sort in which we learned our language, then any failure to follow the convention for saying of something what color it is will be a verbal error. In the next section, we can proceed with the examination of the relationship between verbal errors and the justification of our beliefs in adequate perceptual circumstances.

G. Perception, Justification and Verbal Error

We have been examining the relationship between a speaker's perceptual position and his verbal competence with a vocabulary. We need now to formulate the relationship between the verbal competence of a speaker's utterance, when he gives a report in an adequate perceptual circumstance, and the epistemic status of the speaker's report. This relationship is the one displayed in our practice of

dismissing any blatantly false or highly unwarranted statement, made in adequate perceptual circumstances, as a verbal error. That we do engage in this practice is part of the substance of the view that perception contributes to the success of our inquiries, where those successes consist in the acquisition of justified beliefs. To give up the practice is to give up the distinction between unjustified beliefs and verbal errors, and the consequence of losing that distinction is that, then, even the best of all possible perceptual positions will make the same contribution, whatever it is, to justified beliefs and to highly unjustified beliefs. Maintaining the practice of dismissing certain claims as verbal errors is essential to the supposition that perfect perceptual circumstances make it impossible, or at least very difficult, for us to be drastically wrong about what we see.

We saw earlier, with examples extending from Locke through Quine, how it is that philosophers exclude certain utterances, or even speakers, from the class of linguistically competent utterances or speakers by laying charges of verbal error. The practice of excluding utterances as verbal errors, as I have described that practice and as these philosophers have engaged in it, is the practice of inferring from the fact that a report is made in the best possible perceptual circumstances, together with the fact that the report is unjustifiable, to the conclusion that the report embodies a verbal error. The practice can,

then, be formulated as an epistemological principle as follows:

P1: If a speaker's perceptual report, in a language L, is given in an adequate perceptual position, and that report is such that he is not justified in taking it to be true, then that report involves a verbal error in L.

By a 'perceptual report' I mean a report which purports to describe some feature of the perceptual position of the speaker giving the report. For example, "It is a cup" or "It is white", where the use of "It" is subordinated to individuating some object which the speaker sees. The relevant adequate perceptual position is that perceptual position which is adequate for the use of the vocabulary of L in which the report is given; more specifically, our present concerns require that the perceptual position be adequate for the use of those expressions which are uttered in subordination to an achievement of predication.

If my perceptual position includes seeing a red tomato, and includes the tomatoes looking red to me, and my position is perfectly adequate for learning, then and there, the use of the word "red", but I say "The tomato looks blue" (which there is no reason for believing to be true), then I made a verbal mistake when I uttered the word "blue". If my perceptual claim is justified (or I am justified in claiming that it looks blue), then no charge

of linguistic incompetence can be supported by the principle P1; if the perceptual position is not adequate, no charge of verbal incompetence can be supported. But when neither of these conditions is met, and we have an unjustifiable claim made in adequate perceptual conditions, then what has gone wrong, has gone wrong with the wording of the claim. When a speaker's perceptual position is beyond further improvement, when no adjustment of the lighting or change of distance from the object of perception and no change of relevant states of the speaker can contribute to an explanation of why he said that the tomato was blue, then we must conclude that he simply used the wrong word. If, on further questioning, the speaker continues to maintain that the tomato is blue, and we hold fixed the assumption that nothing will count as an improvement of his perceptual position, our only recourse is to reinstruct him in the use of the word "red".

There is a possible objection to the epistemological use of the notion of a semantic verbal mistake which it will be useful to have before us as we proceed. The objection would raise the following rhetorical questions: If we had enough philosophical distance from empiricism, then could we not take the class of semantic verbal errors to be anomalies for empiricist theories of knowledge - in fact, would we not take them to be counterexamples? If a person in the best of all possible perceptual positions can sincerely utter the most preposterous falsehoods, why not take

him as providing a counterexample to the thesis that our perceptual positions provide us with any amount of justification for our empirical claims? Why isn't the proper view of the special class of verbal mistakes that it is introduced in a completely ad hoc way, to preserve the thesis that perception justifies claims against blatant counterexamples? The use made of this special class of errors is to relegate any false or unjustified perceptual claim made in unimpeachable perceptual circumstances to the category of the trivial or uninteresting by labelling the error "merely verbal". But call the error what we will, we don't thereby eliminate counterinstances.

A first response to this complaint points to the fact that we have a clear category of verbal slips in which what a speaker says and what he knows or justifiably believes are very different. A competent speaker of the language will immediately retract his sentence or a word once it is pointed out that he has just called a house "a mouse" or a red sense datum "blue". But this response is inadequate. We still classify a blatantly false perceptual report as a verbal error even if the speaker won't retract it - providing we grant that his perceptual position is beyond reproach. With that much granted, even when a speaker calls all other red things "blue" and calls all other houses "mouse", we write off his error as verbal. So long as he is using the words wrong, it becomes useless to ask him what color he thinks the tomato is (as opposed to what

color he would say it is), since he will then express his sincere belief by saying that it is 'blue', and that he is completely justified in thinking that it is 'blue' because it clearly looks 'blue'. Holding to the stipulation that there is no defect in his perceptual position, we seem to be writing off his claims as verbal errors solely on the grounds that they are blatantly false and unjustified. To do this, we must extend the notion of a verbal mistake far beyond the class of verbal slips which embarrassed speakers retract on a moment's notice. Correcting verbal mistakes in this extended sense of a verbal mistake requires instructing speakers in the correct use of our words.

What seems right about the objection is that the extended sense of 'verbal error' in which the speaker is using his words incorrectly, does seem tied to the assumption that when our perceptual position is adequate, we can't get the facts so very wrong, so we must be getting the words wrong. (That is an inference we are entitled to by P1). But this extended notion only provides an ad hoc means of defending empiricism from clear counterexamples providing that the only criterion for a perceptual claim's being a verbal mistake or involving an incorrect use of a word is the claim's failing to be justified. If we can classify an utterance as a verbal mistake only on the grounds that it is false or unjustified, then our use of the extended notion of verbal mistakes does seem an ad hoc device for preserving a relationship between perceptual

position and justification, by allowing us to remove, a priori, possible counterexamples. The role perception plays in justifying our confidence in perceptual claims will remain obscure until this objection is answered. I will return to it in a later section.

H. Three Theories of justification

So much for the epistemological use of verbal error. But verbal error, it will understandably be said, is a side issue. Only if we assume that the relevant claims are made by competent speakers can we proceed directly to the important issue of how perception justifies those claims. As long as we allow that a speaker may be making an inept use of language, then we must allow that the claim he actually makes may, at face value, be one which he is not justified in making in his perceptual position - that is, there may, then, be no interesting justificatory relationship between what he actually says and the perceptual position he is in, however good that perceptual position is.

As indicated by P1, there are three distinct pairs of options which interact in discussing the justification of perceptual claims:

- 1) The person's perceptual position (which may be adequate or may fail to be adequate).
- 2) The person's claim, report, assertion, etc. (which may be justified or which may fail to be justified).

- 3) The person's use of language in making the claim, report, assertion, etc. (which may be correct or competent, or which may not be correct or competent).

It is in the relationship between (1) and (2) that the heart of the problem of empirical knowledge is thought to lie. If the person uses language incorrectly, in the sense of displaying semantic linguistic incompetence, then we could not expect the report actually issued (e.g., "The tomato looks blue to me") to have any epistemologically interesting relationship to the perceptual position the speaker is in. Hence, the need for the empiricist caveat, offered by Locke, Ayer, Lewis, Quine, and countless others, warning us that even on the supposition that someone is in the best of all possible perceptual positions we cannot be guaranteed a true, incorrigible, or even justified report from them. If there were no such thing as an incorrectly worded report - no semantic linguistic failures possible - these philosophers would be committed to saying that perceptual reports given in adequate perceptual circumstances are all true, incorrigible, or at least justified.

If the key to the justification of perceptual claims lies in the relationship between (1) and (2), it will be useful to see why traditional theories of justification fail to unlock that relationship. Recall from earlier that there is a second part to the 'empirical' or 'in actual practice' content of our high regard for reports brought

back from, or given in the face of, perceptual experience. That is our practice of giving assent to perceptual claims and dissenting from nonperceptual claims when the two come into conflict. How this practice is structured or organized is the subject of foundationalist and coherentist theories of justification. Although either theory can assign positions of special importance to perceptual statements to accommodate the fact that we assent to them while dissenting from other classes of statements, neither can offer a satisfying account of, or defense of, that practice.

Consider the foundationalist strategy: Statements owe whatever justification they have to epistemologically prior generations of statements, and that justification is passed down from one generation of statements to the next through logical relations, either deductive or inductive. To avoid the infinite regress which this invites, a class of basic statements is specified which, though justified, receive their own justification from no other statements. Hence, the need for either a class of self-justifying statements or, perhaps, a class of statements which are justified by their nonlogical relationship with items outside the class of statements. If self-justifying basic statements justify themselves in anything like the manner in which they justify other statements, circularity threatens. If basic statements are to be justified by their nonlogical liaisons with some class of nonpropositional items, then items

perceived seem the obvious candidates, and perceptual claims, or some subset thereof, seem to be the best bet for the class of basic statements. But foundationalism can go no further than this. Infinite regress and circularity may well be avoided by assigning the position of prestige to perceptual claims, but nothing is said to explain how it is that what is perceived, whether macroscopic public items or sensory data, or how the perceiving of these items delivers to the first generation of statements any amount of justification. In fact, the foundationist picture militates against understanding how that could happen, since it makes whatever sense it can of the transfer of justification down the generations in terms of the inferential relationships between statements, and those are relationships which statements cannot conduct with perceptual experience.

The coherentist strategy fares no better; this plan of attack also supposes that our statements owe whatever justification they have to logical relations, of weaker or stronger sorts, that they have with other statements, with the variation that all the interacting statements are, epistemically speaking, of the same generation. Whatever justification one statement hands to others, it eventually gets back from them. But this picture is thought to represent epistemic licentiousness; the product of the inbreeding will be a family of statements which conducts itself with familial propriety, but which lacks a sense of reality. Justification will have to be injected into the body of

statements from outside - and the likely bridge to the outside world is perception. But once again, the model requires us to understand the transfer of justification from one point to the next on the model of inference; and that model requires statements to pass and receive the justification.

The problem with both these traditional strategies, when it comes time to accommodate perception, is that our perceptual experience lies outside our body of statements, however that body is structured, and our perceptual claims stand on the inside. Both strategies suggest, quite reasonably, that the justification of a statement proceeds by providing other statements which are inferentially linked to the first, as premises are linked to a conclusion in an argument. The difficulty they share is that no item perceived, and no perceptual experience of an item, has any such link with a statement about the item, or with a statement about the experience of the item.

This failing of foundationist and coherentist theories of justification suggests the need to view perceptual justification as another kind of link between perception and justified claims - a nonlogical link. The preferred candidate for such a link is causation. A causal theory of the justification of perceptual claims maintains that perceptual claims, or some subset of them, are justified because they have the relevant causal aetiology. This theory has the virtue of being compatible with both tradi-

tional strategies, and so leaves them to iron out their differences between themselves; it has, however, the vice of running afoul of our epistemological tradition which has viewed itself, in Reichenbach's words, as setting out

the reasons that logically justify our beliefs and not a narrative of the causes that in fact led us to adopt them.¹¹

A causal theory of perceptual justification dares to run across Reichenbach's distinction. To justify perceptual claims by giving a "narrative of the causes that in fact led us to adopt them", a causal theory must be sensitive to the fact that we can tell a causal story (though not a story providing reasons) for every opinion we adopt, and many of those stories do not render the opinions more, rather than less, respectable. If a causal history of a belief or claim is cited to demonstrate that our confidence in it is justified, then we must believe the story, and our belief must be justified. But that only brings us back to statements justifying other statements. If a causal theory is to do what it sets out to do, then it must make a case for saying that perception is causally reliable in producing beliefs or claims which are more likely to be true than not. But whatever case it makes for that, it cannot depend on anyone's believing it and, furthermore, must be consistent with its being disbelieved. Unless it can meet this condition, the justification of perceptual claims will be made derivative of the causal story about the claim, and

not derivative of perception. Then the burden of justifying perceptual claims would be shifted back to the problem of how we justify the statements making up the causal story itself, and so back to a view in which perceptual statements are justified by other statements, rather than by perception.

Finally, if we take seriously the fact that, on the causal theory, the causal story does not provide the justification, and so need not be believed for perceptual claims to be justified, then the causal theory of justification becomes an easy target for counterexamples. Suppose, for example, that wishful thinking is in fact a reliable cause, but we falsely believe it to be unreliable. (Do we know that it isn't in fact reliable? That is, do we know that more false beliefs than true beliefs are produced by wishful thinking?) Then a belief, B, produced by wishful thinking is justified, even though we are not justified in believing B. We could not, that is, count B as justified and could not rationally accept B, even though it is justified. The result of such counterexamples is that the beliefs we can rationally count as justified bear no interesting relation to the beliefs which are, in fact, justified.

This consequence leads to the following concession from even the staunchest defender of the theory:

What we really want is an explanation of why we count, or would count, certain beliefs as justified and others as unjustified. Such an explanation must refer to our beliefs about reliability,

not to the actual facts. The reason we count beliefs as justified is that they are formed by what we believe to be reliable belief-forming processes. Our beliefs about which belief-forming processes are reliable may be erroneous, but that does not affect the adequacy of the explanation. Since we believe that wishful thinking is an unreliable belief-forming process, we regard beliefs formed by wishful thinking as unjustified. What matters, then, is what we believe about wishful thinking, not what is true (in the long run) about wishful thinking.¹²

To concede this much is only to concede that the burden of actually providing a justification for our beliefs must be carried by our other beliefs and statements; but it is not to concede that what we are actually justified in believing is a function of those other beliefs or statements. It will yet be a function of which belief producing mechanisms are in fact reliable. Thus, escaping the suggested counter-example does require the causal theorist to claim that he is only providing an account of which of our beliefs are justified and not an account of what we are, or are not, justified in believing to be true. This response to the counter-example will, however, bring the causal theory more worries than comforts.

The distinction between the belief that S and someone's believing that S is such that the belief that S may be shared by any number of persons, but one person's believing that S is separate and independent of any other person's believing it. Suppose a community in which everyone believes that my car is brown, where only one member of the community has seen the car (and he has not spoken to any of

the others). Since everyone has the same belief and, according to the causal theory of the justification of beliefs by perception, that belief is justified, it follows that everyone in the community holds a belief which is justified. Only one person in that community, however, is justified in believing what he does. In short, by using the distinction between a belief's being justified and someone's believing it being justified, the causal theory has lost any interesting correlation at all between the perceptual positions of persons and justification. The only interesting correlation between perception and justification which remains is the correlation between who sees something and who is justified in believing something. The developments which the causal theory must make to meet the proposed counterexamples have the consequence that being justified in believing that S is explained by citing the logical relationships which believing S has to believing other things - in particular, certain causal theories about the reliability or lack of reliability of specific belief-producing mechanisms. So our being justified in believing something is not, on this theory, something which gives a role to perception in its explanation. On the other hand, the counterexample does allow that the causal theory could say something about whether or not a belief is justified (regardless of whether anyone is justified in believing it), and that it could, in principle, give some role to perception as a reliable mechanism in accounting for whe

ther or not a belief is justified. But that option, though left open to the causal theory by the counterexample, is such that it can provide no correlation between what someone perceives and which of their beliefs are justified. I conclude, then, that a causal theory of justification can provide no account of the role which is played by perception in the justification of our perceptual beliefs.

We are left, then, with justification being passed from one statement to another, and with no clear view of how justification gets into the body of statements.

This completes my survey of traditional options and their difficulties. I do not want to suggest that any doubt has been cast on the usefulness of foundationalist and coherentist accounts of justification. One or the other may well be right, or perhaps each characterizes accurately different local areas in our body of beliefs. But it should be clear that disputes about the internal structure of a body of beliefs or statements are not relevant to the issue of the justification of perceptual claims. That is a problem about the role which perception plays in the justification of statements, whether it plays that role directly in the justification of individual statements or only in the justification of large, interconnected bodies of statements. By being told how perceptual claims interact with other statements, we are told nothing about the role of perception in the justification of those claims.

One statement is too much like another, and no statement is enough like perceptual experience.

Recalling why we turned to the traditional theories of justification, we see why it is futile to look there for help with our problem. One way in which the assignment of a high epistemic rank to perceptual statements shows itself is in our practice of assenting to those statements at the cost of rejecting other kinds of statements. But the practice of using one statement to argue against another, or even to argue for another, is business internal to the use of statements, not business which can be conducted between statements and perceptual experience. In fact, we ought to begin to wonder how there could be any dealings at all between perception and statements or beliefs - even dealings of a justificatory sort.

I. Firth: Verbal Errors and Idiolects

My proposal is that we go back for a closer look at the other practice we engage in out of confidence in perceptual claims - our practice of using an extended notion of a verbal mistake to exclude any false or unjustified claim from the class of competent utterances when the speaker has an unimpeachable perceptual position. This is the practice characterized in P1.

The reason for going back to it is that it suggests a theory of justification for perceptual reports quite unlike

the ones we've just looked at. It suggests a theory of justification in this way: Granting an adequate perceptual position, the use of principle P1 would allow us to infer from the premise that a speaker is not committing a verbal error, to the conclusion that his claim is justified; and it would allow us to infer from the premise that a speaker is committing a verbal error, to the conclusion that his claim is unjustified. That would indicate a possible route, through an examination of the notion of semantic linguistic competence, into a theory of the justification of perceptual claims. The proposal is, in short, that the epistemic confidence we place in perceptual reports might be made to rest on the linguistic competence of the speaker who gives the report. This suggestion should not be balked at on the grounds that perception has been left out of the picture; unless we assume that the speaker's perceptual position is adequate, we cannot level the charge of verbal error against a false or unjustified report. If a speaker's perceptual position is not adequate, and he means to say what color something is, then it is no display of linguistic ineptitude if he should utter the wrong word. His mistake can be explained in other obvious ways - he was too far away, the lighting was too dim, he is color blind, etc. But the point of principle P1 is that when no such explanation of his failure is possible, then we must write it off as a linguistic failure. But, then, if his utterance had

been verbally competent, the confidence we place in his claim would be justified.

There is, however, a relatively recent philosophical doctrine which would, if correct, make this kind of approach to a theory of justification impossible. We need to look at two examples.

The first example is an epistemological use of a principle of linguistic competence found in Roderick Firth's "Coherence, Certainty, and Epistemic Priority".¹³ There, Firth points out that when a child begins to use a word like 'red', "with any consistency", he applies it to all things that look red, whether or not those things are in fact red. This leads Firth to suggest that

It would not be unreasonable to assert that the child is using 'red' to express a primitive form of the concept 'looks red'. To call this a 'primitive form' of the concept 'looks red' is to acknowledge that in some sense the child cannot fully understand adult usage.¹⁴

As a point of background to this passage, we should keep in mind that in perceptual situations in which something looks red to us, we are justified in asserting that it is red. Our knowledge that nonred things sometimes look red bars us from claiming to be justified in thinking that something is red solely on the ground that it looks red. But this is only to say that we do not accept the inference from "X looks red" to "X is red" without the addition of further premises.

So, when Firth says that "the child cannot fully understand adult usage", he seems to be employing our first principle, P1. The children are often in an adequate perceptual position to say that something looks red to them - no improvement of their position is necessary; but they do not say that; they say that something is red, and that claim is, in their perceptual position, often unjustified. So we could conclude, using P1, that they are linguistically incompetent. In this case, the children's consistency presents us from writing these utterances off as mere verbal slips - so instead, the children are said to lack fullfledged membership in our linguistic community. If we simply suppose that the sentences the children utter when they say "It's red" means that something is red, when nothing is, and explain the falsity (and unjustifiability) as a result of linguistic incompetence. But Firth is committed by a theory of meaning to going further than that - further than P1 will explain.

Firth has suggested, in the passage quoted, the option of supposing that what the children's utterance expresses is true and, even, justified. The price he pays to do that is the consequence that what the children's utterances mean is not what an adult speaker's utterances of the same sentences would mean. This move requires an account of meaning under which what a speaker's use of the word "red" means depends upon the concept being expressed by the speaker. Such theories of meaning are notoriously fraught

with difficulties, but our present concern lies elsewhere. When the children use the word "red", they express a different concept than adults; via the theory of meaning employed, what a child means when he says something "is red", is that something looks red. But if that is what children mean, then of course they are right (and they are justified in saying that).

The problem with their utterances is not, after all, that they are incompetent speakers of Adult-English, but they are speaking another language - Child-English. Furthermore, what they are saying in Child-English is true and justified to them. In other words, Firth does not, as P1 allows, simply exclude the children from the community of competent speakers of English, but he goes further and attributes to them membership in another linguistic community where they are fullfledged members.

Let us see what additional premise would allow Firth to move beyond excluding them from our linguistic community to putting them in another. According to the principle P1, if we assume that a speaker is in a perfectly adequate perceptual position for saying what color something is, then he can only give a report which is either justified or which embodies a verbal mistake. If it is justified, then there is no verbal mistake; if there is a verbal mistake, then it is not justified. Firth complicates the matter by interpreting the utterances of the children as utterances of another language in which they do not embody a verbal

mistake. But if they are not verbally mistaken, then they are justified. To get this result, Firth can either infer, from the adequacy of the children's perceptual position, that the children are linguistically competent, or he can infer from the adequacy of the perceptual position that the children are justified in claiming what they do. Whichever inference he draws, he can then employ the conclusion of his inference together with P1 to reach the conclusion of the other inference. (That is, assuming adequate perceptual position and concluding from that that there is no verbal mistake, then, using P1, we can conclude that the claim is justified; alternatively, assuming adequate perceptual position and concluding that a claim is justified, then, using P1, we can conclude that there is no verbal mistake.)

But the first option is wildly implausible - to draw that inference would be to suppose that in an adequate perceptual position we cannot commit verbal errors. The second option is both initially plausible and very familiar. Firth does not reinterpret the utterances of the children because he could suppose that they can't get their words wrong; rather, he must be reinterpreting their utterances so as to make them come out justified because he supposes that they can't be getting the facts so very wrong. But this, notice, is a cheat from the perspective of principle P1 alone. That principle allows us to infer only that the children, from an adequate perceptual posi-

tion, will either make justified claims, or give inept linguistic performances. P1 could not entitle us to reinterpret utterances so that they come out justified, since it leaves the option that someone will say something grotesquely false and unjustified by way of linguistic failure. Firth reasons from P1, but to get his conclusion he must also use another principle:

P2: If a speaker's perceptual position is adequate, then that speaker's perceptual claim or report is justified.

P2 will, together with P1, get Firth to his conclusion. P2 will, given the assumption that the children are in an adequate perceptual position (in this case, for saying what color something looks), entail that the children's perceptual reports are justified. But since the reports under consideration are not both justified and competent utterances in English, Firth is required (now by P1) to make them competent utterances in some other language. In other words, since what they say is justified (from P2), and what they say is not justified if understood as a sentence of English, then it simply follows that what they say cannot be understood as a sentence of English, but only as a sentence of a language in which it is justified. Then, since the perceptual position is adequate and the claim is justified, the utterance is not incompetent (in the language in which it is given an interpretation under which it is justified; this follows from P1).

Important for our purposes is the use of P2. According to that principle the adequacy of a speaker's perceptual position alone enables us to infer that he is justified. In a more familiar idiom, this is the view that perception confers justification, or that justification comes from, or derives from, perception. This view, we will see shortly, lies at the heart of any theory of knowledge which can call itself empiricism. But this view, when combined with the use of an extended notion of verbal mistakes as employed by principle P1, leads to a situation which seems quite perverse. P1 allows us to say that if a person makes a blunder in a perceptual position which is beyond improving, then he blunders his words; but he still makes a blunder. When we can't get the facts so wrong, we can still use our words wrong. We are guided in our use of the extended notion of a verbal mistake by the assumption that we can understand someone's utterances well enough to recognize that what they've said is false and unwarranted, and so infer from their good perceptual position that they must have made a verbal error. But if we also use principle P2, then we can't understand the utterances of a speaker whose perceptual position is adequate in such a way that it is false and unwarranted; in fact, the correct interpretation of their utterance must be one which renders it true or, at least warranted. Although, as Firth's example shows, this will require understanding the prima facie untoward utter-

ance to be a sentence in another language or idiolect, that is the price we pay to use both P1 and P2.

The upshot of this is that we are blocked, by the kind of move Firth makes, from drawing a distinction between verbal incorrectness (in a language) and verbal correctness (in the same language). This defeats the possibility of finding an account of justification along the route proposed earlier. That proposal was to use an account of verbal correctness and incorrectness to explicate the notions of 'justified' and 'unjustified'. But if being in an adequate perceptual position is sufficient for one's reports or claims to be justified, and if one's claims or reports being justified rules out charges of verbal error, then charges of verbal error or incorrectness cannot be supported. The use of P2 forces the result that any acceptable interpretation of a perceptual claim must render the claim justified, and so verbally correct in some language. So, for perceptual claims made in adequate perceptual positions, there is no such thing as a verbal error in the extended sense of that notion we have been employing. There are only correct and incorrect interpretations of perceptual claims, and the correct interpretation can only be reached through a theory of justification - in particular, through the view that our perceptual claims are justified by our perceptual position. Firth uses P2 to force a different interpretation of the children's perceptual reports; but, without the use of P2, Firth might just as well

have said that our charges of verbal incompetence against the children cannot be supported because of the possibility that they are speaking an idiolect. As the consistency of their linguistic practice indicates, they may well be speaking an idiolect. Where we extended, earlier, the notion of verbal error to account for the systematic utterance of false and unjustified perceptual claims, we could as easily have supposed that the claims were true or justified claims in an idiolect. If we do not follow Firth in employing P2, or find some other principle which will do the same job, then there is, in practice, no distinction we can draw between linguistic behavior which indicates systematic verbal error and linguistic behavior which indicates the use of an idiolect. What stands between us and a theory of justification built on a theory of linguistic competence (of the relevant semantic sort), is the doctrine W.V.O. Quine has made famous as the doctrine of the radical indeterminacy of translation. Although I have had to squeeze Firth's example hard to get the problem of translation out of it, Quine makes the same point as Firth in its most general and explicit form, and proposes the same solution to the indeterminacy of perceptual reports as Firth employs. Quine provides our second example of the doctrine which blocks the theory of justification proposed earlier.

J. Quine: Verbal Errors and the Indeterminacy of Translation

In providing an illustration of the doctrine of radical indeterminacy of translation, Quine offers the following:

Take the case where, rather than charge someone with an altogether absurd belief, we conclude that his use of a crucial word differs from ours. This is, on the face of it, to conclude that our disagreement with him is verbal rather than factual. Still, our conclusion is no more than a trimming of our speech community to exclude our well-meaning but ill-spoken friend.¹⁵

Thus far, Quine's example is only a case of the inference we are entitled to by P1. In the right perceptual circumstances P1 requires us to conclude that the problem is verbal rather than factual; if the perceptual circumstances are not adequate, then our friend's belief will not appear absurd (unjustified). But Quine goes on in the next sentence, having cut the speaker from our speech community, and proceeds to establish him in another: "Our friend's utterance counts then only as a foreign homophone of the absurd sentence".¹⁶ Quine makes no bones about what leads him to proliferate languages in this way, and that's the place we need to direct our attention.

Translation suffers, on Quine's view, from a very extensive form of indeterminacy. A translation manual of a language is a theory about what sentences in that language mean, and all theories are underdetermined by the evidence for them. But translations are underdetermined in more

serious ways than other theories simply because they are theories about what sentences mean. The example just given illustrates the difficulty. Our decision about what someone's utterance means depends upon a decision about whether it's true or false (or, whether it's justified or unjustified); but our decision about whether it's true (justified) or false (unjustified) depends upon our decision about what it means. We are, initially, at a complete loss about how to understand the utterances of others, even the utterances of putative speakers of our own language - simply because they are only putative speakers of our own language.

Firth did not hesitate, however, to offer a translation of the children's utterances. But neither could Quine hesitate about that example, because on his view there is a way to escape indeterminacy of the radical form for some classes of sentences. If there were not some such escape route, we would never find an inroad into the translation of other classes of sentences. Prominent among the sentences which we can translate are those Quine calls "observation sentences". We can translate observation sentences when we find them because we can assume them to be true. With respect to the observation sentences of their language, speakers are infallible. As Quine says in Word and Object, with observation sentences there is no "scope for error and dispute".¹⁷ If we assume that what the speakers are saying is true (or justified); then our observation of their linguistic behavior and its attendant circumstances

will yield a correct translation of their utterances. But why are we allowed to assume that the observation sentences of a speaker are true, rather than suppose that they might be false and unwarranted because they embody a systematic verbal error?

Quine's answer is that, with observation sentences, "witnesses must in general be able to appreciate that the observation which they are sharing is one that verifies the sentence".¹⁸ That is, we can translate observation sentences because the fact that they are observation sentences is sufficient to leave us justified in taking them to be true. The mutual interdependence of determinations of truth value and meaning is tamed with observation sentences, because the fact that they are "keyed directly to observation"¹⁹ is sufficient for us to conclude that whatever a speaker's observation sentences turn out to mean, they will be justified. Quine breaks out of the circle of meaning and truth value by employing what we have expressed as principle P2. This is, importantly, not the assumption that perceptual reports are justified; it's the assumption that they are justified by our perceptual positions. If perceptual reports were justified, but along some indirect route through other statements, then observing a speaker's linguistic behavior and its attendant perceptual circumstances would not justify us in taking his sentences to be true, and so would not yield a translation.

There is not, among the considerations we have raised, a defense of P2; nor is there an account of how the required justificatory relationship works between perception and statements. That there is such a relationship is an assumption which, we saw in our earlier consideration of traditional theories of justification, some difficulty in making sense of. The import of the position which Quine and Firth have adopted can now be seen.

Recall from earlier that we saw a possible route to a theory of justification for perceptual claims and reports. Taking for granted an adequate perceptual position, principle P1 would allow us to infer from a decision as to whether a speaker's perceptual report is a verbally competent utterance to a conclusion about whether it is justified. That is, to use an old-fashioned idiom, a theory of (semantic) linguistic competence could be seen as conceptually prior to a theory of justification, and the theory of justification could be seen as conceptually derivative of the theory of linguistic competence. But it is also true that P1 would allow us to derive the conclusion that an utterance is verbally competent from the assumption that it is justified. It is this second route which Quine and Firth have followed by adding the principle P2, which in effect embodies a theory of justification. From their perspective, the perspective of the indeterminacy of translation, we can't reach a theory of justification via a theory of linguistic competence because of our completely

general inability to make a stable distinction between linguistically competent performances in another language. Do the children of Firth's example display incompetence in Adult-English, or do they display competence in Child-English? The thesis of the radical indeterminacy of translation says that we can't come up with an answer to that question, except in the way in which Quine and Firth rid observation sentences of indeterminacy by using a prior theory of justification to yield a translation of their sentences. So, on their view, we can't travel from a theory of linguistic competence to a theory of justification, because a theory of justification comes first. If they are right, there is no hope of finding a theory of justification of the kind suggested earlier - recalling the difficulties of traditional theories of justification, there may be no possibility of there being a theory of justification for perceptual claims. ²⁰ The situation is not, however, so gloomy. There is, I will argue in the following section, a revealing 'error' in the premises which generate the spectre of indeterminacy of translation. Noting the false premise will lead directly to the rudiments of a theory of semantic linguistic competence, and from there to a theory of justification.

K. Eliminating Indeterminacy

Radical indeterminacy only arises because of the interdependence of our ways of assigning truth values to sentences and our ways of interpreting them. Since there are optional assignments of truth value and alternative interpretations possible, and the assignment of a truth value depends upon the interpretation, and vice versa, we seem to have a situation in which there are two variables, neither of which can be fixed without first fixing the other. If there were not more than one competing translation possible, or if there were not competing truth value assignments, or if there were no interdependence of truth value and meaning, then indeterminacy of translation would not extend beyond the inductive uncertainty of theories and hypotheses in general.

The dubious premise needed to make a language, or even substantial parts of a language, radically indeterminate, is seen clearly in the words of Donald Davidson. Davidson concurs with Quine on the methods of reducing the extent of indeterminacy by locating sentences which can be justifiably held true independent of any determination of their meaning. He says:

A good place to begin is with the attitude of holding a sentence true, of accepting it as true. This is, of course, a belief, but it is a single attitude applicable to all sentences.²¹

Davidson's final clause is important. He means to capitalize on the feature that sentences have truth values, and that any sentence may be held true, to reduce the indeterminacy of translation; but it is also this feature of sentences, that they have optional truth value assignments, which is needed to generate indeterminacy. Indeterminacy of the radical sort will only arise in so far as sentences have truth values. But that is not very far at all. Most of the sentences which speakers utter are neither true nor false; or, if not most, then enough of them. Consider the imperatives, the interrogatives, and the expletive sentences of a language - all sentences not in the indicative mood. (There is, I suspect, the same point to be made about the expressive uses of indicative sentences - statements of intention, rather than predictions, promises, avowals, expressions of regret, and so on, but the case is strong enough without depending on these.) The class of nonassertoric, nonindicative utterances is large enough, relative to the class of indicative utterances; to provide a broad base of linguistic practice from which to extract an interpretation of the utterances of others. The important advantage gained by extracting our interpretation from the nonindicative utterances lies in the fact that, because they lack truth values altogether, our interpretation can't suffer the severe indeterminacies which result from the interdependence of meaning and truth.

Doing this does not leave the class of indicative utterances under the shadow of indeterminacy, since the vocabulary of indicatives and nonindicatives is a common vocabulary, and in special cases where it might not be a common vocabulary, there remains extensive overlap. The use of indicatives and nonindicatives alike calls for semantic linguistic competence with a single vocabulary, so that competence can be manifested in contexts in which questions about the truth or justification of the sentences uttered cannot even be raised.

A study of linguistic competence of the relevant sort - of the correctness and incorrectness of speaker's use of their vocabulary - must be a study of a much wider area of linguistic activity than a study of the truth or justification of their utterances. Truth and justification belong to, and are denied of, only one class of sentences. But all the sentences of a language must be the subject of a theory of linguistic competence. It is a mistake from the beginning to suppose that a theory of linguistic competence could ride on the coattails of a theory of justification; once the indeterminacy of translation is dispelled, the latter theory must come second to a theory of using a language correctly.

Although the problem of radical indeterminacy of translation would arise if our language consisted entirely of sentences which have truth values, it does not arise because our language does not consist entirely of such sen-

tences. Since our queries, commands and curses share a referential and predicative vocabulary with our claims and reports, we can understand the latter through our understanding of the former. Once we have our understanding of the language a speaker is using (in particular, once we can speak of the correctness and incorrectness of a speaker's use of the referential and predicative vocabulary), then principle P1 provides us with all we need to make sense of saying that his perceptual reports are justified or unjustified.

Since, on this view, a theory of justification waits on a theory of semantic linguistic competence, and that ordering of theories depends upon the elimination of the radical indeterminacy of translation, I will proceed to a discussion, in general outline, of what those theories must look like. But first, in case it is thought that my removal of indeterminacy is too hasty, I will consider briefly some possible objections to it.

First, it might be pointed out that all the nonindicative sentences carry with them, when they are used, a large number of indicatives and their accompanying truth values. The command "Close that door!" carries with it presuppositions and implications like "That is a door" and "That is open", and the correct translation of the command would be a function of the truth or falsity of those sentences. Thus, the indeterminacy of those indicatives will immediately infect the translation of the imperative. There are

a number of things to be said to this. First, this move does not eliminate the command from the class of sentences to be translated; it simply increases the size of the class of indicative sentences to interpret. Even this grants too much, since unless the presuppositions and implications are actually uttered, they provide us with no materials, and so no difficulties, for translation. In practice, we must work in the other direction - moving from an understanding of things said to an understanding of things assumed or believed but left unsaid. Finally, it should be said in response to this criticism that it confronts a dilemma which renders it either too weak to force its conclusion, or question begging. The presuppositions and implications which are thought to generate the difficulty for the proposed elimination of indeterminacy are either logical presuppositions and implications, or they are not. If they are taken to be logical implications and presuppositions, then the criticism is question begging, since it would presuppose that the implications stand in a truth functional relationship with the command; but that assumes that the command bears a truth value, which is precisely what is being denied. If the presuppositions and implications are not in a logical relationship with the command, then the criticism is attempting to make the translation of the command depend upon the prior translation of sentences which may or may not be uttered, and may or may not be believed, by anyone. But then the relationship between the

command and its implications or presuppositions is too weak to make the interpretation of the latter a prerequisite for the interpretation of the former.

A second form of answer to my criticism of the thesis of radical indeterminacy might be thought to come from a view of imperatives, interrogatives, and expletives which takes them to be in some way equivalent to indicatives, or reducible to indicatives. Thus, indeterminacy would accrue to the nonindicatives because they are, in fact, only 'disguised' indicatives. The problem with this response goes far beyond the fact that no one has ever succeeded in producing a plausible reduction of utterances of one mood to utterances of another mood, or an analysis of one mood in terms of another. Like the previous response, it suffers from the fact that any equivalence or reduction of the sort it requires would express a truth functional relationship between the nonindicatives and the indicatives, and would therefore require the question begging assumption that the nonindicatives have truth values.

A third form of response might claim indeterminacy of translation for the nonindicatives themselves in spite of the fact that they lack truth values and have no logically interesting relationship to sentences that do have truth values. Imperatives could be shown to be subject to indeterminacy on the ground that even though they lack truth values, they nonetheless have the property of bivalence. Though neither true nor false, they are, for example,

obeyed or disobeyed. But this form of 'bivalence', and others like it, will not produce indeterminacy of the radical sort. It is simply one of the features which results in the normal uncertainty which belongs to inductive theories and hypotheses. This is because whether a command is obeyed or disobeyed is a function exclusively of a listener's response to the command, it is not a feature independent of actual audience 'uptake'. Perhaps the correct interpretation of a command is intimately tied to understanding the behavior which would count as the appropriate form of obedience to the command, but it can't be a function of whether the command is obeyed. The radical indeterminacy of meaning results from the fact that truth value and meaning are semantic properties of the same utterance; only inductive uncertainty results from the fact that sometimes our claims and reports are, for example, heard and sometimes they are not heard by others.

I conclude that it is a mistake to suppose that radical indeterminacy infects the interpretation of languages. It is a mistake because it depends upon the assumption, which is false, that there are optional truth value assignments for sentences of a language. Contrariwise, there is a significant part of languages consisting of sentences which have no truth values. Since a vast number of sentences in a language can be understood without a prior determination of their truth value, we can exploit the relationship between those sentences and the indicative sentences of the

same language to produce a determinate interpretation of the indicative sentences. But since, as I have already argued, there are no relationships of logical presupposition and implication between the indicatives and the non-indicatives, no truth functional relationships, it is necessary that there be a relationship between the sentences of a language which obtains between sentences of any and every mood, and which will allow us to move from an understanding of some set of sentences to an understanding of other sentences. But we already have such relationships - and we have been employing them all along in our use of the central notion of a verbal mistake. I will return to a discussion of these nontruth functional relationships between sentences in the next section.

The difficulties I have raised for the claim that translation suffers from a radical form of indeterminacy force a certain strategy on us when we begin to acquire the language which is spoken by our linguistic elders - or, rather, it forces a certain strategy on anyone who would attempt to explain how we go about acquiring the language. What this strategy is, and what its merits are, can perhaps best be seen by contrast to the general strategy, adopted by Quine and Davidson, to which it is an alternative. The strategy of Quine and Davidson is well summarized by Davidson:

This method is intended to solve the problem of the interdependence of belief and meaning by holding belief constant as far as possible while solving for meaning. This is accomplished by assigning truth conditions to alien sentences that make native speakers right as often as plausibly possible, according, of course, to our own view of what is right.²²

This passage clearly indicates the dependence of an interpretation of what the alien sentences mean on beliefs about what sentences we are justified in taking to be true; a central class of those sentences will be perceptual claims and reports. But Davidson is not blind to the fact that this method could not be applied directly to nonindicative utterances. As he concedes in his essay "Truth and Meaning",

When we depart from idioms we can accommodate in a truth definition, we lapse into (or create) language for which we have no coherent semantical account - that is, no account at all of how such talk can be integrated into the language as a whole.²³

Among the sentences which create "difficulties and conundrums" for this theory of meaning, he lists "the imperatives, optatives, interrogatives, and a host more".²⁴ He remains, nonetheless, optimistic that it will be possible, with enough semantic spadework, to provide a truth functional or quasi-truth functional characterization of such sentences. The strategy with which Davidson and Quine approach the nonindicative parts of a language is, then, to accept as true (or justified) certain statements of the indicative mood and then work out a translation of those sentences.

Sentences which lack a truth value, and so which will not submit to this direct approach, are eventually to be shown to stand in a complex array of connections with the indicatives, and they can then be translated through the connections they have with the already understood indicatives.

But that strategy, if it would work at all, would work both ways. If we can first understand sentences of the sort which lack truth values, then we could construct a partial dictionary of the language to begin our interpretation of sentences which have truth values. It is this second strategy which I have suggested (with one important difference - I do not suppose that the complex array of connections between the indicatives and the nonindicatives would be truth functional connections).

But the reverse strategy which I adopt is, I suggest, the better strategy for the following reasons. First, the fact that they have no truth values renders the nonindicatives immune to the indeterminacy of translation, and so our understanding of those sentences will be more secure, as will be our transition from understanding them to understanding the indicatives, than could be hoped for on the Davidson/Quine strategy. Second, and most important for our general investigation of theories of justification, the route to translation envisioned by Quine and Davidson depends upon the assumption of a theory of justification which no defense is offered of, and which, we have seen, is in itself problematic - we have not even the inkling of an

understanding of how it is that perception could confer justification on a statement of belief. Their strategy assumes that we do, by and large, get our perceptual reports right. But the question before us is how being in good perceptual circumstances contributes to a speaker's saying things which we are justified in taking to be true.

My third reason for rejecting the Quine/Davidson strategy in favor of the reverse strategy brings us to the heart of the disagreement I have with their approach to language in general, and to the theory of justification which that view of language carries along with it. Understanding the utterances of others and coming to speak a common language with them by beginning with their use of expressions in the nonindicative parts of their language is not an ad hoc proposal for escaping the indeterminacy of translation. It is, rather, an essential part of teaching and learning anything that can properly be called a language. Languages are conventional, in the sense I insisted on much earlier, and have been assuming to be uncontentious all along. To utter an expression in such a way that one can be thought to be speaking a particular language is to follow a conventional practice. But to follow a convention or a precedent is to use an expression in conformity with others, whatever that expression may be. I have characterized the conventionality of language by saying that expressions in a language are counterfactually, arbitrary, but noncounterfactually, nonarbitrary. Though any expression

could have been used in a very different way than it is, once the precedent is established or the conventional practice built up around the precedent, we cannot use an expression in a very different way than others and still be viewed as speaking the same language as they. If such anarchistic utterances attract a following, then an idiolect may arise, internal to which the same demands for conformity must be enforced; otherwise we have a verbal error - a failure to follow the convention. For a sound or mark to move into a language, to be a word in a language, the use which one speaker makes of the sound or mark must become regulated by the use made of it by other speakers - or, in an extreme case, by his own usage. Linguistic conventions are those linguistic practices of a community to which conformity is demanded. We can use a word the way our linguistic peers use it, in which case we use the word correctly relative to the conventions of our linguistic community. To fail to follow the conventions is to use the word incorrectly relative to the conventions of that community. If, in that sense, we use a word correctly, it is because we utter the right word; if we use a word incorrectly, we utter the wrong word.

It is essential to the notions of "correct" and "incorrect" usage, thus characterized, that a speaker is not held responsible only to the way that word is used in indicative sentences. In pointing and uttering the sentence "This is a table", my use of the word "table", if it is to be in

accord with the conventions of my linguistic community, must stand in the relevant relationships to the use of that word in what is generally a vast number of sentences not in the indicative mood - curses, commands, and queries make the same contribution to whether my use of a word is conventional as do claims and reports. "Damn table", "Get your feet off the table", and "Whose table is this?" make as large a contribution to fixing what will count as the correct use of "table" as does a claim that something is a table.

As mature speakers, we use a word in the context of a sentence. Whether we are putting a question, giving an instruction, or making a claim, we aim to be understood by maintaining conformity in our usage - at the level of individual words - with the practices of our linguistic peers and elders. The linguistic novice, whose tutelage consists largely of eavesdropping and imitating, uses all these acts of speech as examples to follow. But when an initiate to our language copies our use of words in an indicative sentence, by not yet understanding what we have said, he remains uncorrupted by any false opinions we may express. It is possible to follow the practices used with words even in the expression of false opinions without thereby following the opinions of others. If this imitation is not corrected, an idiolect may begin. We can thus count it as a blessing that only a small part of the speech a language learner aims to copy consists of claims or

expressions of belief. But an even greater blessing is that we have the practice of correcting someone's speech - of showing them what to do in order to be following a convention.

In learning to speak a language, an initiate - whether he be a child, or a linguist in the field - is made responsible for learning to do what others do - for, that is, imitating. But this responsibility to conform operates only at the level of individual words. What the initiate must learn to do before he can be thought to have mastered a language and before he can be thought to understand the utterances of others is maintain conformity with others at the level of individual words. What he must be thought to be aiming to do, in so far as he is thought to be learning a language, is simply to follow the conventions. This much has been said differently already. But now we can see the relevance of this to the heart of the disagreement I have with Quine and Davidson - to follow a convention or a precedent for the use of a word is not to try to find a rendering of a sentence which may turn out to be either true or false. Conventions and precedents are beyond that kind of criticism. They operate in very much the same way as commands - "use the expression X in this way". "This is white", viewed as an utterance at a ceremony baptizing the colors, is a stipulation. Conventions and precedents are prescriptions for ensuing behavior, not descriptions of present facts. To view an utterance as a convention is to

view it as a kind of prescription - as an example to follow; but children setting out to acquire a language have no other view of our utterances.

The strategy I have proposed as a counter to the strategy of Quine and Davidson is not, then, simply a device for evading the indeterminacy of translation (though I think it works well enough as that). Our coming to understand a language depends on taking a view of the utterances we are exposed to as prescriptions, examples for us to follow. The present, ongoing verbal utterances of speakers of the language is all we have to serve as those prescriptions, whether what they say is true or false or neither. Understanding a language begins with the nonindicative utterances and moves from there to an understanding of the indicative utterances, because it begins with learning to follow conventions and moves from there to being able to subordinate the use of those conventions to saying what color something is.

Any act of utterance which can be viewed as an indicative utterance in a language, L, must be viewed as an activity of following a series of conventions for the use of words making up that utterance. But those conventions are themselves comprised of other acts of utterance. The indicative utterance must, then, be viewed as owing the meaning it has in L to its relationships to utterances which serve as prescriptions for the use of the words which comprise it. But this is just the strategy I have sugges-

ted, in opposition to the Quine/Davidson strategy, for escaping the indeterminacy of translation. The heart of the disagreement that I have with those philosophers is, then, that they do not view language as conventional, in the trivial semantic sense, or, at least, that they do not consistently view it as such.

This shows itself clearly in the Quinean move of "trimming error" from our language and then elevating the error to the status of a foreign language or idiolect - a move which I have been resisting. The new languages thereby proliferated lack a set of trivial semantic conventions such that the use of the questionable word not only failed to conform with our usage, but also failed to conform to the usage of any other linguistic community. No one can say in what language our "ill-spoken friend" has gotten his usage right, or to what community of linguistic practices he has conformed.

But there is a deeper and more significant way in which Quine fails to remain consistent with the assumption that languages are, in the trivial semantic sense, conventional. That is the topic of the next section, which returns us to the problem of justification of perceptual reports.

L. Empiricism and the Conventionality of Language

There is an apt metaphor, or set of metaphors, drawn from physics and chemistry, which will help us to express

the defect in the principle of justification, P2. There is a "normative force" or "pressure" on a speaker to conform to the conventional practices of a community if he hopes to be understood. But this normative force is exerted only on the atomic parts of our sentences. In our indicative and nonindicative moods alike, we employ the same vocabulary, and so the demand for conformity in the use of our words applies across and between these categories. In this way, our commands and curses tend to keep the language honest; we can summon the usage of a word in them to help exact obedience from the usage of the same word in an indicative.

No sentence exerts on any other sentence any normative force other than what the individual words used in one sentence exert on the individual words of the other. The normative force operates, so to speak, at the atomic, but not the molecular level. We are, then, under no pressure from our desire to be understood to conform with the opinions of others.

When speakers use the words provided by their linguistic community, they are engaging in an activity which is continually open to criticism by their linguistic peers. But though the usage of each speaker is open to criticism and correction by his community, the community itself is not subject to criticism or correction of its usage by any external authority. Since the standards an individual must copy, or the prescriptions he must follow, in order to use a word correctly, consist only in the conventional prac-

tices of his community, there are no standards, external to those practices, which the community itself can strive to imitate. The pressure to choose one word rather than another is exerted by the community, and no pressure is exerted on the community, or on an individual speaker, from outside. To suppose that the world itself, or our experience of it, provides any standard against which to measure the correctness or incorrectness of a speaker's use of a word is to leave the correctness or incorrectness fundamentally undecidable. It is only in light of the actual linguistic practices of a community, however broadly or narrowly that must be specified, that we can make sense of a word being used correctly or incorrectly - the trivial semantic conventions are the standards with which our usage must accord, if our usage is to be correct, and with which our usage must fail to accord, if our usage is to be incorrect. We are under pressure to follow the present practices of the community - to follow the trivial semantic conventions; at the atomic level of our language, the pressure to follow those conventions is the only pressure we are under.

Classical empiricism metaphorically asserted that all knowledge or justified belief 'comes from', 'derives from', or 'is based on' perceptual experience. But that view, together with the claim that there actually is some knowledge or justified belief, entails that some perceptual experience or set of perceptual experiences is logically or causally (depending on how those metaphors are cashed out) sufficient for someone's having knowledge or being justi-

fied in believing something. It is this thesis which is expressed in P2. But P2, as applied to perceptual claims or reports, is clearly false unless the proviso is added that we must exclude verbal error. Without P1, P2 would admit of straightforward counterexamples - each verbal error would be an unjustified perceptual report made in the best of all possible perceptual circumstances. But, as we saw in the examples of Firth and Quine, when P1 and P2 are both employed, the result is that no one can make a verbal error from an adequate perceptual position. Putative verbal mistakes are not the use of the incorrect word in one language, but are the use of the correct word in another language. But it should be evident that the languages proliferated by this move lack the support of any linguistic community - there are no conventional practices which the speaker's usage can fail or succeed to be in accord with. The only sense of 'correct' usage which empiricism has at its disposal is what can be provided by supposing that the standard for correct usage is given in the experience. To that extent, empiricism is in conflict with the view that our languages are, in the trivial semantic sense, conventional; it needs a view of languages which severs them, at the level of the usage of individual words, from any restrictions imposed by actual linguistic communities.

We can summarize this point by looking at an example. Suppose, in the best possible perceptual circumstances, someone sees a red object which looks red and then states,

or believes, "It is blue". He will not, suppose, retract his statement (or change his mind) after further reflection and observation. Principle P2 says that from such a perceptual position we are justified in taking his statement or belief to be true. Turning to principle P1, we find that the second conjunct of the antecedent is falsified by this use of P2. We are not, then, entitled to infer the truth of the consequent. The fact that what he says or believes is justified prevents us from supposing that he has made a verbal error. This last part conforms to our usual practices well enough - we cannot lay a charge of verbal error against someone who has just uttered a sentence which we take to be the truth about the matter. What he has said is, because justified, verbally unobjectionable. But now, since his utterance is verbally objectionable, we should be entitled to ask, on the assumption that languages are conventional in the way described earlier, in accord with the conventions of what language, or what linguistic group, has he used the expression "blue" correctly? To this, no answer is offered - there is no such language, no such linguistic community. The best that might be hoped for is something like what is provided by the following passage from A. J. Ayer (which he himself states only to take objection to):

If, provided I am not lying, my statement must be true however I express it, then even though I am using words which belong to a public language, and using them correctly, there is a sense in which my

use of them is private. It is private in as much as the meaning of my words is supposed to be fixed entirely by the character of the experience I am using them to indicate, independently of any public standard of usage. 25

Saying that we are "using words which belong to a public language, and using them correctly" suggests one clear sense of "correctly" - namely, that we are using the words in accord with the conventions of the public language. To fail to do that much is to use the words incorrectly. But, there seems to be yet another sense in which we may use our words correctly, which is given by saying that the meaning of the words is fixed "independently of any public standard of usage". This gives us a second standard for the correct use of the words. If there is a second standard, then it ought to make at least prima facie sense to speak of correct and incorrect usage as judged in light of that standard. It is this position, that we could be using our words correctly, as judged in light of some standard despite, even, the fact that we are using them incorrectly in light of the standard provided by public usage, which conflicts with the view that language is, in the most trivial sense, conventional. To say that our use of words is conventional in this sense is to say that there is no standard, other than that provided by public usage, which our usage may meet or fail to meet.

It is tempting to respond, on behalf of the empiricist theory of meaning, that we must read the claim that the

meanings of our words is "fixed entirely by the character of the experience I am using them to indicate" as saying that the words under consideration are given their meaning by their being used to indicate a particular experience. In this case, the usage of the words in this act of speech is genuinely immune to criticism which might be offered in light of the already prevailing public standards of usage. But such a response will not give the empiricists what they need - although the usage cannot, then, be incorrect, neither can it be correct. The stipulative act creates a precedent against which verbal correctness may be measured; it does not meet a standard. The only relevant standard already present is the experience itself.

We find expressed as clearly as we need the view which I am claiming empiricist theories of justification must be committed to - that the meaning of his words is "fixed" by the experience he is having.

As we have seen, the theory of justification expressed by P2 is open to counterexamples unless it is coupled with the principle of linguistic competence expressed by P1. But when the theory of justification is coupled with the principle of linguistic competence, we conflict with the assumption that languages are, in a trivial sense, conventional - that is, with the assumption which may be expressed by saying that "verbal correctness" makes sense only in light of some set of linguistic practices which we may either succeed in following or fail to follow.

1. r

The assumption that language is convention in this trivial sense is too deeply entrenched and, perhaps, too obvious to give up. But principle P2 cannot stand without P1, whereas P1 will stand without P2. We have no option, then, but to reject P2. To reject that principle is to reject the view that any perceptual experience, or any set of perceptual experiences, is sufficient to confer justification on a statement or belief. It is, in short, to reject any theory of justification which deserves to be called empiricist.

This leaves us with a need for an account of the rôle which a speaker's perceptual position plays in his saying things which we are justified in taking to be true. In the next section, an account of the sort suggested earlier, which accounts for the justification of perceptual reports on the basis of a theory of linguistic competence, is sketched.

M. The Justification of Perceptual Judgements

An account of the justification of perceptual judgements can be derived from principle P1, as follows.

Principle P1 says that:

- 1) If a speaker's perceptual report is given in adequate perceptual circumstances, and that report is not justified, then that report involves a verbal error - letting 'P': the speaker's perceptual report is given in

an adequate perceptual position; 'Q': we are justified in taking the report to be true; and 'R': the report involves a verbal error - the principle P1 may be expressed as $(P \ \& \ -Q) \ > \ R$.

Which entails:

- 2) If a speaker's perceptual report does not involve a verbal error, then it is not true both that the report is given in adequate perceptual circumstances, and it is not justified; $[-R \ > \ -(P \ \& \ -Q)]$;

which in turn entails:

- 3) If a speaker's report does not involve a verbal error, then either his perceptual position is not adequate, or the report is justified; $[-R \ > \ (-P \ \text{or} \ Q)]$;

which finally entails the principle we need, that:

- 4) If the speaker's report does not involve a verbal error and his perceptual position is adequate, then we are justified in taking his report to be true; $[(-R \ \& \ P) \ > \ Q]$.

That this conclusion does follow from P1 is, I think, obvious. What is not obvious, however, is how this conclusion can be understood to be the basis of a theory about the justification of perceptual reports. It becomes seen

as an account of justification when it is viewed in the light of our prior conclusion that P2 is false.

Although this conclusion, 4), is equivalent to the original principle P1, which has been employed widely by philosophers who would maintain that perception itself provides some justification for our beliefs - whether because perception 'verifies' our belief, or because it 'confirms' them - the rearrangement of the parts of the principle allow a realignment of the way in which we view perception to be related to the justification of judgments.

The original formulation of P1 allows us to infer from the assumption that a person's perceptual position is adequate, and their belief or report is not justified, to the conclusion that they have made a verbal mistake. Using the principle in this way, we must first make a judgement as to whether a speaker's report or belief is justified before we can reach a conclusion about whether or not he has made a verbal mistake. In the application of the principle, we move from an investigation into the justification of a report to a discovery about the linguistic credentials of the utterance. No account of what our investigation into the justification of the report is an investigation into is suggested, and none is ruled out.

The reformulation of P1 as 4), on the other hand, allows us to infer from the assumption that a person has not made a verbal error - that is, from accepting his

report as a linguistically competent utterance, to the conclusion that it is justified. Using the principle in its reformulation, we may begin with a judgement as to the correctness of a speaker's use of language to a conclusion about the epistemic worth of his report. In the application of the reformulated principle, we infer from the fact that a speaker has established his linguistic credentials to the conclusion that we are justified in taking what he has to say to be true. The infinite regress which troubles foundationalism, and the threat of circularity which troubles coherentism, disappear, providing only that verbal correctness can be explicated, and the judgement that a report is verbally correct can be made, without reference to a prior notion of justification, or to prior statements or beliefs which are in need of justification.

The discussion of the conventionality of language provided an account of the correctness and incorrectness of the use of a word which was independent of any theory of justification. We have, then, an account of semantic linguistic competence which can be used to account for the justification of perceptual claims and reports in light of principle P1.

Verbal slips, verbal errors, displays of linguistic incompetence of the sort which have done so much epistemological work behind the scenes, are failures on the part of a speaker to maintain conformity with the conventional practices of a linguistic community. The errors are only

verbal, in the sense that the wrong word is uttered when, in fact, there is another word which the community actually uses for that purpose. There are some perceptual positions which are good enough for us to learn the use of words like those in our color vocabulary. But if, after having learned the proper use of our color words, we return to a perceptual position sufficiently like the one in which we learned them, and we use the word "red" to describe green things, then we display a failure of our competence with the vocabulary. If there is no room left for improving a speaker's perceptual position, and we still reject his perceptual claims and reports, then we reject them solely on the grounds of linguistic incompetence.

We are now working with the assumption that a speaker's perceptual position is adequate - it is, if you like, a perfect one. For these kinds of perceptual positions, P1 expresses a relationship between something's counting as a justified perceptual claim, and its counting as a linguistically competent utterance. Every linguistically competent claim is a justified claim, and vice versa, because we exclude from the class of justified claims every verbally inept utterance. There cannot be any unjustified perceptual reports of adequate perceptual positions simply because all such reports have been ruled out as inept utterances.

Our confidence in perception is not a confidence that it will somehow deliver up truths or justified beliefs; it

is a confidence that we will have an error of a special kind - a verbal error - if we do utter any falsehoods or unjustified opinions. We can infer, from the fact that a speaker is not making a verbal error, that the speaker's claim is justified; and we can infer from the fact that a speaker is making a verbal error, that the speaker's claim is unjustified. But then, the confidence we place in a speaker's perceptual reports is justified in so far as his sentences withstand the scrutiny of linguistically competent speakers (his own included), when that scrutiny is focused on the success or failure of the report to conform, at the atomic level of the sentence, to the trivial semantic conventions of the linguistic community. The high epistemic ranking of such reports in no way derives from the perceptual experience had in that perceptual position. -> there is, across the bridge of perception, no epistemic traffic.

This is not to belittle the role of perception. Our seeing something red, and its looking red to us, in adequate perceptual circumstances, serves to distinguish our utterance of the word "red" for saying what color it is, from many other utterances of the word "red" for saying what color something is in inadequate perceptual circumstances. The decision to use the word "red", when that decision is made in clear sight of a red object, is a decision we can, epistemically speaking, respect on the single assumption that the decision was made by a competent

speaker. And that assumption is testable in a myriad of ways. The point of requiring the competent speaker to reach an adequate perceptual position is that only then are we allowed to rule out accounts of how he may have been led to make the wrong choice of a crucial word in spite of having exercised linguistic competence. As long as it is possible that he is too far away, or that the lighting is too dim, or that he is drugged, then we cannot take the perceptual position he is in to make possible a fair test of his linguistic competence. It is, in short, the same reason we cannot give a fair test of a language student's mastery of our color vocabulary when the lighting is too dim or the objects he is to describe are too far away - his failure to answer with the correct word when we ask "What's the word for this color?" does not indicate a failure to have mastered the assigned vocabulary.

So a speaker's perceptual position is, on this account, relevant to the assignment of an epistemic ranking to his claims about his perceptual position. In seeking knowledge or justified belief, the role of modifications and improvements in one's perceptual position is to remove conditions which would otherwise defeat an evaluation of the utterance of a claim for its linguistic competence. We want to so construct our perceptual situation that we can eliminate factors which would defeat the supposition that our perceptual situation is good enough to test for linguistic competence, where the test for linguistic competence

is reduced to an inquiry into the conformity or lack of conformity of the utterance to the trivial semantic conventions of the language in which the utterance is couched. The possibility of carrying out this inquiry with the desired results does not itself rest on an assumption that the utterances with which the speaker's usage must maintain conformity are either true or justified - they may be neither, but be the utterances of words in questions, commands and curses, as well as the utterances of those same words in other claims and reports. The relevant conformity obtains, when it does, at the level of individual words, where inferential relationships cannot operate. The difficulties of foundationalist and coherentist accounts of justification, however serious they may be, cannot even arise at this level, and so are avoided.

Finally, it should be noted at what level, on this account, the justification of perceptual claims and reports occurs. The practice of justifying a perceptual report, given adequate perceptual circumstances, consists in inquiring into the correctness of a speaker's use of crucial words in his report. The relevant notion of 'correctness' which operates in the account is explicated as conformity to the trivial semantic conventions of a linguistic community. But these are not conventions which can govern what speakers say, when that expression is taken to refer to what has been reported, stated, or claimed; they govern only our ways of saying what we do. So the operative notion of correctness

is verbal correctness. The distinction, as it is commonly drawn, between what a speaker says and the saying of it, allows that what two speakers say may be the same though the one person's saying of it cannot be the same as the other person's saying of it. It is then, on this account, only in their saying of something that speakers can be justified; what they say cannot be justified. This consequence of the account of justification is, I suggest, one that is to be hoped for as a consequence of any theory of justification which hopes to make sense of the justification of perceptual claims and reports.

26 Consider two speakers who say the same thing, namely that the car is brown. One is looking at the car in an adequate perceptual position, and the other is not (and never has). If what the first speaker says is justified, then what the second speaker says is also justified, since what they say is the same thing. But this consequence is clearly unacceptable. Insofar as the perceptual position of the first speaker is to be thought to make any contribution to our giving his report a higher epistemic ranking than we give to the other speaker's report, then we are compelled to say that the first speaker is justified in saying something, while the other speaker is not justified in saying that.

Some philosophers who have employed a distinction between what is said and the saying of it have also argued that truth and falsity are notions applicable only to what we say, and not to the saying of something. If they are

right, that would force from the account of justification I've offered a sharp distinction between justification and truth; truth could not, for example, be explicated as rational acceptability. But, although we would generally hope there to be some important and interesting relationship between what we are justified in saying and what is true, there are no special difficulties which confront this theory of justification with regards to saying what that relationship is. On the contrary, the theory strongly suggests that the relationship between justification and truth should be explored through an account of the relationship between saying something and what one says.

The kind of view which I've offered some criticism of says, in effect, that some modification of our utterances, at the level of our usage of individual words, is called for by our perceptual experience; what I've argued is that that isn't possible - the pressure to modify what we say is not exerted on our sentences by experience, but by the demand for conformity with the conventions of our language.

In the face of our experiences, we assemble a sentence which can be uttered to characterize those experiences; at each word pressure is exerted on us to choose one word rather than another. At the crucial junctures - when we reach the position in the sentence to utter the word for saying what something is, or the word for saying what color it is, we are under pressure to utter one noun or one adjective rather than another. Empiricism, as I have char-

acterized it, supposes that the normative pressure exerted on us to utter the word "red", when confronted with a red thing, is exerted by the red thing, or our experience of it. The better picture, I have argued, is that the normative pressure exerted on us is exerted only by the current linguistic practices of our community.

N. Conclusion

Earlier (pages 330 - 332), we saw a possible criticism of the practice of ridding ourselves of unjustified reports given in adequate perceptual circumstances by calling them verbal errors. This practice appeared to be an ad hoc method of eliminating counterexamples to the thesis that our reports are justified by perception. We saw, then, that the charge of ad hocness could only stand up if the criterion we need to establish that a verbal mistake has been made consists in the fact that the report given was unjustified. P1, we can now see, withstands this criticism, since the criterion for verbal error or verbal correctness is the report's fit, or failure to fit, at the level of the usage of individual words, with community practices.

We also saw in earlier chapters that the view of language that has been taken here, which places so much emphasis on the need for conformity of usage with one's linguistic peers, seems to invite the unpleasant companions of

antirealism and skepticism. Antirealism appeared to accompany the claim that there are no basic or primitive relationships between language and the world struck and maintained by establishing precedents for speech and by following them. The sense to be given to saying that our statements and our conversations are about the world will, on this account, bottom out with the notions of following precedents or conventions and subordinating acts to achievements. But those are relations between one act of speech and others, or between acts and the achievements they result in, not, like the more traditional notions of word reference or denotation, relationships between words and things. The spectre of antirealism is removed by showing how, with the notions of 'following' and 'subordination', our judgements can still be understood to be judgements about the world - in the cases I have emphasized, judgements about the things we perceive. But since the account of perception is realist, such that we perceive the spatiotemporal particulars of the world, our judgements are about the spatiotemporal particulars.

The possibility was left open, by the account of the intentionality of our talk, that though we speak about the world, what we have to say about it may be false - in this way skepticism threatens the account. But in this last chapter we have seen how the threat of skepticism is removed, despite the conclusion that the objects of perception, and our perception of them, stand in no justificatory

relationship with our judgements about the world. Relationships internal to a community's use of a language suffice to explain how our perceptual judgements are justified. The theory, then, tells us how we come to have beliefs which are about the world, and how we come to be justified in taking those beliefs to be true. Our language is not tied to the world by 'primitive' semantic relations (like denotation) at the level of individual words; but what we say is sometimes about the world, and we are sometimes justified in taking what we say to be true.

To say that language is not 'tied to' the world is not very helpful; but, then, my argument has been against this way of talking altogether - to wonder about the relationship between language and the world is to ponder a ghost. We do not find it useful to wonder about the relationship between football and the world, and yet it stands to the rest of the world in the same relationship as language - speaking a language, like playing football, is one of the things we do in the world. More useful questions to ask are: In what does the activity of speaking a language consist? What acts, skills and native abilities do we subordinate to speaking a language? And, to what higher objectives do we subordinate the activity of speaking a language?

The banal analogue of football also illustrates the role of perception in the acquisition of justified beliefs. Because we could neither learn to play football, nor play

it well, if we could not perceive, perception is rightly thought to play an important role in that game. But it is not thought to contribute points to one's game. The 'points' - the successes that mark the progress of our inquiries - consist in the acquisition of justified beliefs, and perception has been widely held to contribute those. Defective perceptual positions can make it impossible for us to learn and play football, and poor lighting and eyesight make it difficult; at a certain point in the deterioration of our perceptual position, any points scored are a matter of luck. But when the lighting is good, and our eyesight is beyond reproach, we have an opportunity to learn and play the game. Whether we play it well or poorly is a function of the skill we acquire in practice. This same relationship between perception and justified belief suffices to characterize the role of perception in our inquiries. The burden of acquiring justified beliefs falls on the exercise of our linguistic competence, though perception provides us with the opportunity to acquire and exercise that competence. As a coach or referee might examine the lighting conditions and then concede that any passes caught will be a matter of luck, so an examination of a person's perceptual position might conclude that if any true beliefs are acquired, they will be a matter of luck. But when the lighting is good enough, the evaluation of our athletic, or our epistemic successes and failures, has only our exercise of skills to focus on.

Thinking about the subject matter of epistemology (i.e., human inquiry) teleologically forces a realignment of perception and judgement. It requires us to cite, in our explanations of the activities and achievements of inquiry, relationships which do not come into sight when we try to cast the same subject matter into a causal framework. I claim as the advantage of using the teleological approach that it realigns the activities studied in such a way that they fall into place under a set of relationships which also characterize the rest of human activities - the relationships between skills, the opportunities in which we exercise those skills, and the achievements we manage as a result of the convergence of skill and opportunity.

Descartes' picture of himself, as secluded from society in a stove-heated room and, relieved of the cares and passions of a busy life to go about the project of pure inquiry, which I used in the first chapter to illustrate the form into which his influence cast epistemological problems, is the picture of an inquirer who has cut himself off from the tool needed to "measure the merits"²⁷ of his results. The two features of the Cartesian philosophy which I have argued against are his wholesale rejection of Aristotle's teleological framework in favor of an investigation into the causes of our ideas, and his isolation of the activity of inquiry from society. By restoring the teleological framework and by giving the community its role, we avoid the Humean separation of the life of inquiry and life in com-

munity. Hume's inquirer becomes "some strange uncouth monster, who not being able to mingle and unite in society, has been expell'd all human commerce, and left utterly abandoned and disconsolate";²⁸ only the sentiments of "the spleen and indolence"²⁹ lead him to assert that one opinion is more likely or probable than another when he abandons inquiry to "live, and talk, and act like other people in the common affairs of life".³⁰ I have argued against such a separation of inquiry and social activity. The resources for conducting an inquiry, and the tests for the merits of the results of an inquiry, lie literally in the inquirer's ability to "act like other people" - they lie in the ability to maintain conformity with the conventional practices of a linguistic community.

Notes for Chapter Five

1. I borrow this terminology from Hilary Putman, who uses the expression in numerous essays. I intend for it to have no meaning beyond what is given it in these pages.
2. Essay Concerning Human Understanding, Book IV, Chapter 1, Section 4.
3. The Foundations of Empirical Knowledge. MacMillan, London, 1940, p. 83.
4. Ibid, p. 111.

5. An Analysis of Knowledge and Valuation. Open Court Publishing Co., LaSalle, Illinois, 1946, p. 183.
6. Ibid, p. 184.
7. "The Nature of Natural Knowledge", in S. Guttenplan (ed.), Mind and Language. Clarendon Press, Oxford, 1975, p. 73.
8. "Reply to Chomsky", in Words and Objections: Essays on the Work of W. V. Quine, D. Davidson and J. Hintikka (Eds.), p. 311.
9. The notion of 'relevant similarity' employed here will come up for closer examination before the end of this section.
10. This claim should not be thought to go beyond what has already been said earlier. Maintaining a skill and then reacquiring an opportunity as good as the one in which the skill is acquired (and, hence, successfully exercised) is all that one can do to make possible another successful action. This does not, again, ensure success.
11. Paraphrased here by Anthony Quinton, in Empirical Knowledge, R. Chisholm and R. Swartz (Eds.), p. 544.

12. Alvin Goldman, "What is Justified Belief?", in Justification and Knowledge, George S. Pappas (Ed.), p. 18.
13. op. cit., Empirical Knowledge, pp. 459-470.
14. Ibid, p. 462.
15. op. cit., "Reply to Chomsky", pp. 310-311.
16. Ibid, p. 311.
17. Word and Object, M.I.T. Press, Cambridge, 1960, p. 44.
18. op. cit., "The Nature of Natural Knowledge", p. 72.
19. Ibid, p. 72.
20. This is the conclusion which Quine himself draws, for different but closed related reasons in his essay "Epistemology Naturalized".
21. "Radical Interpretation", Dialectica, Vol. 27, 1973, p. 322.
22. Ibid, p. 324.
23. Readings in the Philosophy of Language, J. F. Rosenberg and Charles Travis (Eds.), Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1971, p. 462.
24. Ibid, p. 465.

25. The Problem of Knowledge, Penguin Books, Ltd., Middlesex, 1956, p. 58.
26. In this way, the difficulty which confronts the causal theory of justification is not a difficulty for this account - it is a welcomed consequence. See pp. 291-292 above.
27. Keith Lehrer, Knowledge, Oxford University Press, Oxford, 1974, pp. 17-18. See quote on p. 15 above.
28. David Hume, A Treatise of Human Nature, Clarendon Press, Oxford, 1968, p. 264.
29. Ibid, p. 270.
30. Ibid, p. 269.

BIBLIOGRAPHY

- AUSTIN, J.L. How To Do Things With Words. Oxford University Press, New York, 1962.
- AUSTIN, J.L. Philosophical Papers. Oxford University Press, London, 1970.
- AYER, A.J. The Foundations of Empirical Knowledge. MacMillan Press, London, 1940.
- AYER, A.J. The Problem of Knowledge. Penguin Books Ltd., Middlesex, 1956.
- BOSLEY, R. On Truth: A Neo-Pragmatist Treatise in Logic, Metaphysics and Epistemology. University Press of America, Washington, D.C.
- CARNAP, R. Meaning and Necessity. The University of Chicago Press Ltd., London, 1947.
- CHISHOLM, R., SWARTZ, R. (Eds.) Empirical Knowledge. Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1973.
- CHURCH, A. On Carnap's Analysis of Statements of Assertion and Belief. In Reference and Modality. Edited by L. Linsky. Oxford University Press, London, 1971.
- COOK, J.W. Wittgenstein on Privacy. In Wittgenstein: The Philosophical Investigations. Edited by G. Pitcher. Doubleday and Company, Garden City, 1966.
- DAVIDSON, D. Radical Interpretation. *Dialectica*, 27, 1973.
- DAVIDSON, D. Thought and Talk. In Mind and Language. Edited by S. Guttenplan. Clarendon Press, Oxford Press, Oxford, 1975.
- DAVIDSON, D., Hintikka, J. (Eds.) Words and Objections: Essays on the Work of W.V.O. Quine. D. Reidel Publishing Co., Dordrecht-Holland, 1969.
- DESCARTES, R. The Philosophical Words of Descartes. Edited by E.S. Haldane and G.R.T. Ross. Cambridge University Press, 1911.
- DONNELLAN, K. Substitution and Reference. *Journal of Philosophy*, LXIII, 21, November 10, 1966.

- DONNELLAN, K. Reference and Definite Descriptions. The Philosophical Review, 75, 1966.
- DRETSKE, F. Knowledge and the Flow of Information. M.I.T. Press, Cambridge, Massachusetts, 1981.
- DRETSKE, F. Seeing and Knowing. The University of Chicago Press, 1969.
- FODOR, J. The Language of Thought.
- FREGE, G. On Sense and Reference. In Translations From the Philosophical Writings of Gottlob Frege. Edited by P. Geach and B. Black. Oxford, 1960.
- FREGE, G. The Thought: A Logical Inquiry. In Philosophical Logic. Edited by P.F. Strawson. Oxford University Press, London, 1967.
- GEACH, P. Mental Acts. Routledge and Kegan Paul, London, 1957.
- GOLDMAN, A.I. A Theory of Human Action. Prentice-Hall, Englewood Cliffs, New Jersey, 1970.
- GUTTENPLAN, S. Mind and Language. Clarendon Press, Oxford, 1975.
- HOCKETT, C.F. A Course in Modern Linguistics. MacMillan, New York, 1958.
- HUME, D. A Treatise of Human Nature. Clarendon Press, Oxford, 1968.
- KAPLAN, D. Quantifying In. In Reference and Modality. Edited by L. Linsky. Oxford University Press, London, 1971.
- KRIPKE, S. Naming and Necessity. In Semantics of Natural Language. Edited by Davidson and Harman. Synthese Library, 40, D. Reidel Publishing Co., Dordrecht-Holland, 1972.
- LEHRER, K. Knowledge. Oxford University Press, Oxford, 1974.
- LEIBNIZ, G.F. New Essays on Human Understanding. Cambridge University Press, New York, 1981.
- LEWIS, C.I. An Analysis of Knowledge and Valuation. Open Court Publishing Co., LaSalle, Illinois, 1946.

- LEWIS, D. Convention: A Philosophical Essay. Harvard University Press, Cambridge, Massachusetts, 1969.
- LOCKE, J. An Essay Concerning Human Understanding. Dover Publications Inc., New York, 1950.
- LINSKY, L. Reference, Essentialism and Modality. In Reference and Modality. Edited by L. Linsky. Oxford University Press, London, 1971.
- MUNDLE, C.W.K. Perception: Facts and Theories. Oxford University Press, London, 1971.
- PEARS, D. (Ed.) Russell's Logical Atomism. Fontana/Collins, London, 1972.
- PRICE, H.H. Thinking and Experience. Hutchinson, 1953.
- ROSENBERG, J.F., TRAVIS, C. (Eds.) Readings in the Philosophy of Language. Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1971.
- RUSSELL, B. Our Knowledge of the External World. George Allen and Unwin Ltd., London, 1914.
- RYLE, G. Dilemmas. Cambridge University Press, London, 1954.
- RYLE, G. The Concept of Mind. Penguin Books Ltd., Harmondsworth, Middlesex, 1949.
- STAMPE, D. Toward a Causal Theory of Linguistic Representation. In Contemporary Perspectives on the Philosophy of Language. Edited by P.A. French, et. al. University of Minnesota Press, Minneapolis, 1979.
- SWARTZ, R.J. (Ed.) Perceiving, Sensing and Knowing. Doubleday Anchor, 1965.
- QUINE, W.V.O. From a Logical Point of View. Harvard University Press, 1953.
- QUINE, W.V.O. The Ways of Paradox. Random House, New York, 1966.
- QUINE, W.V.O. Word and Object. M.I.T. and Wiley, New York, 1963.
- VENDLER, Z. Res Cogitans: An Essay in Rational Psychology. Cornell University Press, 1972.
- WITTGENSTEIN, L. Philosophical Investigations. MacMillan Co., New York, 1953.