

INFORMATION TO USERS

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

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

1. The sign or "target" for pages apparently lacking from the document photographed is "Missing Page(s)". If it was possible to obtain the missing page(s) or section, they are spliced into the film along with adjacent pages. This may have necessitated cutting thru an image and duplicating adjacent pages to insure you complete continuity.
2. When an image on the film is obliterated with a large round black mark, it is an indication that the photographer suspected that the copy may have moved during exposure and thus cause a blurred image. You will find a good image of the page in the adjacent frame.
3. When a map, drawing or chart, etc., was part of the material being photographed the photographer followed a definite method in "sectioning" the material. It is customary to begin photoing at the upper left hand corner of a large sheet and to continue photoing from left to right in equal sections with a small overlap. If necessary, sectioning is continued again — beginning below the first row and continuing on until complete.
4. The majority of users indicate that the textual content is of greatest value, however, a somewhat higher quality reproduction could be made from "photographs" if essential to the understanding of the dissertation. Silver prints of "photographs" may be ordered at additional charge by writing the Order Department, giving the catalog number, title, author and specific pages you wish reproduced.
5. PLEASE NOTE: Some pages may have indistinct print. Filmed as received.

University Microfilms International

300 North Zeeb Road
Ann Arbor, Michigan 48106 USA
St. John's Road, Tyler's Green
High Wycombe, Bucks, England HP10 8HR

78-11,165

NUNBERG, Geoffrey D., 1945-
THE PRAGMATICS OF REFERENCE.

City University of New York,
Ph.D., 1978
Language, linguistics

University Microfilms International, Ann Arbor, Michigan 48106

THE PRAGMATICS OF REFERENCE

by

Geoffrey D. Nunberg

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

1977

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

Dec. 19, 1977
Date

D. Terence Langendoen
Chairman of Examining
Committee,
Executive Officer

D. Terence Langendoen

Arnold Koslow

Robert Schwartz

Supervisory Committee

The City University of New York

ABSTRACT

It is generally assumed that the "lexical component" of a grammar may be treated on the analogue of the syntactic and phonological components, as a device that pairs words with specifications of the things and classes of things that they can normally be used to refer to; if an item can be used to refer to more than one extension, then, it must be accorded more than one lexical entry. We begin by showing how these assumptions come to grief over the problems posed by the phenomenon of polysemy, when one word, like chicken, is used to refer at one time to a kind of bird and at another to a kind of meat. It is argued, first that the inclusion of multiple lexical entries leads to undesirable complications in the syntax, and second that pragmatic schemata that generate multiple uses must be independently available, since the same phenomena appear in ostensive reference. An account of "deferred ostension" is presented, in which it is shown why you can point at some things to identify others; this account is then extended to the multiple uses of "descriptive terms." We proceed to argue that with descriptions, there is no way to tell which of the uses of a word is specified by the rules of language, either in practice or in principle; the semantics/pragmatics distinction cannot be drawn.

In that case, we argue, the idealization to the knowledge of the ideal speaker-listener in a perfectly homogeneous speech-community is considerably more pernicious for semantics than for syntax or phonology, as we must presume as well a community in which extralinguistic beliefs and practices are uniform. We conclude that it is more reasonable to describe the speaker's knowledge of word-meanings as part of his knowledge of the (heterogeneous) collective beliefs of the community, which we can idealize as his "system of normal beliefs," borrowing our model from game theory. Against this idealization, we show, it is possible to explain a number of disparate linguistic phenomena, such as speaker judgments that some word uses are "acceptable," the phenomena associated with "partial motivation," metaphorical word-uses, and change in meaning.

Adviser: Professor D. Terence Langendoen

ACKNOWLEDGMENTS

As chairman of both my department and my committee, Terry Langendoen has gone to considerable efforts, not only to encourage me to continue my own, but also to chart me a course through some highly treacherous waters. And if that much was his official responsibility, he was certainly under no obligation to do it the hard way, without ever making a move for the helm. I don't think this thesis could have been written with another chairman, though that, as he would readily admit, isn't necessarily all to the good.

I owe a special debt to Arnold Koslow, of the C.U.N.Y. Graduate Center Ph.D. Program in Philosophy, not only for the many helpful comments and suggestions that he offered as a member of my committee, but also for more general support and assurances. I am grateful to Bob Schwartz, also of the Ph.D. Program in Philosophy, for a number of comments he made as outside reader; had I had the sense to get together with him sooner, the arguments in the first half of this thesis might have been a good deal tighter than they now are. I am also grateful to Miroslav Rensky, chairman of the Brooklyn College Linguistics Program, who has taken considerable pains on my behalf, so that I would have time to work on this thing.

In the course of writing this thesis, I have

benefited from discussions with a number of friends and colleagues. Chiahua Pan endured endless discussions--or rather monologues--out of which grew a number of the arguments presented here. I am grateful to Fred Katz for planting the germ of the notion of referring functions, of which I have made a good deal more than he might approve of. Mark Aronoff is the only other linguist I know who had independently concluded that the phenomenon of polysemy was worth investigating; mere knowledge of this fact has been enormously encouraging, apart from the helpful discussions we have had on the topic. I am grateful to Ivan Sag, not only for his comments in a number of areas--particularly as regards the discussion of syntactic identity in Chapter One--but also for providing me with a chalet in Vermont and a pallet in Cambridge, where I managed to accomplish a good deal more work than is possible in Manhattan. I am grateful to Charlotte Linde both for linguistic and professional advice ("Get it done, for God's sake.") In discussions with Bob Asahina, Karen Blu, Jack Fuchs, Bob Newsom, and Scott Parker, I have been helped to make more sense of the psychological, sociological, philosophical and literary-theoretical literatures that deal with problems I am concerned with.

In preparing for deposit the final draft of this thesis, in Berkeley, I have benefited from comments offered by George Lakoff and Henry Thompson; unfortunately,

it was then too late to make any but cosmetic changes, adding a footnote or two, for example, to allude to the skeletons that George was good enough to discover lurking in the syntactic closet.

Finally, I am grateful to my parents, for their bemused forbearance, and to my friends, for making allowances. And finally, a note of thanks to Bobby the Horse, for providing me with a constant source of inspiration.

TABLE OF CONTENTS

Chapter		page
ONE:	THE DESCRIPTIVELY ADEQUATE LEXICON	1
1.0	Introduction	1
1.1	The Problem of Polysemy	9
1.2	Polysemy as an Epiphenomenon	12
1.3	Polysemy and Syntax	19
1.4	Problems with Linguistic Treatments of Polysemy	23
1.5	Polysemy and Ostension	40
TWO:	THE MECHANICS OF REFERENCE	51
2.1	Ostensive Reference	51
2.2	Descriptive Reference	85
2.2.1	The Determination of Meaning	85
2.2.2	The Extent of Polysemy	99
2.2.3	The Indeterminacy of Meaning	112
THREE:	KNOWLEDGE OF MEANINGS	124
3.1	Introduction	124
3.2	Meanings and Categories	132
FOUR:	KINDS OF WORD-USES	157
4.1	Normal Uses	157
4.1.1	The Problem of Acceptability Judgments	157
4.1.2	Normal Beliefs	165
4.2	Conversation and Conventionalized Uses	188
4.2.1	Convention and Predictability	188
4.2.2	The Syntax of Idioms	207
4.2.3	"Do so" and "?"	239
4.2.4	Theory-change, Use-change, and Reanalysis	249
4.3.1	Metaphorical Uses	268
4.3.2	Metaphor and Slang	294
4.4	Notes from All Over	309
	BIBLIOGRAPHY	326

Dr. Whittaker always talked like this. He spoke reasonably even to pets. He talked the way a windmill would talk, the way a sentence would talk--as he spoke, English seemed to have been dead for many centuries, and its bones to have set up a safe, staid, sleepy system of their own, in respectable secession from existence.

--Randall Jarrell, Pictures from an Institution

CHAPTER ONE

THE DESCRIPTIVELY ADEQUATE LEXICON

1.0 Introduction

Notwithstanding the title of this essay, or the fact that I will be spending the largest part of it in discussing the ways in which speakers bring non-linguistic information to bear in interpreting utterances, its subject is linguistic knowledge: specifically, speakers' knowledge of the conventions that govern the use of words. Like many other linguistic investigations of pragmatics, and of performance in general, it was undertaken in the hope of simplifying linguistic description; the line of argument follows the familiar form: "If we say that such-and-such a regularity is predicted entirely by a linguistic rule, we must introduce undesirably powerful devices into linguistic description; instead, we can show that the regularity follows from the simpler, more general rules governing other linguistic regularities, taken in concert with reasonable assumptions about speakers' non-linguistic beliefs and faculties." In the end, however, the argument turned out to shave too closely, and led, not to a more economical representation of purely lexical information, but rather to the conclusion that it was impossible to achieve such a representation at all;

knowledge of word-meanings can only be treated as an inseparable part of knowledge of other kinds of conventions and social practices, from which it cannot, even in theory, be isolated. The thesis accordingly divides itself into two parts--though none too neatly. In the first--Chapters One and Two--we will be talking largely about what a representation of knowledge of lexical meanings does not contain; in the second--Chapters Three and Four--we will present a view of meaning that runs counter to those most common in the linguistic literature, and try to show how it can be brought to bear to resolve some fairly concrete problems in semantics and syntax.

Let me begin by characterizing what we may take to be the "standard" view of the lexicon, keeping things general enough to allow a broad range of disagreement over particulars. On this view, the lexicon consists of a list of entries, each of which contains a phonological (or morphological) representation of a sign, a "subcategorization frame," which specifies the classes of syntactic environments into which the sign may be inserted, and a representation of some information that enables the speaker to use the sign appropriately.¹ This information need not be equated with the "meaning" of the sign; for at least some signs, like proper names and "natural kind" words, it may suffice that the lexicon simply pair the form with an individual constant. Nor must we insist that the lexicon specify all of the

information that the speaker requires in order to use the sign; the lexicon need make no mention, for example, of the extralinguistic information that we bring to bear in identifying exemplars of the extension of the sign, or in determining its correct application. So, for example, while it is true that no English speaker could appropriately use tiger to refer to something that was manifestly a spotted canine, we may not want to say that this restriction of use is specified by purely lexical information. Moreover, signs may be associated with certain "usage conditions"--they may be slang or vulgar, for example--which we may not want to list in the lexicon proper, but rather in some rhetorical appanage of the grammar.

It is generally assumed that the lexicon is subject to two further requirements: it must be linguistically autonomous, and it must be descriptively adequate. By "autonomous" I do not mean that the lexicon must be an entirely discrete component of the grammar, though we will assume that view here for simplicity's sake, but rather that the lexicon must contain all and only the purely linguistic information available to the speaker, pace our remarks about "usage conditions" in the last paragraph. It is thus assumed that it is possible, at least in principle, to distinguish among those propositions that are part of the speaker's "knowledge of his language," and those that are part of his "knowledge of the world," even if the former consist of no more than a

specification of the form, "The word w designates the class (individual) a."

By "descriptively adequate," I mean that the lexicon must be responsible to specify all of the distinct uses to which a given form can be put, just as the grammar is usually taken to be responsible to generate all of the acceptable strings in the language.² There has been much disagreement, of course, over how "acceptability" should be determined. On the original formulation of Katz and Fodor 1964, it was assumed that acceptability was determined solely by speaker judgments, but others have since argued that criteria of actual use were more appropriate. There has also been some dispute over whether acceptability should be reckoned solely relative to the presentation of sentences in the null context, as Katz and Fodor proposed, or perhaps relative to a specified context; some linguists have become adept at cooking up examples which sound like gibberish when presented baldly, but which resolve themselves into well-formed utterances when an appropriate context is suggested. We will sidestep both of these issues for the time being, by saying simply that the descriptively adequate lexicon will be responsible to generate all standard word-uses, leaving open the question of how "standard" shall be interpreted, and by avoiding citation of any examples that would be problematic on one or the other formulation.

These two assumptions about lexical content--we will

try to use this phrase rather than "lexical meaning," though we may slip up from time to time--follow from the assumption that lexical semantics can be studied on the model of syntax and phonology, as a part of the semantic component of a grammar. From this assumption follows another about the methodology of lexical semantics: it is usually assumed, either explicitly or implicitly, that lexical semantics can be carried out against the background of what we will call the "Chomskian idealization." As Chomsky 1965 puts it:

Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant considerations as memory limitation. . . . This seems to me to have been the position of the founders of modern general linguistics, and no cogent reason for modifying it has ever been offered. (1965:3-4)

On this last point, of course, there has been more disagreement than on the assumptions of autonomy and descriptive adequacy. A number of linguists, notably Weinreich and Labov, have argued that the grammar should be made responsible to systematic variation in use, and have proposed ways in which grammars could be written which would generate a heterogeneous output. For the present, however, we will ignore these objections, though we will return to the question of variation in Chapter Three.

It will be our object in the first part of this thesis to show, one by one, that none of these principles is consistent with our avowed goal of describing the speaker's

knowledge of the conventions governing word-use. We will begin Chapter One by considering a lexicon that is both descriptively adequate and linguistically autonomous; that is, which generates all standard uses, and which contains no encyclopedic information. We will proceed to argue, following the familiar line we outlined above, that at least some of these uses must be generated by pragmatic schemata. In the first part of Chapter Two, we will give a formal account of the schemata that generate multiple word-uses; in the second half of the chapter, we will argue that it is impossible to determine which uses are so generated, and hence to say which propositions about word-use must be included in the grammar.

All of this is calculated to bring us, at the beginning of Chapter Three, to a point at which we must re-examine the Chomskian idealization, which we will have been assuming throughout the first half of the thesis. There we will argue that the idealization to a homogeneous speech-community is far more pernicious for lexical semantics than for phonology and syntax, since it forces us to assume a community in which all individuals hold identical extralinguistic beliefs (as opposed, say, to a community in which all individuals are identical with respect to vocal anatomy or memory limitation, an idealization that we can easily live with.) In Chapter Three, we will try to present a view of meaning and use which squares with some obvious observations about the heterogeneity

of belief-systems; I might note that we will not be following any of the familiar lines that have been developed to deal with variation in phonology or syntax, and will be at pains to show why they are not applicable to variation in use. At the beginning of Chapter Four, we will try to construct an alternative idealization for describing the speaker's knowledge of conventions, and to show how it can be brought to bear to explain speakers' judgments that certain word-uses are "standard" or "deviant." In the subsequent sections of Chapter Four, our attentions will be somewhat dissipated. In 4.2, we will turn to the description of the speaker's knowledge of linguistic conventions, with an eye towards explaining regularities associated with those word-uses that have been described as "partially motivated," and will try to show how the view of meaning offered here can be used to explain the syntactic irregularity of phrasal idioms. In 4.3, we will move abruptly to the consideration of word-uses that speakers deem metaphors, which have up to now (we will argue) raised intractable problems for more conventional views of lexical semantics; these too we will try to explain in the light of our newly-offered idealization. We will conclude this section with some brief remarks on, of all things, slang, in an effort to show how the problems raised by the existence of usage-conditions, which have not up to now figured even marginally in the literature on linguistic semantics, are in fact both vexing and intriguing.

In earlier drafts of this thesis, I made a number of references to the philosophical literature which I have since expunged. Obviously, the problem of writing a "descriptively adequate, linguistically autonomous lexicon," as we have put it here, is equivalent to a part of the problem of writing a truth-functional semantics for a natural language, and most linguists have assumed that the output of the semantic component of the grammar should be a pairing of sentence-types with a specification of the conditions under which they can be used to make true assertions (or "valid utterances," if we wish to talk about other than declaratives). So we cannot reject this view of the lexicon without attendant consequences for a range of issues in the philosophy of language and related fields. Many of the arguments that follow are either paralleled or directly influenced by arguments to be found in Quine and Wittgenstein, in particular, though I will talk about the former only briefly, at the end of Chapter Two, and about the latter not at all. There are two reasons for these omissions. First, these philosophers have offered their arguments in the service of dogmas which have only a limited relevance to the empirical concerns. A linguistic's theory of meaning, for example, must stand or fall independently of the theory of truth; if this thesis has any metatheoretic aim, it is just to drive a wedge between the two areas of inquiry. Second, and more practically, I know of no two philosophers whose work has been subject to so much controversy and

reinterpretation, by writers much more sophisticated than I, and it would be nugatory to try to cite them in support of some point or another when at every turn the interpretation of their remarks might easily be called into question. It has not been possible to avoid all mention of the philosophical literature, especially in Chapter Three, which contains a fairly extensive discussion of the name-theories of meaning offered by Putnam and Kripke, and of some criticisms offered in rebuttal. I beg the reader to interpret these discussions generously, keeping in mind that I raise these issues only to facilitate the presentation of certain empirical questions.

1.1 The Problem of Polysemy

We noted above that the requirement that the lexicon be descriptively adequate is usually taken to mean that the lexicon must be responsible to generate all of the standard uses that a word could have, however "standard" is defined. This entails that different uses be accorded different entries, or sub-entries, but there is some problem, of course, in saying when two uses of a word are truly different. For the present purposes, we can say that a word has more than one use when it can be used to refer to more than one category or type of thing, either "in the world" (for proper and common names) or relativized to a particular utterance-context (for indexical terms). (I will use "common name"

throughout for the more usual "descriptive term"; see the beginning of Chapter Two for a justification.) On this basis, we say that words like bank and bill are homonymous, and assume that they are twice listed in the lexicon, just as they would be in any standard dictionary.

Yet among forms that have multiple uses, we intuitively distinguish two classes. We would be inclined to say that English has two different words that are pronounced [bæŋk], one of which is used to refer to a kind of shoreline configuration, and the other to a kind of financial institution. But when presented with the two uses of chicken, to refer once to a kind of bird, and once to a kind of meat, or with the uses of newspaper to refer variously to kinds of publications and kinds of publishers, we are more likely to say that these are "different uses of the same word." On this basis we say that forms like chicken are polysemous, rather than homonymous: not two words, but one word with several senses. "Sense" is an unhappy word to have to use here, of course, since it is systematically ambiguous between "Fregean sense" (= "intension") and "Pickwickian sense"--roughly, "a particular use of a word." (Thus one could say without logical contradiction, "Chicken has several senses, but words do not have senses.") In what follows, I will try always to spell the word with a p, as "psense," when I am using it in its "Pickwickian sense" sense.

The phenomenon of polysemy has interested both

linguists and philosophers from time to time, for a variety of reasons. But it is not on its face necessarily a problem for a descriptively adequate lexicon; the fact that two word-uses are perceived as instantiating "the same word" doesn't entail that the uses are not semantically distinct. We may be able to explain the intuition for "sameness" on independent grounds. If the "word" can be given a purely formal definition, like the "morpheme," then the intuition is not at all troubling. On the basis of its regular alternation with -miss, for example, there is every reason for saying that the -mit, of omit, commit, remit, etc., is the "same morpheme" without having to claim that it should be given a unique semantic entry in the lexicon; there are a number of other non-semantic criteria in terms of which we might be able to explain intuitions that a given phonological form is "one word" in its several uses. All that we ask is that whatever definition we select shall enable us to predict the relevant intuitions. (Later we'll be adding other criteria for polysemy that must also be satisfied.) If non-semantic analyses of polysemy fail, however, we have to allow the possibility that polysemous items are one word in virtue of having only one meaning, in which case we may have to abandon the assumption that different uses correspond to different semantic entries in the lexicon.

1.2 Polysemy as an Epiphenomenon

Philosophers, linguists and psychologists have all worried about polysemy at one time or another, though not always for reasons that need concern us here. Philosophers since Plato have worried a great deal about whether the word "good" should be regarded as univocal, homonymous, or polysemous, and what it might mean for it to be the last. (See, for example, Wiggins 1971a, 1971b, Ziff 1960, Alston 1971, Katz 1964b.) And there have been similar discussions about the meanings of words like pain, see, and mean itself. (This aspect of the polysemy problem has been a major methodological concern of ordinary language philosophy, in particular.) Linguists have often addressed this question; Kurlyowicz 1953 and Weinreich 1972 ask on what grounds we identify the eat of "eat steak" with the eat of "eat soup," (or even with the eat of "eat an ice-cream cone," since the last is usually done without utensils.)

But the question of how to tell whether two uses of a word are semantically different (or whether two different kinds of things are members of the same extension), for all its interest to various areas of philosophy, is less important to us than the question of how to deal with uses of a word that are uncontrovertibly distinct. A newspaper copy and a newspaper company don't count as instances of the same kind of thing in anyone's ontology, nor do a plate of chicken wings and a chicken on the hoof. So it will be easiest to

stick to examples of the latter sort, where the different uses clearly cut across orders or categories. (At the same time, we will want to keep in mind Weinreich's 1972 warning that it may be impossible to validate "an absolute distinction between true ambiguity and mere indefiniteness of reference.")

In the nineteenth century, polysemy was usually defined diachronically: several uses of a single form constituted "one word" when they had a common etymon. This is the most important of the "Historical Principles" that inspired the "N" of "NED," and it remains standard lexicographic practice. But by itself, etymology couldn't be an explanation for the intuitions that lead us to ask about polysemy in the first place. ("The people has a feeling for utility," Bréal said, "but does not trouble at all about history.") Moreover, etymology is no sure guide in these matters. On grounds of common origin, we will have to say that sole "fish"/"part of the foot" is a single polysemous item, or even that items that are now orthographically distinguished, such as flour and flower, are "one word." At the same time, the shoot of "shoot the rapids" (< Fr. chuter) can't be identified with the shoot of "shoot the arrows" or "shoot the works"; nor the boil of "bring the water to a boil" with the boil of "a boil on the body." We can't even identify the slug of "slug of whiskey" (<OI slog) with slug of "slug the umpire," despite the parallel with

belt and hit. (Curiously, this view of polysemy is today most frequently encountered in the philosophical literature; see, for example, Quine 1960, Katz 1972, and even Ziff 1960.)

Notwithstanding the defects of a purely diachronic treatment of polysemy, it was an interest in etymology that led many late nineteenth- and early twentieth-century linguists, particularly in Europe, to adopt the "psychologistic" views of polysemy that are now common. It was clear to the earliest investigators of semantic change--Darmesteter, Bréal, Meillet, Greenough and Kittredge, and Erdmann--that the meaning-relations that held between new and old uses of words were very like the meaning-relations that hold synchronically among the psenses of polysemous words, and that the psenses of a word most commonly had the same etymology. Accordingly, they proposed that the "laws" of semantic change were rooted in certain universal cognitive principles, which the science of etymology (in French, simply "la sémantique") could help to uncover. In his appendix to the (1900) English edition of Bréal's Essai de Sémantique, J. P. Postgate wrote that ". . . psychology, and through that metaphysics, will gain advantage from the study of semantics. . . ." And psychologists did become interested in the question at roughly the same time, and for roughly the same reasons; they hoped that an investigation of polysemy and semantic change might shed light on conceptual organization. This interest can be traced in a (somewhat broken) line from Wundt and his

successors to Miller and Johnson-Laird 1976.

Psychologistic accounts of polysemy, whether offered by linguists or psychologists, invariably make use of some notion of "core meaning" ("central concept," "seminal meaning," "Hauptbedeutung," etc.) from which psenses are derived in one of two ways. The core meaning has sometimes been regarded as an abstract concept which specific psenses are generated from by some process of "contextualization." (Versions of this view can be found in Greenough and Kittredge, Firth, Sperber, Trench, and Miller and Johnson-Laird, among many others.) This view takes two forms. "Context" can be understood extra-linguistically, in terms of the expectations that accompany a particular use of a term. Miller and Johnson-Laird, for example, suggest that the single concept represented by words like move and buy may be supplemented by information about "typical scenarios" of moving and buying, each of which involves a de facto opposition with other sets of concepts--here the notion of the "semantic field" becomes relevant. Given a particular context of use, "construal rules" modulate the scenario to provide an interpretation appropriate to that context. Other versions of this idea have been attractive to etymologists (following the motto "Wörter und Sachen") who have used it to explain the ways in which changing contexts of use induce semantic change. Often, the "core meanings" suggested have been quite abstract. (Trench 1859 saw in the varied uses of English

post a common idea of "that which is placed," for example.)

When context is construed linguistically, the approach is basically the same. Weinreich 1972 used construal rules to map the semantic features of verbs onto their subjects, as when house is used to refer to a house-percept in "a red house occurred twice." Bierwisch 1967, Leech 1969, and Fillmore 1972 have talked about the ways in which the choice of preposition can affect the interpretation of its NP head, as in "at the corner" (point)/"on the corner" (locus); "on the island"/"in the island," or "in the grass"/"on the grass." A number of linguists (Jespersen 1924, Christopher-son 1939, Kirsner 1972) have tried to show how a semantic difference between zero and indefinite singular articles could produce the differences in interpretation between "chicken" and "a chicken" from a single item.

Not all of these cases involve uses whose extensions are clearly different--some clearly pattern with the example of good and eat that we talked about earlier. Where extensions are obviously distinct, another version of the notion of "core meaning" has been invoked: it has also been treated as that single use of a word which is somehow prior to the others, and from which "extended senses" are derived by metaphorical and metonymic processes. Here again, the connection to processes of historical change is obvious; the intuition for the relation between the two psenses of shame in "He has no shame" and "America's secret shame" reflects

the same "conceptual association" whereby the second use arose in the first place.

Neither of these approaches precludes the other, and many of the writers I've mentioned held both; contextualization can be treated as just one instance of Wundt's Principles of Association, ("Association by Similarity," for instance). Take the uses of see illustrated in 1-4 (taken from Miller and Johnson-Laird 1976):

1. Isaak saw a fish.
2. We are seeing Pringle today.
3. Did you see Psycho?
4. Ambler was seen to the frontier.

All of these uses, Miller and Johnson-Laird suggest, are instances of a single sense of see, "perceive with the eye." The specific implications that accompany each case (for example, to "see" a movie is generally to see through it) arise when the verb is interpreted relative to typical scenarios (some would say "scripts" or "frames") of movie-going, frontier-escorting, and so forth. Other uses of "see," however, have to be treated differently:

5. I don't see him as president.
6. I see what you mean.
7. I saw we had lost the war.

These uses involve other senses. "to have a mental image of" (for 5) and "to understand" (for 6 and 7). Miller and Johnson-Laird don't have much to say about the relation

between these uses and "perceive with the eye," but there's no reason not to say that they are connected by associations and analogies like those that lead us to talk about "the mind's eye" (or "mental image," for that matter).³

Analyses of this sort, whether directed at problems of polysemy, metaphor, or semantic change, are curiously unsatisfying. I say "curiously" because there is little in them to argue with: clearly there is a conceptual association between chickens and their meat, or between newspapers and the companies that publish them. Yet having said this, we feel that there is more to be told; that we still don't understand why a word can be used in several ways, or why new uses should arise when and where they do; or how new uses pass from metaphors to "standard"; or how the special exploitations of these associations that we find in poetic language are different from ordinary extended uses, and why they bear a particular affect. All of these questions will concern us below, but for now we can simply accept principles of these psychologistic accounts (to the degree, at least, that we can make sense of them). What is immediately relevant is that there is nothing in any of this that creates obvious difficulties for a descriptively adequate semantics. Where we decide that two uses of a word differ only on account of contextual inferences (i.e., that they are instances of the same extension) the semantics can ignore the difference, leaving to pragmatics the problem of saying how and

where the inferences arise. When two "related" uses of a word clearly involve different extensions, we can assign it two lexical entries; the "cognitive association" between the uses will not affect the determination of truth-conditions, even if we should choose to represent the connection in the lexicon by means of some formal device. Thus, if see has three entries, then three readings can be assigned to "I can't see John wearing that hat" (corresponding roughly to "perceive," "picture," and "understand") and it doesn't seem to matter whether these readings are somehow related in cognition.

1.3 Polysemy and Syntax

Psychologicistic accounts of polysemy have been based almost entirely on intuitions that two uses of a given phonological shape were instances of the same word (the "entirely" needs qualification only because etymology has also been taken into account in some discussions.) Linguists who did not credit this sort of evidence, such as the American Structuralists, have had very little to say about polysemy, since these intuitions seemed to have no interesting distributional correlates. While they acknowledged that it was often difficult to say whether a word had one or several meanings (see, e.g., Bloomfield 1933, p. 145) and recognized the importance of metaphorical processes (Bloomfield's "transferred meaning") in semantic change, they didn't pay

much attention to either question. When they tried to distinguish polysemy and homonymy (but most did not) they usually did so in terms of purely distributional criteria.

For example, Algeo 1973 writes:

Two items with the same pronunciation but different meanings may be said to be the same polysemous lexical item if they have the same grammatical distribution, that is, if they belong to the same part of speech and follow the same rules of syntactic concord.

This sort of definition has sometimes appealed to linguists concerned about proper names (such as Jespersen 1924, Hockett 1958 and Chafe 1970) who have been reluctant to say, for example, that English had a huge number of homonymous items John and Mary, and who have preferred to treat all of John as instances of the "same word." Obviously, this definition is quite irrelevant to our interests. Algeo cites iris "flower"/"part of the eye" as an example of a polysemous item; the two interpretations of bank and file would also qualify.

But there are other distributional criteria that distinguish polysemy and homonymy in a more interesting way. When generative linguists began to look carefully at syntactic processes, they discovered that homonymous and (some) polysemous items behave differently with respect to some rules of deletion and anaphora that seemed to require some notion of "linguistic identity," (to use a term that is neutral with respect to different theories about what it was that actually had to be identical--reference, deep structure,

shallow structure, phonological shape, or whatever).

I can't mean that Bill gave Harry a tool, and received a dossier from Jane, for example:

1. Bill gave Harry the file he got from Jane.

But 2-5 are perfectly acceptable for most speakers, though the elements that have been deleted, or the referents of the anaphoric elements, differ in extension:

2. Yeats did not enjoy hearing himself read in an English accent.
3. John's dissertation, which weighs five pounds, has already been refuted.
4. The chair you're sitting in is commonly seen in eighteenth century interiors.
5. The Times has decided to change its format.

Thus, in 2.1.2 Yeats refers to the poet, and himself to his work. In 2.1.3, dissertation appears to refer at the same time to a physical object, which weighs five pounds, and a set of propositions, which has been refuted. In 2.1.4, chair refers to both a chair-token and a chair-type. And in 2.1.5, The Times refers to a publishing company, and the anaphoric it to the company's product.

Examples like these are well-known; they have been discussed by Langendoen 1966, McCawley 1968, Chomsky 1972, Postal 1969, Green 1974 and Borkin 1972. In general, the solutions proposed have involved the introduction of some formal device whereby a single "lexical entry" could be accorded more than one semantic representation. For example,

under Postal's "head deletion" analysis, 5 has a source like this:

6. The publisher of The Times_i has decided to change The Times'_i format.

and 4 has a source like:

7. The type of chair_i (you are sitting in the chair_i) is commonly seen in eighteenth century interiors.

A decapitation rule then deletes the heads "the publisher of" and "the type of" prior to the operation of deletion or pronominalization. Another proposal, due to McCawley and Green, would have these multiple psenses generated by pre-lexical transformations, which would be "opaque" to the operation of rules of deletion and anaphora. Chomsky, (then) working with the notion of lexical entry as a Katzian feature-bundle, has suggested that the relevant markers should be entered disjunctively under a single item; thus the entry for dissertation will contain the features: [physical object] v [linguistic object]. Then the requirement that rules of anaphora and deletion must operate on the "same lexical item" will be satisfied. (Chomsky also suggests that these cases are apparent exceptions to Katz's claim, in a rejoinder to Weinreich, that Boolean functions of specified features are formally equivalent to sets of discrete lexical entries.)

I'm going to assume here that all of these proposals are descriptively equivalent, under the theoretical assumptions that each implies--that there is no interesting set

of cases that could be handled within, say, a head-deletion analysis, but not within a "disjoint-entry" analysis, or vice versa, whatever the loss to "simplicity" one way or the other (a question on which these linguists are always unable, and often indisposed, to agree). I won't even worry about whether the apparatus proposed to deal with polysemy are relegated to the syntax or the semantics (or, for that matter, the semantico-syntax). The important thing is that they all require that we treat these observations about deletion, relativization and so forth in the grammar, as part of what is to be included in linguistic description. In what follows, I'll talk only in terms of a "disjoint entry" approach, where all of the operations involved in the derivation of sentences 2-5 are treated as syntactic rules. But this is mostly for the sake of clarity and convenience, because we want to keep distinct the purely semantic questions about the interpretation of polysemous items, and those questions that are "syntactic" in the broadest sense; i.e., which involve the determination of the scope of operators and the binding of variables in semantic representation.

1.4 Problems with Linguistic Treatments of Polysemy

All these syntactic treatments of polysemy are subject to two related criticisms: they complicate the grammar unnecessarily, and they are unrevealing, giving us descriptions

in place of explanations. The grammar is complicated by whatever formal device is introduced to distinguish polysemous words like newspaper from homonyms like file--in terms of Chomsky's proposal, by the introduction of two kinds of lexical entries. This may seem a necessary complication; how else could we account for the fact that some semantic differences behave differently from others? But in saying this, we make the assumption that the difference between the two uses of newspaper is semantic, and that each use must be represented in the lexicon. There are several excellent reasons, however, for supposing that many of the multiple uses of polysemous words don't have to be listed, since they would be generated by pragmatic schemata in any event.

First, we find these same patterns of polysemy cross-linguistically. We aren't surprised to learn that French journal or German Zeitung behave like newspaper, or that livre and Buch behave like book. In fact, it wouldn't occur to most speakers that things might be otherwise. If you ask someone how to say "file" in French, he may respond, "Do you mean the tool or the container?" But if you ask him how to say "newspaper," it is highly unlikely that he will bother to ask whether you mean the publication or the company; nor would it occur to you that the words might be different. We assume that French journal will behave like English newspaper, because we know that their newspapers are so much

like ours. Or take school "educational institution"/ "building for education." This psense-ambiguity comes up, quite predictably, in most Western European languages, and in Chinese. So does the "author"/"oeuvre" ambiguity; the Chinese sentence "K'ung Fu T'su hen yu i ssu" shares both of the interpretations of its English equivalent. "Confucius is very interesting." Or look at the use of the German, French, Russian and Latin verbs for "write" to mean both "inscribe" and "correspond," or at the use of their words for "door" to refer both to a kind of opening and the thing that goes inside it. These patterns of polysemy aren't universal--it goes without saying that there will be languages in which the words for "inscribe" or "doorway" are not used to refer to "correspond" and "door," if only because the speakers lack the latter notions. And in other cases, one language may have several words where English has only one. But the fact that there is such widespread similarity persuades us that polysemy is not a morphological accident; the relation between the uses of newspaper is not to be compared with the relation between the two verbs stand ("rise to the feet"/"tolerate") which happen to have an identical past tense.

One approach to this problem is suggested by McCawley 1968, in a discussion of cases like the "mind"/"body" ambiguity of person. He concludes that, inasmuch as such ambiguities are productive and universal, they are animated

by word-formation rules that are part of universal grammar. So if you have a word book₁ "inscription," you also get gratis a word book₂ "content of book₁" and so on. In appropriate forms, such rules could be fit into any of the models we are discussing--as rules of zero-derivation, redundancy rules, lexical transformations, or whatever. But this kind of proposal should make us uneasy, even for the cases it does handle. What would it mean to say that the "content"/"inscription" ambiguity of book was generated by rules of universal grammar? Many people take "universal grammar" to be innate. Does that mean, then, that knowledge of books is part of innate linguistic knowledge? Or suppose that universal grammar isn't innate. Then the universal psense-ambiguity of "book" is due to some universal feature of the learning condition--i.e., to the sameness of books wherever they are found. But in that case, why should such rules be made part of linguistic description at all, rather than being left to pragmatics?

Moreover, universal grammar can't be asked to account for all of the cases we're interested in, such as the two uses of newspaper, unless we suppose universal tacit knowledge of newspapers. Nor could this polysemy be the result of the application of some more general rule that gets us from linguistic entities to their promulgators; otherwise, why shouldn't book mean "book publisher," or sheet music, "sheet music publisher"? With similar cases in mind, Green

1974 argues that McCawley's universal sense-creating rules have to be made subject to language-specific constraints, which explain the existence of various "lexical gaps." The absence of an item book₂ "book publisher," or sheet music₂ "sheet music publisher," would thus be treated as a morphological accident of English. But we have already seen that there are problems here as well.

For one thing, livre and Buch pattern just like book, as opposed to journal and Zeitung. And, even more important, we note that the uses of polysemous words are mirrored in cases of what Quine 1971 calls "deferred ostension." The same "ambiguity" that inheres in "Hearst bought a newspaper" inheres in an utterance of "Hearst bought that," accompanied by ostension of a newspaper copy, while the same sentence cannot be used, pointing at a book, to identify a book publisher. Similarly, we can point at a chair-token and say "That was a fixture of eighteenth-century drawingrooms," and so on. Cases like these will be very important in our subsequent discussion of meaning and reference; for the moment, it will be sufficient to observe that it could only be the properties of the physical demonstratum, and the knowledge that the speakers bring to the context, that determine what things they can point at the demonstratum to identify. So the principles involved must be quite independent of the vagaries of morphology.

If a pragmatic strategy (still to be spelled out) can

get us from newspaper copies to newspaper companies, it would seem unnecessary to posit a linguistic rule to get from newspaper to the same referent. Suppose, then, that the lexical entry for newspaper made no mention of companies at all, and that we gave up our insistence that the lexicon stipulate all of the kinds of things that the word could be used to refer to. Then the syntax would not be troubled by a sentence like:

8. I read this newspaper, which Hearst bought last week.

The wh- word here could be assigned exactly the same reading as its antecedent, albeit its referent would be different. Thus in leaving out of the lexicon those uses of polysemous words that can be pragmatically derived, we achieve two economies. The lexicon is shorter, and requires only one kind of entry.⁴

But still more is at stake. Recall the second objection to syntactic treatments of polysemy: they are unrevealing, and even obscure important points. In each case, the device that's introduced, whether a redundancy rule, a transformation, or a disjoint entry, serves only one purpose: to map from things that are syntactically identical to sets of things that are semantically distinct; that is, from one form to several meanings. Now languages clearly have such devices, in the form of rules of morphology and word formation which mediate between morphs and words. We noted above Aronoff's 1976 example of the multivocality of the verb stand,

as in:

9. Just stand here for a while.

10. I can't stand him any longer.

On the basis of a regular alternation with a past tense stood, there is every reason for saying that 9 and 10 exemplify the same morpheme, but there seem to be no syntactic or semantic grounds for saying that they exemplify the same word. Similar conclusions can be drawn from examples like cranberry, or undergo (which has principle parts in -went, -gone), and most linguists have concluded, after their various fashions, that the notion "morpheme" can't be defined either syntactically or semantically, and that a set of rules (or even two sets of rules) must be set up to map from morphemes to words (or from morphemes to words and from words to derived lexical entries).

The necessity of having some such rules is unquestioned, but they contrast in important ways with the devices proposed to deal with polysemy. First, as we have noted, they are language-specific, and often unproductive. As such they are simply part of the facts about a given language, and have no synchronic explanation. The rules exist, in fact, only to facilitate the mapping from words to sounds, and to capture what seem to be significant generalizations about patterns of alternation (e.g. that stand₁ and stand₂ have identical past tenses). Thus the "morpheme" is a construct which can be justified only in terms of the canons of

evaluation within a particular theory of grammar: by introducing a morphemic level, we get a more economical linguistic description.

Syntactic treatments of polysemy accord the same status to the notion "lexical entry." They define a "lexical level," not isomorphic with the level at which interpretation takes place, in order to facilitate the statement of the conditions on syntactic rules. That is, rather than incorporating into the identity condition itself an enumeration of all the semantically distinct items which can undergo deletion when the controller is the word newspaper, we simply constrain it to operate on instances of the "same lexical item," and then incorporate into the item the different psenses that newspaper can have. The existence of a lexical level is justified by the resulting economy.⁵

But as we saw, the relation between the uses of newspaper is not like the relation between the two verbs stand: it can be predicted on the basis of semantic and pragmatic information, and doesn't have to be learned word by word. It follows, then, that the form of the lexical level is determined by universal considerations. This is not to say, as McCawley does, that there are universal rules which relate certain uses of certain words, but rather that uses are always lexically related when they are conceptually related in certain basic ways. Then we'll want to know, understandably, what it is about the language faculty in

general, and about the nature of syntactic operations in particular, that makes them responsive to just these relations. But when we give "syntactic identity" a technical definition, we make these questions hard to answer, and even hard to ask. All we have is a correlation between certain conceptual relations and certain kinds of lexical structures, with no way of explaining why these relations should be reflected as they are in the model.

On the other hand, if the different uses of polysemous words are not represented in the lexicon, we won't have to take syntactic identity as a theoretical construct; we can say that a necessary condition for identity is that the items involved must have the same lexical content. The explanation for the fact that syntactic rules ignore polysemy is then simple and straightforward: the syntax responds only to linguistic distinctions, and these psense-differences are not linguistic. The syntax treats newspaper₁ and newspaper₂ as the same word because they are the same word.

There is much more to be said about the interaction of syntax and polysemy. For one thing, while semantic identity is a necessary condition for syntactic identity, it is not sufficient. In addition to whatever purely syntactic considerations may play a role in determining identity, there are many cases of polysemy which syntactic operations can not ignore, such as when one psense is

conventionalized ("foot of a mountain"/"foot of a person"). Ultimately we'll be able to explain all of these cases without having to complicate the identity condition, by looking at the functions of the relevant rules and constructions. For now, it's enough to have outlined the savings to syntactic metatheory that a pragmatic treatment of polysemy makes possible.

J. Katz (personal communication) has suggested to me another position that could be taken on the syntactic problems that arise when polysemy is treated in the grammar. Suppose we listed all of the uses of polysemous words separately, without recourse to disjoint entries or other such devices. Then all of the sentences that prompted the syntactic treatments under discussion would be labeled ungrammatical; no linguistic distinction would be drawn between sentences like:

11. The newspaper you're reading has come out against gun control.
12. Yeats disliked hearing himself read in an English accent.

and

13. The bank which overflowed is giving away free toasters with new accounts.

The obvious interpretations that can be assigned to 11 and 12, then, would be generated by the theory of "semi-sentences" (see Katz 1964, Chomsky 1964), a part of pragmatics, which assigns interpretations to grammatically deviant strings.

This line is unsatisfactory on a number of grounds, but not all of them are crucial. For one thing, 11 and 12 certainly don't seem to be deviant; not at least like "John is probable to win" or "John wants that Mary will go," deviant but interpretable strings that might be candidates for such a treatment. But a discussion of this question would take us too far afield. So let's suppose either that we could find a group of particularly literal-minded speakers who would reject these sentences out of hand, or that speaker judgments that sentences are acceptable are not a sufficient condition for labeling them grammatical. (See, for example, Langendoen and Bever 1973.)

It could also be argued that there is a certain inconsistency involved in letting pragmatic schemata generate the interpretations of 11 and 12, while insisting that the two interpretations of a sentence like "Hearst bought a newspaper" must be assigned by the semantics, since the same schemata which relate the two uses of an item in the former cases would be equally available to generate those uses, as we have seen. Why should our intuition that a use is semantically non-deviant be an assurance that it is generated by the semantic component, if the judgment that an analogous syntactic construction is acceptable is not an assurance of grammaticality? But I won't pursue this argument either.

Katz's suggestion is most interesting insofar as it forces us to look for non-syntactic arguments against giving

semantic representations to each of the uses of polysemous words. These arguments have basically the same form as the arguments I presented earlier against syntactic treatments: a semantic treatment of polysemy, even when the syntactic evidence is ignored, is both uneconomical and unrevealing.

The economies are easy to see: a lexicon that lists only one meaning for newspaper is simpler than one that lists several. However new psenses are generated, they clearly require the introduction of some new entities. It is true that the economies are not on the same scale as those that we gain, say, by introducing transformations into a grammar; the difference between a lexicon that lists only one meaning for each item and one that lists all its uses, seems to be on the order of the difference between a pocket dictionary and the OED. But the simplicity arguments here are much easier to see than they are in syntax, where the evaluation metric is a function of the theory itself, either in principle (see, e.g. Chomsky 1965), or at least in practice.

Simplicity arguments in syntax are usually used to defend the introduction into the grammar of one kind of device (say, interpretive rules) over another (say derivational constraints) where there is no extralinguistic reason for preferring one or the other. But when we are trying to choose between a pragmatic and semantic explanation, it isn't a question of "either/or," but of "one or

both." If the pragmatic account is well-motivated--if it is based on empirically valid assumptions about language users--then it stands quite independent of any assumption about the language faculty. Saying that newspaper has two lexical entries doesn't free us from having to say that speakers know certain facts about the way newspapers are published; that is determined on independent grounds. But a speaker who knows such facts would be able to generate all of the uses of newspaper on the basis of one meaning; he would have no reason to introduce a second item into his lexicon.

But more than just economy is at stake here as well. In asking the semantics to represent all non-deviant word-uses, Katz and Fodor are following the standard practice of syntacticians. But there are important differences between cases of syntactic and lexical deviance. True, a speaker constantly hears examples of both, in poetical language, in the speech of non-natives and children, and in plain speech errors. But his reasons for judging these instances deviant cannot be analogous. Consider the process whereby a language learner comes to rule out certain utterances from the corpus on which his grammar is built.

A child of six who is confronted by a sentence like "Not is John tall" would calculate the changes that would have to be wrought in his grammar if it were asked to generate the sentence, and on finding them extensive, would count

the sentence aberrant, at least pending much more evidence. By contrast, there can never be any trauma to the lexicon in adding a new word to it; that part of language learning continues past any critical stage. Why, then, should we ever judge that a word-use is deviant? For example, if I read Byron's "The castled crag of Drachenfels frowns on the wide and winding Rhine," why shouldn't I simply assume that in nineteenth-century English, frown had an ordinary use to mean "overlook, overhang," rather than judging the use a "deviant" metaphor? Not, surely because I would have any trouble in adding a new item "frown₂" (obs.) to my lexicon, and not, either, because the use is rare. We do not have to hear a word more than once in order to learn it, provided the use we hear makes the sense clear; and besides, we constantly encounter novel uses in old texts, and we are not tempted to regard all of them as metaphors, even when the texts are poems. And frequent repetition does not necessarily rob a use of its metaphorical character, though its "effect" may be blunted. People have been calling each other "dogs" and "cows" for millennia.

We couldn't say, either, that we judge the use to be metaphorical because we can see how it would be generated in context by pragmatic schemata; because we don't have to add a new item to account for it. For what about the uses that are at issue here, such as with newspaper and chicken? They too are generated by pragmatic schemata, but they are

not judged "metaphorical," or at least, "deviant," even on first hearing (an assumption supported by the ease with which we can judge which now-obsolete uses must have been conventional for Byron). How, then, are the two cases distinguished, and why should only one of them be entered in the lexicon?

If the relevant uses of newspaper and frown are both generated by pragmatic schemata, and if criteria of frequency and setting are not sufficient to tease them apart, then the difference must lie in the particular schemata that generate them. It can't be that the schemata are themselves different in form; that they constitute different kinds of pragmatic "rules"--the fact that the same use can change from metaphorical to standard shows us that. Rather, there has to be some difference in context; specifically, in the kinds of beliefs that license one or the other use, and in the degree to which a speaker may reasonably expect these beliefs to be shared by other members of the speech community.

Until we have looked at these pragmatic schemata in detail, I can only argue from suggestive examples. Newspaper is standardly used to refer to newspaper companies because the use is consistent with widespread beliefs about how newspapers are; frown can be used to refer to the disposition of natural objects only against a background of beliefs about nature that most speakers wouldn't count as part of the set of assumptions that underlie ordinary discourse, though

they can readily accept these beliefs as a basis for the understanding of Romantic poetry. Or take another sort of example. The same process whereby chicken can be used to refer to chicken meat allows us to use any count noun to refer to a substance, provided that it is generally accepted, among other things, that the same kind of substance is characteristic of all and only the members of the class of things designated by the noun. Whether or not a particular count noun can be used to refer to a substance depends then upon what is commonly believed about the designata of the noun, and upon what kind of substance is being referred to. We use chicken and fish to kinds of meat, but not, standardly to kinds of skin; reptile, on the other hand, can be used to refer to a kind of skin, but not, standardly to kinds of meat; water buffalo can be used to refer to either; insect to neither. We commonly talk about eating liver, but not leg or wing. Yet all of these uses reflect matters of taste, in one way or another. We don't standardly use cigar to refer to tobacco (other things are made of it), or mattress to refer to bed-ticking (they can be made of other stuff); but we can use kleenex to refer either to the things or the kind of paper they're made of. Find a new use for that paper, however, or ban pipes and cigarettes, and our judgments about "deviance" will change.

If we make the lexicon responsible to record all "standard" uses, then the grammar becomes something more than the

"theory of the language"--it becomes a theory both of the language and of the uses to which the language may rationally be put in a given community. As we saw above, there can't be any argument that this line is not possible in theory, provided we are willing to forego economy. We can add as many lexical items and redundancy rules as we like without changing the lexical system, if we disregard syntactic problems. But what is gained? We are no closer to an explanation of why some word-uses are "standard"--all we have is a list of them. Nor do we come any closer to an understanding of the nature of the language faculty; if anything, we make it more obscure, in failing to distinguish what is arbitrary from what is not.

I have no further arguments to offer; from here on, I will simply assume that not all of the "acceptable" uses of polysemous words are represented in the lexicon. By itself, of course, this is not inconsistent with the idea of an autonomous semantics; it merely shows that intuitions for "acceptability" are not the grounds on which the semantic theory of a natural language ought to be built. I have not shown that we couldn't do with a "reduced" semantics, in which only one use of each polysemous word was determined by its semantic representation. Thus newspaper could have only one meaning, which gets us in one way or another to newspaper copies, and this meaning would be the input to schemata which determined what sorts of things it could be used to

refer to.

The argument against this second, weakened position can be made, but it must be built up by stages, and we will not come to it until the end of Chapter Two. More immediately, we have an IOU to pay off: it still remains to give an explicit account of the pragmatic schemata that make extended reference possible. In the following section, we will set up the background against which the account will be given, and enter into the account proper at the beginning of Chapter Two, which follows.

1.5 Polysemy and Ostension

We noted earlier that many of the patterns of polysemy that we find with descriptive terms are mirrored in ostension; that an utterance of 1, accompanied by ostension of a newspaper copy, may have the same set of interpretations as 2:

1. Hearst bought that.
2. Hearst bought a newspaper.

These cases of deferred ostension are important to us for two reasons. First, as we noted, they show that the schemata that generate multiple word-uses could not be purely linguistic, since there is no term in 1 in which any ambiguity could inhere. Rather, the indeterminacy of 1 must be due to the pragmatic strategies whereby the speaker and hearer determine the referents of ostensive terms.

Let me belabor this second point. Suppose we wanted to treat 1 as semantically ambiguous, and that we had

already available a disjoint-entry analysis of polysemy.

Then we could establish a convention of reference as follows:

3. a demonstrative term can be used to refer either to the physically present demonstratum a that is salient in the context, or to any object b such that b is named by some term t which is listed in the lexicon disjointly with some term t' that names a.

But apart from the perversity of introducing a rule of this peculiar and unmotivated form (and we fare no better under other semantic treatments of polysemy), and apart from the undesirability of treating polysemy in the grammar in the first place, 3 won't work. Consider a sentence like 4, as spoken by a restaurant waiter, together with 5, spoken under the same circumstances and accompanied by ostension of a sandwich:

4. The ham sandwich is sitting at table 20.

5. He is sitting at table 20.

Obviously, 4 is an example of some sort of metonymy, and could be sloughed off as deviant. But 5 is not deviant, yet 3 could not provide us with a referent; the lexical entry for ham sandwich will have nothing to say about persons. In fact, I can point at a newspaper and say 5 to indicate an indefinite number of persons: the editor, the publisher, a reporter, or the person who has just left the paper sitting on a chair. And 3 will enable me to identify none of these referents. So we will require a pragmatic theory to deal with 5 in any event, a theory which will

equally well account for the newspaper cases. But in that case, why bother at all with a rule like 3?

Let's assume, then, that the indeterminacy of sentences like "Hearst bought that" abides in pragmatic factors. Then we begin to address the second point of interest in deferred ostention: by looking at utterances in which semantics could not play a role in creating indeterminacies, we can investigate the pragmatics of reference without interference. We can imagine three sorts of problems that could arise in identifying the referents of demonstrative terms. First, there could be a difficulty in saying which of several possible objects the speaker intends us to identify--say if we know him to be astygmatic, or if his demonstrative gesture is too vague. Problems of this sort are discussed by Kaplan (ms.) but we will ignore them, and assume a context in which the demonstratum has been satisfactorily differentiated from other objects on the scene. But even then, there could be some question as to which of a set of physically coterminous things was to be taken as the referent of the demonstrative term. For example, suppose I point at my watch and say:

6. This once sold for \$50.

where both "type" and "token" interpretations are possible. Then we might want to say that both referents are physically present, and that their instantiations on the scene are (as it happens) co-extensive; the only indeterminacy involves

which of them is being pointed at. The situation is rather the analogue of pointing at the Statue of Liberty from the rail of an incoming ship and saying:

7. That is a heavily populated region.

where the intended demonstratum can be construed as any of a number of entities of which the Statue of Liberty is the only visible part: Bedloe's Island, New York City, the Northeast, and so on. In this case the sentence is not ambiguous, but neither is it problematic; corresponding to each of its interpretations is a different demonstratum.

Goodman 1968 suggests an account of deferred ostension very like this, in talking about "exemplification." He notes that under certain circumstances, a demonstratum may be taken as a "sample" of something else, as when one points at a tailor's swatch to identify a bolt of cloth, or the set of all pieces of cloth that have some property or set of properties in common with the swatch. In order for a particular a to serve as a sample of b, Goodman suggests, two conditions must be satisfied. First, there must be some convention whereby a and b must both be denoted by some property "p"; that is, a can be a sample of "whatever is tweed" just in case a is tweed. Which means that some part of the referent must be physically present, and picked out by the demonstrative term. Now the first condition--that the reference be "conventional"--is just another way of saying that it is semantic, or at least arbitrary. And it

surely fails in the majority of cases. Swatches of cloth, as it happens, do function more-or-less conventionally in this way, but it isn't at all clear that the relevant conventions are linguistic. In any event, there is surely no convention whereby the Statue of Liberty refers to the North-east, or a copy of The New York Times to its editors. (Goodman introduces this condition in order to explain the fact that a swatch of cloth can't be used to refer to some extensions of which it may be a part, such as "all cloth finished on November 3." But this has nothing to do with convention, as we'll see; in fact, under appropriate contextual conditions, a swatch can be so used.)

Goodman's second condition--that (at least some part of) the intended referent must be physically present, is equally unsatisfactory. When I utter 1, a newspaper company is in no way physically present; a copy of The Times and the Times Company are simply not in the same extensional set. This is even more obvious with examples like 5--"he is at table 20"--where I point at a sandwich to identify a person. In the same way, you can point at a book to identify the film rights to it, saying,

8. MGM bought that last week.⁶

So there is no way to argue that the indeterminacy of sentences of these sorts is simply due to an uncertainty over which of two objects is being pointed at.⁷

At this point, however, someone might say, "Well, all

you've shown is that the referents of the demonstrative term don't have to be physically present. But a referent could nonetheless be "salient" in a context, and a demonstrative gesture could be used to pick out the referent by drawing attention to the particular feature of the context in which it is salient. In which case we are still dealing simply with a question of which of several things that are implicitly salient in the context of utterance is to be taken as the referent of the demonstrative term. This is roughly what Quine 1971 says about these cases (as when we point at the gas gauge to show that there is gasoline). Quine suggests that "deferred ostension occurs very naturally when, as in the case of the gasoline gauge, we have a correspondence in mind."

Now depending on how "salience" or "a correspondence in mind," is defined, I have no quarrel with any of this. But consider its consequences. What is claimed is that a newspaper copy somehow carries along with it a newspaper company--the presence of the former ensures the salience of the latter. So that whenever you can point to a newspaper to identify it, you can also point to it to identify the newspaper company that publishes it. All that is required is that it be evident to the hearer that the physically present demonstratum is a newspaper, and that the hearer know certain things about the way in which newspapers are published. Lewis 1972, for example, has suggested that the

uses of definite descriptions require something like a "prominent objects coordinate," which functions like other indexicals. He notes that an object may be prominent "because it is nearby, or pointed at, or mentioned." Moreover, he allows that the coordinate "will be determined, on a given occasion of utterance of a sentence, by mental factors such as the speaker's expectations regarding the things he is likely to bring to the attention of his audience." It wouldn't be hard, I think, to accommodate this to the notion of "salience" we are using.

But it is important to note that the possible referents of the demonstrative term are not "salient" independent of one another. A newspaper company is prominent only when a newspaper is, and an apple-type only when an apple is. In either case, you can get to the first object only via the second. (And of course, it is the position of the physically present demonstratum that determines whether a "proximal" or "remote" deictic this or that will be used.)

At this point, however, we can ask for a further explication of the notion "salience." Why is a book publisher less "salient" in a book copy than a newspaper publisher in a newspaper copy? Given a physical demonstratum a, just what things can you use it to identify ostensively, and why? What we require now is a theory that explains how it is that knowing the identity of some things enables us to identify others. Obviously this theory won't be a part of grammar--

it will have to take into account a variety of extralinguistic factors, such as the speaker's knowledge of the demonstratum (for example, the property of being published in such-and-such a way) and of the context (such as a given newspaper's property of having been deposited in a chair by Jones). In short, it will be a pragmatic theory of ostensive reference, and by invoking it, we'll be able to explain the indeterminacy of 1--"Hearst bought that"--without recourse to "ambiguity."

But now let's return to 2--"Hearst bought a newspaper." Suppose we have a satisfactory theory of ostensive reference, which tells us when and why we can use a newspaper copy to refer to a newspaper company. We should expect to have relatively little difficulty in modifying our account of ostension to deal with description as well, so that we can explain how the word newspaper can be used to the same end. Then we will have an account of reference, tout court.

This theory will explain, not only the "standard" example of polysemy we have been talking about, but other sorts of cases, which have been equally vexing to semantic theories. For example, "use/mention" indeterminacies are present in ostension as well as description. Alongside of sentences like 9 and 10:

9. "Newspaper" begins with a nasal.
10. "Newspaper" is practically the same word in French and Italian.

we have examples like 11 and 12, accompanied by ostension of a newspaper copy:

11. What letter does this begin with, Johnny?

12. Is this the same word in French and Italian?

In the same way, the pragmatic account of reference will resolve for us problems of "representations," as when we use the phrase "The president of the college" to refer to a painting of the president of the college; note that we can point to a college president and say:

13. He has blue eyes in John's portrait.

to the same effect.

Section 2.1 will be devoted to developing a schematic account of ostensive reference, and 2.2 to its extension to descriptions.

NOTES--CHAPTER ONE

1. The lexicon may of course contain other features and devices, such as word-formation rules. We will have call to mention these again in the following section, but we will not be concerned here with the purely formal properties of signs.
2. We can disregard here the question of whether there may be unacceptable grammatical sentences, as argued by Langendoen and Bever 1973. We will see below that the lexicon could not provide us with any cases analogous to theirs.
3. Why such associations should exist is a fascinating problem for cognitive psychology. The use of perception verbs like see and hear to refer to mental states (cf. French s'entendre bien avec quelq'un, American slang I hear you) seems a universal or near-universal process.
4. Though it should be noted that the interaction of syntactic processes with the processes that produce polysemy are a good deal more complicated than we have heretofore allowed. This is the thrust of Borkin 1972, who gives a number of examples in which deletion or anaphor may be blocked across psense-differences, depending on the kind of rule involved, and the syntactic configuration on which it operates. For example, she presents contrasts like those in i-ii and iii-iv:
 - i. Mailer does not like hearing himself read aloud.
 - ii. ?Mailer will be reading himself tonight.
 - iii. Homer, who was a blind poet, fills two feet of my shelves.
 - iv. ?Homer, which fills two feet of my shelves, was a blind poet.

More generally, George Lakoff has pointed out to me that some anaphoric processes, such as the so does construction, may not operate when the deleted (or interpreted) NP is not identical in psense to the controller NP; thus v does not permit a crossed reading.

- v. John bought a newspaper and so did Hearst.

We could account for all such problematic cases only

in the light of a comprehensive theory of anaphoric processes, which is obviously beyond the scope of this thesis. It is my feeling, however, that explanations will fall out into two categories. Alongside of ii, for example, we have the acceptable v; which suggests that the badness of ii is due to some fairly low-level pragmatic interference:

- vi. Mailer will be reading himself, among many others, at the Arts Center tomorrow night.

On the other hand, there seem to be no such exceptions to the so do generalization; here we may have to look to a more general principle of anaphora. It may be that we will be able to account for the absence of the crossed reading pragmatically, perhaps by application of the "identity condition" on extended reference, to be presented in section 2.1. Or we may have to build a restriction into the syntax or semantics of this and other constructions. Ultimately, however, problems like these will not force any revision of the position on polysemy argued here; what is critical is that there are some rules that ignore psense-differences, not that there are some that may not.

5. This argument would have to be reformulated somewhat to deal with deletion analyses of polysemy, but the conclusion would be unaltered. Note that on a view of the lexicon taken in "classical" generative semantics, we would have to allow that there were two kinds of pre-lexical transformations, which behaved quite differently with respect to later syntactic operations.
6. Actually, you can point at a book on whose jacket the name of a publisher is prominently displayed and say, for example, "They merged with Scribner's last week." See below.
7. In his later discussion of symbolic systems, Goodman makes a number of observations that are much more relevant to the problem of displaced reference than we have at hand. But these observations are not tied directly to ostension, and I have not followed the parallel in what follows.

CHAPTER TWO

THE MECHANICS OF REFERENCE

2.1 Ostensive Reference

Let us take reference to be an act in which a speaker uses signs to enable a hearer to identify something. (Though this is not a sufficient condition for referring. The speaker must intend to perform the act, and the hearer must recognize this intention, and so forth. See Searle 1969; for the most part, we will not be concerned here with those aspects of the theory of reference that constitute a part of the more general theory of speech acts.) A complete account of the referring uses of demonstrative terms will involve the invocation of two theories, as we have seen. One will explain how a given use of a term enables the hearer to identify a physically present demonstratum; it will enable us to derive a function that takes the term as its argument, and the demonstratum as its value. We will assume that this theory constitutes a part of the more general theory of meaning, but we will not try to justify this usage until we come to the discussion of meaning in Chapter Three. The other theory will explain how the hearer's knowledge of the identity of the demonstratum enables him to identify the referent of the referring term;

it will enable us to derive a function that takes the demonstratum as its argument, and the referent as its value. This second theory will be a part of the more general theory of reference, and it will be concerned with relations between persons and things, not with words or meanings.

Let us say that the hearer has been enabled to identify something when he can give a description that is true of it and false of everything else.¹ In ostension, this requirement is satisfied when the hearer is made to know that the intended referent stands uniquely in a certain relation to the demonstratum; i.e., that the intended referent is the value of a particular function on the demonstratum. We will call this function the referring function (RF) of a given use of a term, and say that an act of ostensive reference is successful just when the hearer can identify the referring function.

It will be convenient to talk about reference in this way, for we can then formulate set-theoretically the schemata that the hearer brings to bear in determining the referent of a use of a demonstrative term. Before proceeding, however, we should stop to disclaim any psychological reality for schemata presented in this way. Obviously, functions are not the sorts of things that are customarily

regarded as objects of knowledge; the best we can say is that these schemata are reasonable models of the assumptions that speakers and hearers might actually make.

This modeling procedure is quite familiar to linguists, especially syntacticians, who are accustomed to asking questions about the "psychological reality" of objects like rules and syntactic categories, without worrying too much about whether the speaker is, or could be aware of knowing these objects. In doing pragmatics, however, we have to be very careful that models are interpretable in terms of concepts that we know to be available to speakers, because we have to justify the models, not only by showing that they are economical and descriptively adequate, but also by invoking plausible hypotheses about what we would expect a rational speaker to do.² (This is not to say that the working out schemata are less a part of "tacit" knowledge than the rules of syntax. Speakers certainly aren't aware of the processes whereby they interpret, say, some questions as commands. The most we can ask is, "What would be a rational course of action in situation S if we were aware of all of such-and-such facts?" If this then is what speakers do in that situation, we will say that they have behaved rationally, whatever their actual mental state may have been.) So we will want to be able to interpret "knowledge of functions" as propositional knowledge.

Let us assume that all speakers come armed with a finite repertory of prime (or basic) functions, such as "type of," "cause of," "possessor of," and so on, defined over "natural" ranges and domains. We need not say that knowledge of such functions is innate, or even universal; it suffices that they shall be commonly known to be familiar to most members of the speech-community. Then given a demonstratum a, and a prime given function *F*, hearer can derive a referring function from *F* when he can assume that *F* yields a unique value at a (other conditions will be added below). Prime functions will also allow the derivation of composite referring functions, such as "source of type of"; these are used in referring as well, when they are uniquely satisfied at the point picked out by the demonstrative term. It follows that for any demonstratum, there will be an infinite number of functions, prime or composite, from among which the hearer can derive a referring function; for the rest of this section, we will be concerned to give the schemata whereby this number can be winnowed down.

Let's take the set *G* to be the (finite) set of all prime given functions and the set *G'* to be the (infinite) set of all given functions and composites of given functions. Then the job of the working-out schemata will be to enable the hearer to pick out one member of *G'* for a given use of a demonstrative term. In the first part of

this section, I'll describe and motivate several conditions on reference, which will tell us when a member of G' can be used as the basis for an RF.²

These conditions can be addressed in two ways. First, they make reference possible, and to this extent they can be stated as absolute constraints. I will be talking about them in this way throughout the first part of this section; saying, for example, that "given a demonstratum a, and an intended referent b, F allows derivation of an RF only if C." But in practice, the degree to which a condition of reference is satisfied will depend upon a variety of factors; and given a set of possible RF's, all of which satisfy absolutely a given set of conditions, some will be "better" than others, in that their use will increase the probability of successful reference. In the second part of this section, I'll be talking about the conditions of reference in a relative way, as factors that determine which of a set of "possible" RF's is most likely to be used; this discussion will be important when we come to talk about the lexicon later on.

All of the pragmatic conditions on reference arise out of a single principle: the value of the RF at the point picked out by the demonstrative gesture has to be determinable. (Later, we will say that the "goodness" of the RF depends on the degree to which the value is determinable at the demonstratum.) Which is only to say that the RF has to

be such that, given the argument, the hearer is enabled to pick out the value. This will be the case only when several sub-conditions have been satisfied. One consideration that constrains the derivation of the RF is obvious: a function can be used in referring only if it yields values at the demonstratum that are members of the "range of reference" determined by the nature of the predication, and by the conversational context. We won't talk about "co-occurrence restrictions" here; for the moment, let's simply accept that, given an utterance of "Hearst bought that for \$50 million dollars," accompanied by ostension of a newspaper copy, we know on the basis of the predication that the RF must be some function that takes us from newspapers to expensive things. And in the same way, given an utterance of "he is a conservative," we know that RF has to take us from newspapers to male persons, in virtue of the inflection of the pronoun. And finally, that in the midst of a conversation about newspapers, it is more reasonable to expect that a speaker will be referring to W. R. Hearst than to Chester A. Arthur. In each of these cases, the range of RF has to be defined over the intersection of the natural range of some given function and the range of reference. Given "he is a conservative" for instance, accompanied by ostension of a newspaper copy, we will derive a function from newspapers to male publishers from the given function "x is published by y," which takes simply persons, or groups

of persons, as its range.

The corollary of this condition is that the demonstratum has to be manifestly in the domain of the given function. These conditions can be stated as Ia and IIa. Given a demonstratum a, and an intended referent b, a member of the range of reference B, a given function $F: X \rightarrow Y$ allows derivation of an RF only if:

Ia. The hearer would be expected to know that $a \in X$ (and would know that he was expected to know, and so on; hereafter I won't bother to mention the Gricean regressus in stating these conditions).

IIa. The hearer would be expected to know that $Y \cap B \neq \emptyset$.

I'll take up these two conditions in turn: first, when can we say that Ia is satisfied? Obviously just when it is "manifest" that $a \in X$. But the fact that $a \in X$ may be manifest in any of several ways. X might be defined in terms of the "intrinsically" manifest properties of a. If a is a newspaper, for example, and has not been rolled up to look like a log, or turned into a papier-maché puppet, then the RF can be derived from any function that contains the set of newspapers in its domain (provided that the other conditions on the choice of the RF are satisfied). Or, X might be defined in terms of properties that are "extrinsically" manifest in a. A ham sandwich will be manifestly in the domain of the function "x was ordered by y" only with respect to certain contexts of utterance. And finally it may simply have been stipulated that $a \in X$; the

newspaper-log can be used in ostensive identification when it has already been identified as a newspaper. (In which case it must be manifestly identical to the individual of which such a stipulation has been made.) But nothing of importance turns on these distinctions; it is immaterial how you come to know that $a \in X$.³ When I say that membership in the domain has to be "manifest" I mean only that a has to be such that a rational speaker would expect his hearer to know that $a \in X$, on whatever grounds. And this depends on some notion of "total context"--the beliefs of the speaker and hearer as well as the physical properties of the demonstratum, contribute to the satisfaction of Ia. For example, I might point at a half-dollar and say,

1. That's how much they get for a subway ride now.
and I would probably present no more problems if the demonstratum were a 500 lira note; it requires only that the hearer be able to recognize that the demonstratum is in the domain of the relation "x has a (fixed) monetary value of y," or some such. But now suppose I point at a large stone wheel, and utter 1; successful reference here depends on my hearer's ability to recognize that the demonstratum is a Yap fifty-cent piece, which in turn depends on his knowing that he knows that I know this about him, (and so on), so that he can apprehend my intention to use this function in referring. So "manifest" winds up with "relevant" and "recognize," one of those inevitable pragmatic primitives,

beached on the shores of epistemology.

However, Ia is clearly inadequate by itself. Consider the given function "x was built in y" which, for example, takes cars (among other things) into years. This allows derivation of a good RF in some cases--I can point at a 1957 Chrysler with Batmobile tailfins and say,

2. We were in Toledo then.

But I would clearly have difficulty in using 2 with a Volkswagen. Similarly, the function "x was fermented in y" takes both beers and wines into years, but it can't be used with beers to identify a year. Yet both the VW and a glass of beer are manifestly in the domain of the function. The problem, of course, is that there is in general no way for non-specialists to distinguish a 1973 model VW from a 1975 model, or a 1974 Budweiser from a 1975. In which case there is no way to determine the value of the RF for these arguments. So we will add a second condition on the domain of the RF, which will ensure that, for any given function F that satisfies Ia and b for a certain context of utterance, we can derive an RF only if we can tell which object in the range is the value of f on the demonstratum:

Ib. Given a demonstratum a , and a range of reference B , and a given function $F: X \rightarrow Y$ such that $Y \cap B \neq \emptyset$, and an intended referent $b \in B \cap Y$, F allows derivation of an RF only if the set A of all values of $F^{-1}(b)$ is manifestly discriminable from the set of things that are not values of $F^{-1}(b)$.

Let's call Ib the "Inverse Image Condition." What it says

is simply that, for any function on any demonstratum, you have to be able to distinguish that demonstratum from other things like it for which the same function has different values; if you can't, there is no way to say what value the function yields on this demonstratum. (Note that Ib entails Ia: we can determine whether or not a falls in a given subset of X only if we can determine whether $a \in X$.)

The inverse image condition has widespread application. For example, we noted earlier that book copies, unlike newspaper copies, could not be used in ostensive identification of publishers. This follows from Ic; for any publishing company, it is difficult to determine that a given book is in the inverse image of the function from copies to publishers. While we assume that there will be manifest resemblances among all the newspapers put out by the New York Times Company, which enable us to distinguish them from the newspapers published by the Washington Post Company, we make no such assumptions about the books published by Scribners' and Macmillan. (To be sure "manifest" is back again, this time in a counter-factual context. To say that it has to be determinable whether something is in the inverse image set for a given value is the same as saying that, say, given any copy of a newspaper it would be manifest to us that this newspaper either was, or wasn't, a New York Times. But the problems created by counter-factuals are beyond our concerns; and we cannot introduce any

notion of speaker expectations without using them.)

The important point is that here, as with Ia, we determine whether the conditions of reference are satisfied in the light of a "total context," which includes the beliefs of the speaker and hearer. Some books do allow identification of publishers--books in the Everyman Series, for example, or Little Golden Books--precisely because similarities in format do allow subsets in the domain of the RF to be distinguished. Often, the question of whether reference to a publisher can be brought off may depend solely on how prominently the name of the publisher is displayed on the dust jacket. I don't mean to suggest that books and newspapers are distinguished absolutely in this regard, but only that the likelihood of successful reference to a publisher is greater with a newspaper than a book.

Let me give a couple of other examples of the way in which the Inverse Image Condition constrains the choice of the RF. We noted before that the relation "x was printed on paper made by y" doesn't allow derivation of a "good" RF from newspaper copies to corporations. And this follows from the fact that for ordinary speakers, newsprint made by one company is pretty much the same as newsprint made by another; sub-sets of the domain of the function can't be distinguished. Or consider the function from equivalence classes of Chevrolets to the factories in which they were built--here again, it would not be satisfied, excepting for specialists.

Condition Ia, IIa, and Ib enable us to explain how the RF is derived from a given function (or composite of given functions). Starting with a function $F:X \rightarrow Y$ a demonstratum $a \in X$, and a range of reference B such that $B \cap Y = \emptyset$, we will derive the referring function f_r by restricting the range of F to the intersection of $B \cap Y$, and the domain of F to the largest subset A of X such that $a \in A$, and such that A contains every element y in X such that if $a \neq y$ and both $f(a)$ and $f(y)$ are in $B \cap Y$, and $f(a) \neq f(y)$, then a is discriminable from y . Only when we can derive a function of this form from a given function, or composite of given functions, do we have a possible referring function.

We can turn now to the second problem that the working-out schemata are required to handle: how do we choose one from the (potentially infinite) set of possible RF's in a given act of reference? It will be easiest to come at this problem obliquely, starting with a special case: an absolute condition on reference involving a special function, the identity. Consider the following cases, in which a possible f_r --a function from demonstratum to range of reference that satisfies Ib--cannot be used in referring. Suppose we are talking about the virtues of monarchy, and I point at George IV and say, "He went mad" to mean "George III went mad." The functions from sons to fathers satisfies condition Ib: the values of the inverse

(sons) are quite easily distinguished from one another. But the function is clearly impossible as an RF; in context, he can refer only to George IV himself. The second example is a variant of the first: suppose we are talking about nineteenth-century laws, and I point at a copy of Bleak House and say "He was sent to debtor's prison," to mean Charles Dickens' father, John. Here again, Ib is satisfied, but reference fails. In either case, to determine the value of the function for the argument, you have to determine the identity of some other object which is itself a possible referent: George III in one case, and Charles Dickens in the other. Obviously, there has to be some principle by which such functions are disallowed; otherwise, you could point at anything to identify anything else in the same set. So there has to be a condition that says that the only function that can be used to map from one member of a set to another member of the same set is the identity. Thus, if the physical demonstratum is itself a member of the range, it has to be the referent--if I point at a book and say "This is made of paper" in an ordinary context, this can refer only to the book itself. It's hard to see how the condition could rationally be otherwise--the identity is the only function whose value is trivially computable for all domains.⁴ And if the identity cannot be ruled out as an RF, then allowing any other function to the same range that yields a different value would

make it impossible to pick out a unique referent, which is the whole point. Condition II formalizes this, and extends it to the case of composite functions:

- II. (Strong form) Given a demonstratum $\underline{a}$, and an intended referent $b \in B'$, and given a possible referring function $\underline{f}_r$ such that $a \in A$ and $\underline{f}_r \rightarrow B'$, $\underline{f}_r$ can be used in referring only if:
- a. if $\underline{f}_r$ is a prime function, and $A \cap B' \neq \emptyset$, then $\underline{f} = I$.
 - b. if $\underline{f}_r$ is a composite function $g \circ h$ such that $h:A \rightarrow C$ and $g:C \rightarrow B'$; and $C \cap B' \neq \emptyset$, then $h = I$. (i.e. $g \circ h$ cannot then be a non-trivial composite.)

(This "strong formulation" of II will have to be modified to handle Donnellan's "referential" uses of demonstrative terms; see below.)

What II says is simply that when the demonstratum itself could be the referent, it must be. And if the value of a function $\underline{h}$ on the demonstratum could be the referent, then it must be the referent if $\underline{h}$ is to be used in referring; no function can operate on the output of $h(\underline{a})$ that yields a different value. In other words, once you have computed a value that falls within the range of reference, you have to stop.

The formulation of condition II, which I'll call the "Identity Condition," suggests several questions: First, how do we know whether a function is prime or composite? Obviously any prime function can be written as a composite function, and vice-versa. But we are claiming psychological reality for these functions, so that we start with a finite

set of prime given functions from which all composite functions must be derived. In most cases, we will be able to decide intuitively when a function is composite; the functions "father of," "author of," and "publisher of," for example, will be primes; the functions "father of the author of" and "president of the publisher of" will be composite. But I don't mean to suggest that all functions must have a monomorphemic lexical realization, or that all lexically represented relations are primes. The status of "uncle of," for instance, is tricky; ultimately, we'll have to look to the kinship system of a culture to determine whether such relations are effable only as composite functions. There are cultures in which the father's eldest brother takes the child from birth, and the child bears the uncle's name, and so forth, so that for a given child, the function to "uncles" may be "better" than the function to father, and representable non-compositely. But I'll talk about all this later. For present purposes, the decision that a given function is or isn't composite doesn't much affect the way II works. If for example, we should decide that "father of the author of" is cognitively representable as a single function "fauthor of," then II will not apply; rather, we'll go to other strategies which select the "best" of two functions to a set from the same argument, and these will tell us which of the two functions--"fauthor" or "father"--to choose. So II applies where it applies, and where it doesn't, there are

other ways to choose the RF; see below.

I mentioned earlier that the formulation of condition II is different according as reference is "attributive" or "referential." Recall that Donnellan introduced these terms to distinguish the two interpretations of statements like "the man who murdered Jones is insane." On the attributive interpretation, the intended referent is roughly "whatever x is such that x murdered Jones is insane." On the referential interpretation, the description is used to pick out a certain individual a, such that the speaker believes that a murdered Jones; to get the interpretation in these cases we would substitute an individual constant: i.e., "a is insane."

It has been suggested (as by Cole 1976 and Partee 1972) that this distinction is semantic; Donnellan himself, as well as Hintikka 1973 prefer to treat it pragmatically. Hintikka argues that the difference between the two cases depends simply on whether or not the interlocutors have further information about the individual in question; that is, whether there is some independent means of identifying the referent. If the speaker knows, for example, that Smith murdered Jones, then he can substitute "Smith" for "whatever x is such that x murdered Jones." On this view there is no sharp line that distinguishes the interpretations of referentially and attributively used descriptions; rather, we would say that the more identifying propositions a hearer knows

about an intended referent, the more "referential" the use.

Hintikka's pragmatic analysis, it seems to me, is sufficiently persuasive on its own merits, but the clincher is that the same ambiguity of interpretation is available in ostension. For example, I can point at a car and say "He has blocked my driveway" to mean (roughly) "whoever put that car there has blocked my driveway." But I might also know of the person--and expect my hearer to know--who put that car there that he is my neighbor, that he is named Sam, and so forth. In which case the interpretation "Sam is blocking my driveway" is available by substitution.

(Note, by the way, that this observation makes it very difficult to accept Kaplan's suggestion (ms.) that referential uses of definite descriptions should be analyzed as containing a concealed demonstrative, which he calls "Dthat," because this demonstrative would itself allow either interpretation, even when the RF is the identity, as we'll see.)

Since the distinction between the cases turns on the amount of information available to speaker and hearer, we are not surprised to see that it may affect the statement of the conditions of reference. Consider the following situation: suppose we have a demonstratum $a \in A$, and a possible RF f_r such that f_r takes A into B' , the intersection of the range of reference B and the natural range Y of the given function F from which f_r is derived. For

example, let a be an automobile, and let the range of reference be the set of male humans, as determined by the pronoun morphology in an utterance "He is rich." And let F be the given function "x is driven by y," which is defined over a natural range of drivers, so that B' is the intersection of the sets of male humans and drivers. Now suppose as well that the hearer does not know who drives a; that is, that there is nothing such that he knows that it is the driver of this car. Then, in particular, he cannot know of the driver whether or not (it) is male, i.e., he doesn't know whether $f_r(a)$ falls in the intersection of the sets of males and drivers. And suppose that the speaker is aware of the hearer's belief-state with respect to the driver of the car, and that the hearer is aware of the speaker's awareness, and so on. Then both know the hearer has no grounds for ruling out the possibility that the driver of a is the intended referent. So the speaker could not then rationally intend that the hearer should identify the referent as the value of, say, "brother of the driver of a," for there would then be two possible referents--the driver and the brother of the driver--such that the hearer would be unable to choose between them. That is, there would be a function other than the identity mapping from a (possible) member of B' to another possible member of B' .

For this case, we have to state II in its strong form; so long as the range and domain of the RF (or the

range and domain of the last-computed element of a composite RF) intersect, the RF (or the last-computed element of a composite RF) must be the identity. But now consider another version of the same example: what if the speaker and hearer are both familiar with the car, and know that it belongs to the sister of a wealthy neighbor. Then, while they know that the range of the function to drivers intersects the set of male humans, they know as well that the value of the function at the ostended argument hasn't fallen within this intersection--the driver of this car isn't a man. Under these circumstances, the speaker can use a function like "brother of the driver of a" in referring, since the hearer will not take the value of "driver of a" to be a possible referent. Here, the relevant consideration is not whether the range of reference intersects the domain of the possible RF or its last-computed element, but whether the value of this function at a falls in this intersection. Accordingly, we will rewrite II in a weaker form.⁵

- II. (Weak form) Given a demonstratum $a \in A$, and an intended referent $b \in B'$, and given a possible RF f_r such that $f_r: A \rightarrow B'$, f_r can be used in referring only if
- a. If f_r is a prime function, and $A \cap B \neq \emptyset$, and $f_r(a) \in A \cap B'$, then $f = I$.
 - b. If f_r is a composite function $g \circ h$ such that $h: A \rightarrow C$ and $g: C \rightarrow B'$, and $C \cap B' \neq \emptyset$, and $h(a) \in C \cap B'$, then $g = I$.

The range of situations in which II' applies is not quite the same as those in which "referential" reference is

intended. It is possible to know a good many identifying propositions about the value of $f(a)$ without knowing whether it falls within the range of reference. For example, in the example we have been using, I might know of the driver of a certain car such that he is the person who bashed in my fender last week, and that he frequently parks on this block. Then one could substitute "the person who bashed in my fender is rich"--i.e., this is a "referential" use. But I still might not know whether the driver of the car is male or female, so that I couldn't use the function "brother of the driver of" in referring. The relevant consideration is that I have to know of the value of $f_r(a)$ whether or not it is a member of the range of reference; if I don't know this, then no non-trivial composite $g \circ f_r$ can be used.

A word here on the special case of the identity. It might seem that, when the physical demonstratum is possibly the intended referent, it makes no sense to talk about "attributive" reference--we have the referent available, and can always tell whether it falls within the range. But in this case, as well, we may know more or less about the value of $f_r(a)$. Consider the case of the game-show host who reaches into the box a contestant has selected, and says, "Well this will help you get through the winter." Then the contestant knows of the demonstratum only that it is a physical object smaller than the box; he has no way of saying

whether it is itself or is not a member of the range of things that help people get through winters. It could be a muffler, but it could also be a pair of plane tickets, in which case the RF would not be the identity, but a function from the tickets to a vacation in Bermuda. In this case, the hearer knows only that $A \cap B' \neq \emptyset$; not that $a \notin B'$; and we would apply II in its strong form.

We could state II otherwise, in terms of the conditions that the range of reference has to satisfy in order that a certain RF can be used. The resolution of the game-show case may depend upon how certain the contestant can be that the object in the box is not itself the referent. If this is a show that doesn't give out booby prizes like mufflers, then it is less likely that the physical demonstratum is itself the thing that will help to get through the winter. The determination that $f(a)$ is or isn't in a given range can be made only on the basis of knowledge about a , f , and B . The determination that II is satisfied, then, depends on the identification of the range of reference--the extent to which the predication, the inflectional morphology, and the context enable the hearer to say whether or not the domain and range of the function (or of some component of a composite function) intersect.

Consider a (somewhat more plausible) example. Suppose I point at an army private's neatly pressed jacket, hanging on a hook, and say "He is meticulous." Then if we

know only what we can assume from the rest of the sentence --that the referent is a man who is charged with certain chores, or whatever--"he" can refer only to the jacket's owner. But if we have just been talking, say, about the degree to which various officers take care that their men are well turned-out; if say, you have been arguing that none of the officers are meticulous, and I offer my statement in counter-evidence, then "he" may refer to "the C.O. of the owner of X," since the range has already been constrained to the set of officers on the base, and the owner of the jacket is not a possible referent. Similarly, the range could be constrained by the choice of a different predicate; if I had said "He must be cracking down on his men," the C.O. would again be a candidate for reference, because the private could not be.

It might appear that uses of composite functions in referring are sufficiently rare, and the examples so marginal, that we might better throw them all out as involving computations that are too complex to allow successful reference. But in fact, they abound in ordinary use. Consider, for example, how the referent is determined in a case like 3 (following page). The working-out schemata here are bogglingly complicated; the referent of "this," as best I can figure, is "the events that led to the state (of depravity) exemplified by the aspect of the girl pictured by the demonstratum." (The further assumption that all



*This might
have been
prevented*

**BUT SHE STAYED
IN THE RUMBLE
SEAT TOO LONG!**

- Is spooning dangerous?
- Does a petting party
stop with a kiss or
does it go further?

At last the question is answered

SEND \$1.00 TODAY FOR THIS 200-PAGE ILLUSTRATED BOOK
(1931 EDITION JUST OFF PRESS)

**Sex and Love Relations
FOR THE YOUNG ADULT**

BOSTON, MASSACHUSETTS

Copyright Key Publishing Co.

such dissolution might be prevented is, I think, a Gricean inference from the "literal meaning" of the statement. Obviously, it is not of general interest that just this one woman's fall might have been prevented.) Getting to the referent here involves a progressive restriction of the range. It could not be the writer's intention to say that this picture could have been prevented; nor simply that her having been photographed in this pose might have been prevented--if the photo had been taken by a peeping Tom, the connection with petting would not be relevant, though that could be the interpretation if this were an ad for window shades. So we assume that her posing must have been deliberate. But the fact of her having once so posed could not be the referent--while that could be prevented, say by fortuitous police action, the causal connection with petting would again be lost. So we need another function to take us from this act to a disposition to perform it, and then, since it is the formation of dispositions, and not dispositions themselves, that can be "prevented," to the circumstances which led to the adoption of the disposition to behave habitually as the girl pictured by the demonstratum. But the details of reference at this point become rather inscrutable, and I leave it as an exercise to the reader to figure out what the poster is about.

To a large extent, then, the determination of the range of reference involves some complex computation over

contextual factors, such as the topic of conversation, the speaker's opinions, and so forth. In many cases, it will be hard to say whether or not II is satisfied, since there may be disagreement about the extent to which the range has been limited, for any context of utterance. But obviously the hearer can reduce the range of reference only insofar as his apprehension of the speaker's intention allows him to; the range has to be defined as the largest set of things satisfying the predication such that the speaker might reasonably be expected to refer to them in the context.

There is some advantage in stating II as an absolute condition, but it's really just a special instance of the strategies that enable the hearer to pick out one of a set of possible RF's to a contextually determined range of reference--the second part of the job that the working-out schemata are responsible for. The identity is simply the "best" available function, and it will be on the basis of the "goodness" of the functions to a given range that the selection of the RF is made. This notion of a "good" function will be crucial to our discussion of the lexicon, where we will want to be able to distinguish conventional uses of words from non-conventional ones. But to talk about the lexicon, we have to turn the problem around a little. Rather than asking "given a demonstratum a, and a range of reference B, which function to B is most likely being used?",

we will assume the intended referent, and ask, "Given an intended referent b, and a context of utterance U in which the set of objects A is prominent, which of the members of A is such that ostensive identification of b could most easily be made?" When we are talking about demonstratives, the answer to this question isn't too interesting, because A varies widely from context to context. For a randomly chosen referent b, and a given context, it's unlikely that b could be ostensively identified at all. How can you refer to Socrates in an average post office, using only a demonstrative term? But with common names, the universe of elements that can be the arguments of the RF is not limited in this way, and they are always available for use. A speaker always has newspaper at hand. In which case the question of "convention" arises. Given that there could be two members of A--the set of designata of natural language terms--such that each was the argument of a different function to the same member of B, then why should one or the other be consistently used? Doesn't this then mean that this particular use is an arbitrary convention, whether or not we call this convention part of the "meaning" of the term?

In some cases we will have to say that it does. But the fact that a use is regular doesn't mean that it is arbitrary; the use may be entirely rational. The use of newspaper to refer to newspaper companies is of this sort,

I believe. And even if there is some other term that might be so used, and if there would be some rationality to its use, the argument that these uses are arbitrary requires that we show that use of the other term would be equally well-motivated. It is certainly rational that we should call certain parts of chairs "arms," but it wouldn't be irrational to call them "wings" instead. Does this mean that the use of arms is arbitrary? No--I think we can argue that the function from arms to those things is "better" than the function from wings to those things, given the conventions of reference, and our "theories" of human anatomy, birds, and furniture. So we'll want to know how people choose among possible functions, because the definition of l'arbitraire du signe will ultimately turn on the answer.

Moreover, we'll require some notion of a "best" RF in order to be able to specify the "meaning" of common names at all. This is a second problem that arises with common names and not in ostension. When a demonstrative term is used, we always know what the argument of the RF is --we can pick it out in context. Whereas descriptive terms are always at a remove from reality, and there may be a problem in distinguishing the meaning--what is designated by the term--from among the members of the set of things that the term is often used to refer to.

For the moment, I want to continue the discussion using only examples of ostension, but with an expanded

hypothetical context, one that is more like the universe made available by the lexicon. Suppose we are in Macy's, and that we can transport ourselves instantly from counter to counter (like Harpo on his roller skates in The Big Store). Now suppose I want to refer to some b by pointing at something in the store; what thing might I most reasonably point at? Of course, this depends on my knowledge of the referent. If I know nothing about it save that it is the value of a function from a given argument--"whoever designed this glass," or "whatever the price of that chair is," then I have no choice among possible demonstrata, and the question of a "best" function doesn't arise. So for the moment, we'll keep to "referential" ostension: uses of demonstratives to refer to things with which both the speaker and the hearer are well-acquainted. And, of course, it will be more interesting to talk about referents that are not themselves in the context of utterance; i.e. not floorwalkers or bassinets, since condition II would require that we point at these very things to identify them.

Let's take as an intended referent something that might be identified by pointing at any of a number of items in the store; for example, "the game of golf." There will be a number of things in the sporting goods section that I could point at, each of which is a member of a discriminable inverse image set of the given function "x is used in y" for $y = \text{golf}$: tees, clubs, balls, hats, shirts, etc. And

similarly there will be golf books in the book department, golf-motif tumblers in kitchenware, as well as--grant for argument's sake--pictures of Lee Trevino in men's slacks, and of President Eisenhower in the executive boardroom. All of these are the arguments of other functions whose value is golf: "x is about y," "x is decorated with pictures of things common to y," "x plays y," and so forth. So the choice of a demonstratum is made on consideration of two things. First, given several different functions to sets of things of which golf is a member, which do we choose to derive the RF? And second, given any one of these functions such that there is more than one member in the inverse image of the referent (golf) for f--to which we point; i.e., given that "x is used in golf" might be true of any of several things in the sporting goods department, which counter do we skate to?

We can start with the second question; the answer to the first will follow trivially. Given the function "x is used in y," which of the things in the sporting goods department gives us the "best" function to golf? Which is simply to say, "by pointing at which of these objects will I be most likely to enable my hearer to identify the intended referent?"

In assuming that both speaker and hearer are "well-acquainted" with the referent, we assume that they are familiar with a fairly extensive set of propositions about

it. We might make reference here to the notions of stereotypes or prototypes, as used by cognitive psychologists, or to the notion of the "broad concept" of a category, but the term "theory" seems to be most generous. Let me take the theory of the referent to be simply a (finite) set of propositions "about" the referent that we can attribute to a particular individual; let us assume also that speaker and hearer have identical theories of the referent, and that they are aware of one another's theories, and of this awareness, and so forth.⁶ Then a common theory of golf might include all of 1-6:

1. Golf is a game.
2. Golf is played with clubs of shape s, composition c, and so forth.
3. Golf is played with balls of shape b, etc.
4. Players of golf often wear shirts of design d.
5. Golf was a hobby of President Eisenhower.
6. Golf is played professionally by Lee Trevino.

and so on. (Note that not all of these propositions would necessarily be listed under what we would think of as the "stereotype" of golf: 5 and 6, for example, might more appropriately be entered in cognition under the headings "Eisenhower" and "Trevino." But this is only a reason for preferring the broader term "theory"; we can't worry here about the organization of cognition.)

Allow me now to talk about 1-6 as properties that

golf is believed to have. Some of them are believed to hold of all instances of the game, such as 1-3; others are believed to hold only typically, such as 4. (We will not worry here about which properties are held to be criterial of golf, and which only contingent; or over which propositions are true of golf "analytically," and which only "synthetically.") Second, some properties characterize only golf, such as 2 and 3, while others, like 1 and 5, characterize other sports as well. Let me borrow here the psychologists' notion of cue-validity, and say that a property p has a high cue-validity for a given referent b to the extent that the proposition "x has p" is likely to be believed to be true for x = b, and false for everything else. All other things being equal, it follows that the greater the cue-validity of the property of being the value of f(a) for b, the higher the probability of successful use of a in ostensive identification of b. So inasmuch as the property of being the value of the function "x is used in y" has a higher cue-validity when x is a golf club than when x is a golf shirt, the function from clubs is the better of the two. (From here on, we can simply talk about the cue-validity of functions, avoiding the awkward circumlocutions that talk of properties forces on us.)

At this point, the two problems involved in choosing a demonstratum--picking a function, and picking from among the arguments that allow derivation of an RF for a

particular function--can be collapsed. For each of the possible demonstrata $\underline{x}$ in the store--clubs, pictures, books and golf shirts-- $f(\underline{x})$ for the relevant function to the intended referent golf will have a different cue-validity, and we simply choose the highest among all these. Necessarily, the identity has the highest cue-validity of any function; given $\underline{a}$, the probability that we can pick out $I(\underline{a})$ is 100%. So the requirement that we use the identity whenever it is a possible RF, which is condition II, follows from the way in which the RF is chosen.

For functions other than the identity (and other than some other cases, like the "hypostatic" function from tokens to types), cue-validity is determined by two theories: one for the things in the range, and the other for things in the domain. It is in virtue of our theory of golf that we know that it is always played with clubs of such-and-such a shape. And it is in virtue of our theory of golf clubs that we know that they are intended to be used for no other purpose than to play such-and-such a game. (Identifying the demonstratum as a "golf club" may somewhat prejudice the issue here. So let's call it a niblick, and say that it is part of our theory of niblicks that they are used only in golf.) To the extent that the relevant properties are believed to hold categorically of golf and niblicks, the function from one to the other is "good."

Provided, of course, that the hearer has the same

theory of these things that we do. When I say that these things are part of "our theory" of golf and niblicks, I am talking about the members of "our" speech community--and in particular, the participants in the talk-exchange. Obviously, theories vary from individual to individual, and the cue-validity of a given function $f(\underline{a})$ for a given referent $\underline{b}$ varies accordingly. To the extent that theories of niblicks and golf are invariant throughout a speech-community, niblicks can always be used in ostensive identification of golf. And the invariance of these theories, with respect to the relevant propositions ("Niblicks are used in golf"/"Golfers use niblicks"), must itself depend on a number of other factors, such as what kinds of things golf and niblicks are ("functional" properties will be more usually invariant for artifacts than animals, for example), and on the degree to which the argument and value are familiar to all members of the speech community. All of this interacts in a complicated way. For example, the function from persons to their professions is generally "better" than the function from persons to their avocations, which are therefore less useful in identification. It is unlikely that there is anyone who could identify the photo of Lee Trevino who would not know that he is a golfer (assuming that he is not shown carrying a club, in which case we would simply be pointing at a picture of someone playing golf, a wholly different matter). But it is quite possible that someone

should be able to recognize Eisenhower without knowing that he was a golfer. On the other hand, the likelihood that Trevino will be recognized at all is much smaller than the probability that Eisenhower will be recognized, so that you have to give Eisenhower a couple of strokes by way of handicap. It would be hard to say whether the cue validity for golf of "x plays y" is higher for one or the other picture.

In this way, the choice of the "best" RF depends on the speaker's ability to estimate what beliefs his hearer has about the referent and demonstratum. And among the things he has to take into consideration is the hearer's estimation as to what the speaker's estimation of the hearer's beliefs will be--the cooperative regressus. The probability of successful reference is increased as the participants in the talk-exchange are able to apprehend one another's beliefs, and act accordingly. And in the same way, the consistency with which members of an entire speech community use certain terms to refer to given kinds and individuals--that is to say, the extent to which "conventions of reference" are rational--depends on the extent to which the culture of the community determines a uniform world-picture, to which all of its members have access.

It is clear, accordingly, that heterogeneities in speaker-beliefs must play an important role in determining the way in which terms are used to refer, and ultimately, in the form of the representation of the speaker's knowledge of

the conventions that govern the uses of words. But it will be easier to discuss the application of our theory of ostensive reference to descriptions in 2.2 against the assumption of absolute homogeneity of belief-systems within a community, and to reserve our discussion of heterogeneities for Chapter Three.

2.2 Descriptive Reference

2.2.1. The Determination of Meaning

Now we can turn to the problem of generalizing the account of ostension presented in the last section so that it will cover the use of proper and common names, as well as demonstrative terms. As a preliminary, let me clarify my use of terms. What I'll be calling "common names"--words like tiger and newspaper--are often called "descriptive terms," and it has sometimes been argued that proper names as well should be analyzed as "concealed descriptions." And it's common, even among linguists, to use "description" to refer to phrases like "the king of France." But I want to use "description" here only to refer to acts, keeping it parallel to "ostension"; the two will exhaustively subcategorize the ways in which natural language terms can be used in referring (though not necessarily all possible kinds of reference). I'll use "name" to refer to terms themselves, and "descriptive phrase" to refer to phrases like the king of France, independent of their use. Finally, I'll be using "referent" as equivalent to "intended referent," without the ontological

commitment that the term has sometimes assumed in the philosophical literature; this usage follows the standard linguistic practice.

The difference between ostension and description, of course, is that in descriptions, the argument of the referring function is not before us. To generalize our account of ostensive reference, then, we have to be able to say what sort of thing it is that we must take as the descriptive analogue of an ostended particular; i.e., what stands in relation to a use of newspaper as a newspaper copy may stand in relation to a given use of a demonstrative term like that. Depending as we answer this question, we will have to revise the conditions of reference given in 2.1, replacing terms like "manifest" with others more apposite to the use of names.

For the present, we can do with saying simply that names like Dickens and newspaper designate things and kinds of things, leaving for a later chapter such questions as what sort of relation holds between the name and its designatum, or how the designata are defined; these are the responsibility of the theory of meaning, not of reference. We will further assume that designata are described by theories that are uniform for the entire speech-community, so that the cue-validity of the RF will be the same for all speakers on a given occasion of use of a name, provided that all have access to the same contextual information. We can then replace such conditions as "It must be manifest that $a \in A$ " with more

general statements like "It must be generally believed that $a \in A$," and so on.

Before proceeding, two notes. Terms like "designatum" and "nominatum" are often used interchangeably with "referent" or "extension"; let me stress that on our view the two are quite distinct. We could say that newspaper designates a kind of publication even on those occasions when it is used to refer to a company; the RF would then take us from a kind or class of publication to one of its members, and from there to an organization. (Note that the argument of the RF in description, unlike in ostension, is not a particular, unless a proper name is used.) It should also be noted that we will have no call to distinguish among properties that are criterially ascribed to categories, and those that are ascribed only characteristically; we do not have to believe that newspapers are criterially published as they are in order to use the term that designates them to refer to companies. But as with ostension, descriptive reference will be most likely to succeed where the referent is criterially the value of the RF, since the cue-validity of the RF is then 1.0.

The fact that the argument of the RF is not present in descriptions leads us immediately to another point. If we use a name once to refer to one thing, and another time to refer to something else, how do we determine which of those instances the thing or category designated by the name? On what grounds, for example, do we say that newspaper designates

a kind of publication, rather than a kind of company, and so specify one category rather than the other in its lexical entry? Moreover, how can we be sure that we can always select one of the referring uses of a name as its lexically-specified meaning?

The problem of selecting the meaning of a term from among its referring uses has two aspects, which correspond to the two conditions of reference outlined in the previous section. In both cases, we have a situation as follows: a given term t is commonly used to refer both to a and b, and, let us say, to nothing else. Then there are three analyses available to the speaker:

1. t could be homonymous--that is, there could be two forms, t₁ and t₂, which designated a and b respectively. For example, if there were two people in a room named "Tom Jones," we would assume homonymy.
2. t might designate a, while being used to refer to b by means of some RF not the identity. (When Tom Jones is used to refer to a kind of record, as in "I just bought a Tom Jones," we would presumably assume this analysis.) Conversely t could designate b, etc.
3. t might designate the set {a, b}--that is, the RF could be always the identity. (If a and b are the sets of male and female dogs, for example, we would say that the multiple uses of dog were of this sort.)

The problem of determining whether 1 is the case--whether a term is homonymous--will be left to a later section. For the moment, we're interested in cases where a univocal analysis is possible. The problem raised by 2, of deciding which of a disjunctive set of referents is the designatum, is

resolved by making certain assumptions about the way in which speakers are satisfying condition I--the inverse image condition. The problem raised by 3, which is the problem of saying whether t is simply "neutral" in meaning with respect to the difference between a and b, involves assumptions about the satisfaction of condition II--the identity condition. But the connection between the problems and the solutions is easier made in terms of some simple examples.

Let me begin with some examples of "condition II" situations, where we are trying to figure out whether all of the things a term is used to refer to are members of the same set. This is a problem that has been discussed in both the philosophical and linguistic literatures; for example, the argument over good is of this sort. But it won't be convenient to use the traditional distinction between "univocity" and "multivocity" to describe these cases, because this makes no provision for the situation described above in 2, where we may preserve a univocal analysis without having to grant a single extension. Thus, it is conceivable that good should designate only a moral property, while being used to refer to other properties. I won't defend this analysis--I think it's wrong--but it is a possible alternative.

The question of whether or not two referents can be counted as part of a single extension depends in large part on what kinds of ontologies we accord speakers (the study of which Strawson calls "descriptive metaphysics"); it will be on

these grounds, finally, that we decide whether or not a set $\{a, b\}$ constitutes a "natural class." But this kind of question is better left to others to worry over; I'll confine myself here to showing how some simple cases can be incorporated into the framework I'm using, and pointing out a couple of purely linguistic difficulties that may come up along the way.

Take a trivial case. The word bear may be used to refer either to males or females of the species, and no one would want to say that these were different uses; i.e., that the RF was different in either case. Still, a speaker who has heard the word used only once, say of a male, would have no way of telling whether bear was more like horse or stallion. On hearing bear used for a female, however, the hearer will know that it is the species name. For suppose that bear designated only male bears, and that its use to refer to females involved a function from male bears to male-bear-like things. This function could not be as good as the function from bears to bears--the identity--which would be the RF if bear simply designated the species. All things being equal, then, the hearer will assume that the speaker is using the identity as the best function to the referent; he will try to define the domain of the RF as the largest single "epistemically natural" category that includes all of the things the term is used to refer to. Whether he can do this depends on whether he can come up with a unified theory of the argument

of the RF; that is, whether there is a set of propositions P such that the cue-validity of P for $\{a, b\}$ is as high as the cue validity of any other set of properties P_1 for a or b alone. And this is clearly possible for bears, which do constitute a natural class.

Consider a slightly more elaborate example: The word cub is used to refer to the young of foxes, wolves, tigers, lions, bears, etc. Again, if we heard the word used only for young foxes, we would not know whether to analyze it as designating simply the young of that species, on the model of kitten or puppy, or as designating some larger class. (Let's assume we have a way of telling that the word doesn't designate foxes, or canines--that we know its designation is young somethings.) But subsequent uses to refer to young lions, tigers, and wolves would allow the construction of a single class of "young wild carnivores," so that all these uses could be treated as involving the identity as the RF.

Suppose, however, that the picture were more complicated; supposed there were a word simblet which was used to refer only to young lions. Then there would be no set of properties that has a high cue-validity for just the remaining species; whatever theory generates tigers, foxes, wolves, and bears must generate lions as well. But I don't think the existence of this word would force us to re-analyze cub as homonymous; we would still take it to designate young wild carnivores, the superordinate category of "simblets." On the

other hand, the introduction of a sufficient number of terms like simblet might force such a re-analysis; if cub were used only for young foxes and tigers, we might be less disposed to grant it superordinate status. It would then be more like cow, used for female cattle and whales, where there is little temptation to assume that cow is the "unmarked" term for female mammals. It's hard to say at what point to re-analysis is forced; we would require some notion of "economy in a cognitive system" which is well beyond our present reach. But I can suggest a rough operational test for such re-analysis. So long as cub is analyzed as designating young wild carnivores, we can safely assume that it will be the productive form--that we will use it for young hyenas and ocelots, for example--because the RF will still be the identity in these cases, and that is the best available function. But once cub is re-analyzed, say, as the homonymous realization of several words that designate young foxes, tigers and wolves respectively, then its use to refer to young hyenas and ocelots--while still possible--will involve another function, from young tigers, say, to young tiger-like things. In which case it will be in competition with words like wolfing and simblet, which designate the young of other wild carnivores, and the choice of cub for other species will be nowhere near so uniform--the function from wolves to hyenas could be better than the function from foxes, for example.

As cases get more complicated in this way, however, it becomes increasingly unlikely that all speakers will come up with the same analysis of the terms, or that we will be able to say for sure whether all of the referents of a term are parts of the kind it designates. I'll talk about the effects of such heterogeneities in the following section; let me now give a typical example (borrowed from Bolinger 1975). The word cell is used to refer to the constitutive parts of a number of things: honeycombs, living organisms, batteries, political organizations, matrices, cigarette filters, and certain buildings, such as monasteries and prisons, which are divided up into small, largely uniform compartments. And I suppose we could say that all of these form a single extensional set; that they are all instances of the uniform sealed compartments of some larger structure. It's true that there are things which might then be called cells, but aren't, such as the alveolae of the lungs, the pigeonholes of a desk, or the cabins of a steamship, but in each case we can explain the gap. Alveolae are like our hypothetical "simblets"; they happen to be designated by another term (with good reason; otherwise talk of the "cells of the lung" would create intolerable ambiguity, since there would almost always be two plausible candidates for reference within a contextually-determined range.) Pigeonholes, as it happens, have been called cells, according to the OED, but here again, another term is available. In fact, I can think of nothing such that it could be

called a "cell" that is not so called, unless there is another term which designates that kind of thing in particular.

I don't think it is correct to say that cell designates a single superordinate class, but this kind of analysis has often been attractive to linguists. (CF, Trench, cited above, who proposed a single meaning "that which has been placed" to underlie all of the uses of English post.) The problem with cell is that even though it is true that all of the things that we call "cells" do happen, usually, to be the uniform sealed compartments of a larger structure, they are not all criterially so. We use cell to refer to unicellular organisms, for example, which are not part of any structure. Yet the "part of a structure" requirement is clearly needed; the word cannot now be used (though it once was) to refer to a single one-room house. Conversely, some of the categories referred to by cell require additional definiens. The cells of a political organization are criterially clandestine; the cells of prisons and monasteries are criterially rooms to or from which access is restricted (which is why, I think, it would be odd to use cell to refer to library carrels, except in jest.) So there is no proposition or set of propositions such that it would individuate the class of things called "cells" from other classes, and hence no reason to suppose that cell is always used to refer to its designatum.

At the same time, many of the uses of cell do seem to be connected by referring functions. If we knew what a matrix was, for example, and were familiar with all of the other uses of cell, then we would have no trouble in figuring out what parts of the matrix are being referred to when we talk about its "cells." And a cigarette manufacturer can confidently advertise a filter as being divided up into "cells," without worrying whether listeners have heard this particular use of the word before. So cell cannot be homonymous, as in case 1 above; rather, we must be dealing with case 2, whereby a word that designates one kind of thing is being used to refer to another. But then, which of the things that cell is used to refer to does it designate?

Cell is not a good example to begin with, because various kinds of conventions intrude--a problem we'll discuss in the next chapter. Rather, let's consider a much simpler kind of case. Names like Dickens and Homer can be used to refer both to certain authors and their works, which gives us the ambiguity of "Dickens is interesting." There's no question here of the term merely being non-specific--authors and books simply don't form a natural class. And in this case, we can rule out homonymy because the process is so productive, though in other cases the decision won't be so easy. That leaves us with the problem of saying whether Dickens designates a man or an oeuvre. Of course, we'll say that it names the writer, but on what grounds?

Take an even more stripped-down version of this case: suppose we know that there was a man living in Anglo-Saxon England, called "Caedmon," who wrote a certain set of poems, which are also called "Caedmon"; suppose, also, that we knew no further proposition about either the man or the works. We would nonetheless feel secure in saying that Caedmon designated the singer, not the song, in virtue of our theories of persons and poetry.

We can subcategorize literary works according to similarities of form (sonnets), function (odes) or provenance (Augustan); "oeuvres," let us say, are bodies of works among which a consistency of individual style or sensibility can be perceived. An oeuvre does not have to be the work of one person; we may talk about the collective oeuvre of the Brontës, or identify as a single oeuvre a group of works--such as the "Juliana group"--whose authorship is unknown. In general, however, we explain the idiosyncratic properties that differentiate one oeuvre from another by reference to the circumstances or intentions of its author (or the common circumstances and intentions of a group of authors). So we would assume that whatever properties distinguish one oeuvre from others, it has as a result of the properties of its author. Even if we don't know what those properties are, we assume that it could not be as it is if he were not as he is. Given the author, then, we have a good function to the oeuvre; its individuating properties will all be expressible

as the values of functions from his. By contrast, Caedmon will be best differentiated from other men by his lineage or circumstance of birth, not by having written such-and-such a body of poetry. It is true that there is a function from its properties to certain of his, but they do not define him. Caedmon (that very man) could have been other than pious and trite, but the Caedmon (that very poetry) could not have been witty and irreverant without constituting a "different" oeuvre.

We are more likely to succeed at identifying an oeuvre as the work of a given man, then, than to identify man as the author of a given oeuvre. If we assume that speakers are behaving rationally, then we must assume they are following the former course in using the word Caedmon. The general procedure for determining the meaning of a term in these cases works out as follows: given a term t that is used to refer to both a and b, and a pair of "good" referring functions, f and g, such that $f(a) = b$ and $g(b) = a$, we assume that t designates that individual that is the argument of the function that has the highest cue-validity for the other (classes of) referents.⁷

But this determination can be made only while keeping in mind the beliefs of the community in which the term is used. I've been assuming that we are talking about the usage of present-day English speakers, among whom it is widely held that the properties of a literary work are explained by

reference to the character and intentions of its author. If we were trying to figure out what Caedmon designated a thousand years ago, we might come to a very different conclusion. As Bede tells the story, the poems were recited to the illiterate shepherd Caedmon by an angel who came to him in a dream. In which case the function from the properties of Caedmon (vassal vessel) to the work would be nowhere near so good, and we might decide that Caedmon designated the works, or even the angel.

Let me give two more variants on this story. Suppose instead that we have a name Hey Jude that is used alternatively to refer to a song, and to a certain type of phonograph record on which this song appears. Here, we would go the other way, and say in all likelihood that the term designated the (abstract) work, rather than the record. Because we will explain the individuating properties of the record in terms of the properties of the song, but we will not explain the individuating properties of the song in terms of the properties of the record, but in terms of its tune and content. Again, there will be a difference in cue-validity between a function and its inverse, and we will choose the better of the two as the RF, and its argument as the designatum of the term.

Finally, an indeterminate case. Suppose that we have a body of verse called "Fungoids" which is known to have been produced by a computer program also called "Fungoids." Which

way we go here, I think, will depend upon what further facts we know about the software. If it is simply a black box that does no transformation on the input, then it would be the precise analogue of the phonograph record, and we would say that Fungoids designated the poems. But suppose the program generated the poetry from a random input according to certain heuristics. Then, depending on the power and complexity of the heuristics, we might well conclude that Fungoids designated the program, and was being used to refer to the output; this, would be more like the Caedmon case.

The determination of meaning in cases of this sort, then, involves the assumption that speakers are behaving in such a way as to maximize the probability of successful reference, and are therefore choosing the best of the possible array of functions that could get them from kind to kind. And cue-validity is calculated, as we saw in section 2.1 on the basis of the degree to which the inverse image condition is satisfied. Needless to say, then, that the determination of meaning is just the evaluation of the calculations that enter into the determination of reference, and that it is pointless even to try to distinguish the processes.

2.2.2. The Extent of Polysemy

This section will be largely an aside, in which we try to give some idea of the range and pervasiveness of the different functions that are ordinarily used in natural language description. Against such a background, we can more

conveniently mount the argument in the following section; moreover, the description may have some practical linguistic interest for its own sake. Up to now, it has been sufficient to argue from a few anecdotal cases of multiple word-use. In most of the literature on polysemy, in fact, it seems to have been assumed that most words have a fairly small number of psenses, on the order of the number of sub-headings given under a single entry in the OED. But this is in some ways a misleading assumption (it may have contributed in part to the conviction that polysemy could be dealt with on the model of homonymy).

To describe the range of possible word-uses other than anecdotally, we would need no less than a general account of cognitive structure, so that we could enumerate all of the ways in which categories can be perceived as related. In what follows, I'll try to do no more than to sketch out a rough taxonomy of the functions that can be used in referring, and to talk about the uses of a few words in a little detail. I'll further confine myself to talking only about nouns, and avoid cases where a use or set of uses has become in some way "conventionalized." We won't take up the question of kinds of convention until the next chapter; for now, let us say that a use is "free" when it could be generated by most speakers solely on the basis of the encyclopedic information available to them. We could use a rough heuristic for "free" uses the fact that they are mirrored in ostension, or that

the pattern of uses is often found cross-linguistically, or that the two uses can be related syntactically by anaphoric devices. None of these tests is infallible, or always applicable. They will suffice, however, to disallow the use of tongue to refer to a part of a wagon, or the use of book in "make book on the races," where an element of convention seems obviously to play a role.⁸

For the time being, we can do with a rough taxonomy of the functions that can be used in referring. Let us say that kinds of word-uses can be "normal," "local," or "metaphorical," according to whether the use would be judged rational out of context, rational in context, or irrational but interpretable. (Section 4.1 will be given to an extensive discussion of such judgments.) I will try to discuss only normal uses here.

We can further distinguish among normal uses that are related by "general" and "culture-specific" functions, though here again the distinction is only rough. 1-3, for example, instance a general function:

1. John bought the calendar watch.
2. Bulova is introducing the (new) calendar watch in July.
3. John invented the calendar watch.

In 1, calendar watch refers to a token, in 2 to a sub-class, in 3 to a type. This "hypostasizing" function may include virtually any argument in its domain, and is

universally used. (There are a few cases in which its application is pragmatically odd; see Vendler 1967 for discussion.)⁹

Another such function is illustrated in 4 and 5:

4. The committee is our only obstacle.

5. I was surprised by his obstinacy.

The underlined NP's here can refer to "the fact that the committee exists," or "the fact of his obstinacy." His obstinacy can also refer, of course, to "the extent to which he is obstinate"; any noun that can be used to refer to a kind of thing can be used to refer to the dimension or extent of some instance of it, as in:

6. He can hit the ball two football fields.

7. The turnout astonished the mayor.

8. John is six foot two.

(8 is of particular interest, because it points the way to a univocal analysis of the copula as the identity; see the conclusion for further remarks on this.)

Two other general functions merit special attention.

Any noun can be used to refer to a description or representation of an instance of its designatum, as in:

9. No one can read the death of little Nell without laughing.

10. That's my last duchess hanging on the wall.

11. There's a wooden Indian on the porch.

12. Her Ophelia was affecting in last night's performance.

The "use/mention" indeterminacy is also treatable as a general function; any term can be used to refer to a name of its designatum (which is not necessarily the same thing as being used to refer to itself, as 14 makes clear):

13. "Plato" has five letters.

14. "Plato" has six letters in French.

The referent here can be of several types. In 13 and 14 "Plato" is used to refer to orthographic objects, in 15 it is used to refer to a phonetic object.

15. "Plato" begins with a stop.

Similarly, a word can be used to refer to phonological objects, morphological objects, and so forth.

There have been a number of attempts to deal with both of these cases in the semantics. Notice, however, that they behave just like other pragmatic indeterminacies with respect to ostension, as we noted at the end of section 2.4. One can point at a man and say 16:

16. He has blue eyes in John's portrait.

Or one can hold up a fountain pen and say:

17. What letter does this begin with?

Similarly, anaphonic devices can ignore these distinctions, as illustrated by 18 and 19:

18. Plato, who was Greek, is wearing a top hat in Marcel's portrait.

19. "Glass," which is a solid, begins with a stop;
"water," which is a liquid, begins with a glide.

(19 is impossible to punctuate, but easy to say--another relevant example; I will leave it as a problem for Derridians to worry over.)

Other general functions are limited in application. Words for linguistic objects, for example, can always be used to refer either to forms (orthographic, phonetic, etc.) or meanings:

20. The word may be obscene, but it is easy to spell and to pronounce.
21. John's dissertation, which was written in a hurry, has recently been refuted.

This indeterminacy bears an obvious relation to the use-mention indeterminacy. We could say that 13 and 14 involve composite functions of the form "spelling of name of" etc. Similarly, words for directions (left, right, east, etc.) can always be used to refer to the places that these directions pick out relative to an origin, either fixed ("Carter comes from the South") or contextually determined ("I looked to the South"); and the names of places can always be used to refer either to the places themselves, or to the things, climates, or populations found there:

22. Hawaii is volcanic/ growing in population/
hot in July.¹⁰

In these cases, we know that the designatum of a word falls within the domain of a particular function simply in virtue of the superordinate category to which it is assigned. The function "x is pronounced as y" yields a unique, discriminable value for every linguistic object (I should better say

"linguistic act-type"); the function "x is the locus of y" discriminates among classes of things or persons for all places. More often, however, contingent factors enter into the determination of whether a particular thing or kind of thing falls within the domain of a particular function, and of whether that function satisfies the inverse image condition on that argument. We will then say that a function is culture-specific. Take the function "x is made of/ from y," which takes us from sets of individuals to kinds of substances. Given our belief that most "natural" kinds are defined on the basis of internal constitution, it follows that the stuff that chickens are made of should be manifestly different from the stuff that geese and ducks are made of, so that we can use this function in referring to "masses," as in:

23. I like to eat chicken (duck, turtle, etc.).

In the same way, we can use this function to refer to kinds of wood or wine:

24. The table is oak.

25. He was drinking Pinot Chardonnay.

Where members of a kind do not all instance the same substance, however, or where the same substance characterizes more than one kind, then the cue-validity of this function will be lower, and it cannot be used in referring. So we say:

26. Fruit is good for you.

27. His bracelet is made of hair.

But not,

28. ?Vegetable is good for you.

28. ?His purse is made of pancreas.

(Note that liver is O.K. in 28, since we assume that livers are made out of a characteristic substance.) Of course, we have only intuitive evidence for saying that a use of chicken to refer to a kind of meat involves a function other than the identity. And while that is plausible in these cases, it sometimes seems that the inverse of the function, from kinds of things to substances, is at work. Thus we might want to say that brick designates a kind of substance, and that when we say "a brick," we invoke a function from substances to the things characteristically made of it; likewise tissue, tile, and canvas.

The relation between kinds of things and their characteristic substances may be simply a special case of a more basic "source/product" relation, which is at work in various ways in examples like:

- 29. I bought a KLH (Cadillac, Flexible Flyer, Steinway).
- 30. The lamp is bright (The stereo is too loud, etc.).
- 31. You need a tweeter to reproduce the cymbals accurately.

Each of these cases has its wrinkles. (Why, for instance, can we say "a Cadillac" but not "a G.M.": why can a "G.E." be used to refer to a washer or a radio, but not a light bulb? We'll come back to this question in 4.1.) We may note another interesting property of examples like 29. We might want to say here that KLH named a company, and that a function other

than the identity was involved in its use to refer to a kind of stereo component. There are a number of functions like this, which take us from individuals to other things (that is, there are several ways in which proper names can be used in referring). The classic examples involve words like Casanova--what Peter Quenell called "one of the runaway names that have become dissociated from their original owners and indissolubly attached to a certain kind of human conduct." (Likewise Mata Hari, Judas, or Job.) These have a certain anecdotal interest, but the process has instances that are more interesting and productive. Example 22 (with Hawaii) was one such; consider also 33-38:

- 33. There was a rush on IBM.
- 34. Burgundy goes well with meat.
- 35. New York won, 21-17.
- 36. He owns a Vermeer.
- 37. We sold sixty Buddenbrooks last week.
- 38. He danced a brilliant Swan Lake.

and so on.

On the other hand, it could be that KLH names a kind of component, and that its use as a company name is pragmatically derived. This is a still more interesting kind of function, whatever its applicability in this instance. Consider 39-47, all of which seem to involve a "derived" name or indexical term.

39. I'll see you on Tuesday (at lunch).
40. Go to Hell.
41. God is on our side.
42. Mother is in a bath.
43. Congress convenes next week.
44. He finished Chapter Two in March, 1965.
45. It's ten degrees above freezing.
46. Queen-six bets.
47. Don't forget to touch second.

Each of these cases has its own intricacies; suffice it to note that all of them involve at least moderately productive processes.¹¹

At this point, we are getting into functions that are quite restricted in application. Rather than classifying the functions according to type, we might look at a few representative words, and try to enumerate some of the things they are standardly used to refer to.

Take radio. It can be used to refer to a physical object, of course. It can also be used to refer to a method of transmission, as in:

47. They got the news by radio.

(Likewise television, train, airplane, slow freight. Sometimes the medium is the method: by sea, by air. Or the instrument: by hand, by foot.) Radio can also be used to refer to an industry, as in:

48. He made a pile in radio.

(Likewise television, ice cream, and almost anything else around which some branch of commerce has evolved.) This use is not quite the same, I think, as:

49. He's one of the most creative DJ's in radio.
where reference seems to be a collective enterprise, rather than a business. At any rate, both uses are clearly distinct from uses like:

50. Radio has gone downhill since TV came in.
where reference is to (the quality) of the product commonly transmitted over radio sets. (Radio used to have a use to refer to a certain range of the electromagnetic spectrum, like microwave and CB, but FM changed that. FM is now used to refer to a frequency range, however. TV, which is transmitted over the FM range, has no such use.)

Or take game. Suppose it designates a kind of activity. Then it can be used to refer to the number of points needed to win a game, or the rules of a game, or the outcome of a game, or a community involved in a game, or the way in which someone plays a game, or the equipment with which (certain) games are played. In each of these uses, it patterns with another, different set of words, as in:

- 51. Game is 21 points (cf. match).
- 52. The game is simple: you try to make the other player spell a word. (Cf. language, taxpaying, or any other rule-governed activity.)
- 53. He's one of the best in the game today (cf. radio, above, and many other words.)

54. His game is not what it was. (Cf. French, driving, etc.)

55. You left the game all over the floor. (Cf. dinner.)

Game can also be used to refer to periods of time, in at least two ways. Like other words that can be used to refer to bounded activities, game can be used to refer to the duration of a game:

56. He's been pitching badly the whole game (inning, trip, meeting, etc.).

And like other words that can be used to refer to regular events, game can be used to refer to the interval between games, as in:

57. He hit a home run two games ago. (Cf. beer, meeting, heartbeat, etc.)

We could go on with this, without getting into more "conventionalized" uses ("the real-estate game," "give someone a game," "the game is up," etc.). And obviously the designatum of game can be the argument of the more general functions described above; it can be used to refer to a game-type or game-token, and so forth.

This brings up another point, which I should note in passing. The processes by which word-uses are derived are generative and recursive. As we noted in our discussion of the mechanics of ostensive reference, given functions can be freely combined to form composite RF's. Consider 58 and 59, for example:

58. IBM went up three points last week.

59. In John's painting, the tenor is wearing a wig. 58 involves an RF derived from a composite that could be rendered as "price of stock of"; in 59, the composite involves, first a function from voice ranges to persons, and then a function from persons to representations of persons. Note that there is no constraint on the order in which functions shall be computed, as shown by 60:

60. That saxophone was a hallmark of 30's swing. 60 could be used to describe two situations. The subject NP could refer to the use of a certain type of saxophone (say, a soprano). Then the RF would take us from a particular saxophone to its type, and from there to the use of that type. Or it could refer to a way of using a saxophone, in which case the RF would first take us from a saxophone to the way it was being played, and then to a type of playing. We may even find the same given function used twice in a composite: a photograph of a statue of a saint can still be called "a saint." It should also be noted that RF's can apply cyclically to complex constituents. The phrase "a rare virus," for example, can be used to refer either to a virologically-caused disease that is rare, or a disease caused by a virus that is rare. In the second case, the function "x causes y" must operate on the entire constituent, not simply its head.¹¹

2.2.3 The Indeterminacy of Meaning

The discussion in the last section, while inadequate in some respects, will enable us to address the question we asked at its beginning: is it always possible to compute the meaning of a term? Recall the procedure by which we do this. Given a term $\underline{t}$ which is used to refer to $\underline{a}$ and $\underline{b}$, we compute the cue-validity of the "best" function $\underline{f}$ such that $f(a) = b$, and then compute the cue-validity for $\underline{a}$ of its inverse $\underline{f}^{-1}$ on $\underline{b}$. Then if the cue-validity of $\underline{f}$ for $\underline{b}$ is greater than the cue-validity of $\underline{f}^{-1}$ for $\underline{a}$, we assume that $\underline{t}$ designates $\underline{a}$; if $\underline{f}^{-1}$ is greater, that $\underline{t}$ designates $\underline{b}$. In some cases, like those we discussed in the last section, this computation yields a clear result. But there are other cases of polysemy such that the cue-validity of possible RF is either necessarily or contingently the same as its inverse.

Consider some of the general functions we talked about before. Should we say that calendar watch designates a set of particulars, or a type of thing? Which of its uses, that is, is "metaphorical?" Obviously, there is no linguistic or epistemological grounds for going one way or the other; whatever we said it designated, its referring uses would be the same. The function from a set of tokens to a type has necessarily the same cue-validity as its inverse.

Or consider the indeterminacy of words used to refer to linguistic objects. Does sentence designate a phonetic object, an orthographic object, a semantic representation,

a surface structure, or a deep structure? Or does it designate an ordered triple (quadruple, pentuple) of these things, or the relation that holds among the members of such sets? We might answer in any way, but there would be no empirical basis for our choice. Given the beliefs that determine the uses of sentence in the English-speaking community (which are relatively homogeneous, as these things go) the word will have the same set of uses whatever it is taken to mean.

There's a difference between the cases, of course: the type/token indeterminacy is as immutable as the laws of logic, while the form/content indeterminacy of a word like sentence is owing only to the nature of linguistic signs. When we turn to the indeterminacy of a word like book or dissertation, things seem to be even more labile; it would want only a few minor changes in our beliefs about the world--say, that it should not be required that a dissertation be written down--to resolve the indeterminacy. But that is not how the world is, and given our actual beliefs, there is no way to determine which kind of thing is the designatum of the word.

There are many other such cases, some of which we touched on earlier. Does window designate a kind of hole, or the thing that goes in it? Is sponge a kind of substance or a kind of artifact? Does Cadillac name a kind of car, or a car manufacturer? Do captain, mother, and doctor designate roles, or the people who fill them? Or consider the a and b pairs in 58-60:

- 58a. He is interested in foreign policy.
- b. We need a new foreign policy.
- 59a. More force is required to open the jar.
- b. Linguistics can be a force for good or evil.
- 60a. I am in favor of vice.
- b. Smoking is a vice.
- 61a. Fire is hot.
- b. He set a fire.

Here again, there seems to be no grounds for taking either the "mass" or "count" use of the term as prior.

We said earlier that game designated a kind of activity, but there is no reason for not saying that it designates the rules--whether fixed or flexible--that govern such activities. After all, the category of activity is defined by the presence of such rules, and the rules are distinguished from other kinds of rules by the purposes of which the activities that they generate. Which of 62-63 is "conveyed" or "metaphorical," then?

62. The game lasted an hour.

63. The game is simple.

Or take gossip. Does it name an activity in which a certain kind of information is conveyed, or the kind of information that is conveyed in that way? (Once we answer that, we can decide whether the verb or noun is prior; similarly report, talk, etc.) Indeterminacies like these pervade the lexicon.

Any mention of "indeterminacy," of course, must recall Quine's arguments for the indeterminacy of translation; but the point I am making is not directly connected to the problems he raises, as I read him. Quine's well-known example involves a Kabala native who utters "gavagai" in the presence of a rabbit; he argues that there is no principled basis for deciding whether what the native is referring to is a rabbit, or a rabbit-stage, or a collection of rabbit-parts. And since he accepts that words used to refer to different classes of things must be different in meaning, he argues from these cases to the indeterminacy of the meaning of gavagai in the native's language. Our argument is quite distinct from this one. We could accept that the native's intended referent was invariably a rabbit particular; in any event, as Quine has repeatedly emphasized, the "inscrutability of reference" under some set of clever permutations is no stumbling block in the path of a linguist who is willing to embrace "the appropriate fictions." But even if we are quite certain that the native is referring to a rabbit-particular with this utterance of gavagai, we still couldn't say with any assurance that gavagai designated a class of rabbit-particulars. For on another occasion, the native might use gavagai to refer to a rabbit sub-species, saying (in Kabala) "That gavagai has been dying out for some time now," and there will be no way to determine when the RF is the identity. Note, moreover, that we cannot ask the native any questions that

will resolve this indeterminacy, even if we are willing to assume that he cognizes the world exactly as we do. If we ask, "Is a gavagai a species of rodent?" he will assent; but if we then turn around and ask, "Well is a gavagai the same as that little animal over there?" he will also assent. Nor can we ask for help from what Quine calls the "apparatus of reference"; the plural gavagaim could be used to refer to sets either of rabbit-types or rabbit-tokens.¹³

We can anticipate several sorts of objections to our proposal that meaning is indeterminable. First, it could be argued that the indeterminacy could be resolved given a sufficiently comprehensive list of word-uses; that we have not cast our net wide enough. I do not see how this argument could be brought to bear on critical cases like the type/token indeterminacy, since all possible referents must be of one or the other order. In other cases, however, this line of attack has a certain plausibility. For example, we might want to argue that since book can be used to refer to a bound volume of blank pages (say a sketch-book), and since this use would be licensed only if schemata operated on the "volume" psense of book, we can conclude that the "content" psense is derived. But we also note that book can be used to refer to a non-physical representation of a book-content; thus we could say, "His book is now on its way to Alpha Centauri, by radio." Any use of book that is licensed by the assumption that it designates either category, as it

,

happens, will be acceptable, and so on for the other terms that we have used as examples.

Note also that the indeterminacies in question have nothing to do with heterogeneities in the beliefs of members of the speech-community. For some indeterminacies (type/token, content/inscription), there is every reason to suppose that the relevant beliefs are in fact uniform for all speakers in all communities, Quine's permutations notwithstanding. For other indeterminacies, it is true, differences in belief may lead to different analyses; thus there is no necessity that window or game be indeterminate. But it is a matter of fact that virtually all members of the community have the same relevant beliefs about windows and games; even if there were no variation at all, the indeterminacy would remain. As it happens, variation does intercede to make the determination of meaning harder in heterogeneous practice than it under the idealization we have adopted here. But such variation is not the source of indeterminacies, just a confound to the computations whereby meaning is determined.

Finally, it could be argued that although there are no linguistic grounds which will allow the resolution of indeterminacies, there may yet be grounds in metaphysics or epistemology. Thus a philosopher who refused on metaphysical grounds to countenance at all the hypostasis of types might argue that the "token" uses of common names must be prior, and a linguist who refused to allow the existence of either

deep structures or meanings might insist that sentence designates a phonetic form. Strictly speaking, of course, such objections are irrelevant to our claim that there are no empirical grounds on which meaning indeterminacies can be resolved. But we should also note that this sort of objection invites confusion over the difference between putative reference and actual reference. It is obviously of no empirical concern whether the things that speakers believe they are referring to actually exist; we are interested only in the assumptions that a speaker makes in determining the meaning of a word, whether they are in fact true or false.

This still leaves us open to a psychological argument; someone could claim that knowledge of tokens was "cognitively prior" to knowledge of types, for example, so that the "type" uses must be derived. I do not know how such arguments could be countered in the abstract, except by saying that if the notion of "priority" involved were not developmental--in which case the objection would be irrelevant for all but the youngest speakers--it would require considerable fleshing-out before we could make sense of it. In any case, this kind of objection is not available for most cases; we would clearly not wish to argue that knowledge of window-frames was in any way "prior" to knowledge of window-panes.

Granting all this, the linguist who finds the notion of indeterminacy intolerable could still argue as follows:

"You have shown that there are no grounds on which one use of

a word can be selected as its meaning, but not that speakers do not make such selections anyway, in some arbitrary fashion. Thus it could be that one speaker should decide that window designates a kind of hole, and another that window designates the thing that goes in the hole, and so on." On the face of things, we cannot counter this claim at all. Whichever analysis a speaker picks, his use of the word will not be affected; since there is a good function from either of the possible designata of window to the other, any use of the word that is justified by the former analysis will be justified by the latter. We may ask how we can determine which of the analyses a given speaker has made, since intuitions, even if they are consistent, will likely be influenced by criteria of frequency, or by some prior metaphysical assumptions. But the proponent of this view does not have to assume that speakers have intuitions for the analyses that they make, any more than they have intuitions for the form that the passive rule takes in their internalized grammar. The claim seems to be unfalsifiable, and hence irrefutable.

If we could accept this last claim--that the (ideal) speaker chooses arbitrarily from among the possible analyses of use available to him--then we could preserve, at least in principle, the notion of the autonomous lexicon. It is true that we would never be able to determine which of the analyses to incorporate into the grammar, but that is a limitation only the theories of linguistic knowledge that we can hope to

construct, and we do not have to renounce the principle that linguistic and non-linguistic knowledge are discrete. Yet there are reasons for rejecting even this weakest version of the thesis of lexical autonomy, which grow out of a paradox that is forced on us when we adopt the Chomskian idealization in talking about lexical semantics. In articulating this paradox, we will lay the groundwork for the approach to use which will be developed in the second half of this thesis; the following section 3.1 will thus serve as a bridge between the negative arguments which have so far occupied most of our attention, and the positive arguments which are to follow.

NOTES--CHAPTER TWO

1. There is no reason to require further that the reference be "fully consummated" (Searle 1969) or "referential" in Donnellan's psense, though the difference will be of some interest below. Nor do we have to rule out cases in which the only description available to the hearer that uniquely describes the referent is, "the thing that the speaker is talking about" we may say that a speaker has successfully referred to John with an utterance of "John is a bachelor" even if the hearer doesn't know who John is. The status of such examples we can leave to the theorists of speech-acts to worry over.
2. This distinction is often ignored by linguists who talk about the "rules" of a "pragmatic component" as if the difference between pragmatic and semantico-syntactic accounts were simply a question of "terminological variance." It isn't.
3. The only reason I bother to point out the distinction is that a defender of the position that only the normal uses of words must be lexically specified is going to wind up having to grant a central position in his account to the difference between intrinsically and extrinsically manifest properties.
4. In fact, this condition is no more than the formalization of what has often been observed about metaphorical word-uses: that a metaphorical interpretation is available only when a literal interpretation has been ruled out. Cf. Bloomfield 1933, p. 149: " . . . we understand a form (that is respond to it) in the central meaning unless some feature of the practical situation forces us to look to a transferred meaning." Where previous accounts of metaphor have gone astray is in assuming that the apparatus that grinds out "deviant" readings must be different in kind from the apparatus that grinds out the normal uses of polysemous words. See section 4.3 for an extensive discussion of the problem of metaphor.
5. The weak form of II can be represented more simply as follows:

Let $\mathcal{G}$ be the set of all given functions.
 Let $\mathcal{F}$ be the set of all functions f derived from members of $\mathcal{G}$ which have been restricted to contextually determined domains and ranges A and B , such that $A_f \cap B_f \neq \emptyset$, then $f = I$. Then the set R of all

possible RF's must consist entirely of members of $\mathcal{F}$, or composites of members of $\mathcal{F}$. (I am grateful to Chiahua Pan for suggesting this formulation.)

6. If the reader wants to be able to attribute to any individual an infinite set of beliefs, we can call the theory that finite set of propositions from which the individual's beliefs can most economically be generated. I should also note that we will later see that there is reason to be chary of requiring that the beliefs against which reference is made be attributed to any of the interlocutors, but we can let this slide for the moment.
7. In passing, it is worth noting that nothing in this prevents us from saying two terms are "synonymous"--that they designate the same thing or category--even if their referring uses are different. A number of extraneous factors may influence use, so that equivalence of "meaning" does not ensure equivalence of reference. Most obviously, the phonetic shape of a word must influence its use, not only with respect to the "use/mention" function (see the following section), but also in more subtle phonaesthetic ways (see the discussion of pop in section 4.2n). Moreover, we will see at the end of section 4.3 that "usage-conditions," such as the evaluation of a word as slang, may be determined entirely by extralinguistic considerations, and so need not be listed in the lexicon in any way; thus we may conclude that marijuana and grass are entirely synonymous, without having to express the reservations that linguists often feel called upon to make in such cases. But I should also note that none of this means much; by the time we get to Chapter Four, our view of what meaning is will have become so altered that it is doubtful whether there will be any sense to be made of the notion of synonymy in the first place.
8. That is, a second convention, over and above the convention where the term is used as if it designated anything at all--i.e., the conventions whereby tongue is used to refer to a certain body part, or whereby book is used to refer to a bound volume.
9. Where a noun designates an extremely "general" or "abstract" class of things, for example, there may be some oddness in using it to refer to a generic type at all; Vendler contrasts sentences like The Incas did not have the wheel with sentences like ?Apes do not have the tool. Even so, he notes, we may conceive of discourse in which such uses are not misplaced--as, for example, when we talk about "the way the object is represented cognitively." And in any case, we can use any word to refer to a

sub-type, as when we say, "Napier introduced the slide-rule, and the tool (i.e., 'that type of tool') became widely used."

10. Just as the names of times, such as Christmas, can be used to refer to the events that happen at those times, as in Christmas depresses many people.
11. These cases are particularly interesting, as they force us to ask whether there is really any class of word-types that we can label "indexicals," and to speculate as to how these are related to terms that are indexical only in certain uses. How is now different from Tuesday? Or consider the difference between age and before, where the former can be used only to pick out a point relative to the context of utterance. Contrast:

i. I saw her some years before
 ago.

ii. I arrived in 1956 to stay, but I had already
 visited the place two years before
 ago.

All of this is utterly untouched ground, and these questions go right to the heart of the problem of reference.

12. Space does not permit us to follow this line, or to examine the ways in which this view of reference may lead us to look at compositional semantics in general. Without elaboration: words used as adjectives "refer" to functions, unless in predicate position, in which case they refer to things. So we will analyze the red house as follows: the description is being used to refer to an instance of a kind of thing that is representable as the value of a function referred to by red at an argument referred to by house. Of The house is red we will say that the house is used to refer to a color-type, and red to the same kind of thing, and that the be is here, as always, indicating the identity; there is no "be" of predication.
13. I have trouble with Quine--the appearance of Words and Objections was a great consolation to me--and so may have misrepresented his position here. If one takes his position in the indeterminacy of reference only as an argument for the indeterminacy of meaning, and assumes that the second position follows generally from his notion of the "web of belief," then a lot of what will be said in the following chapter is simply a restatement of his views, with other goals in mind. That's what Bob Schwartz says, anyway.

CHAPTER THREE

KNOWLEDGE OF MEANINGS

3.1 Introduction

As we have so far presented the problem, the indeterminacy of meaning is not materially different from analogous indeterminacies in the syntactic or phonological components of the grammar. Though there are no undisputed cases in the literature, it is easy enough to imagine a situation in which either of two distinct phonological rules P_1 or P_2 should be available which would grind out exactly the same output, and which would be such that no criteria of simplicity, markedness, or constraints on the form of grammars should enable us to choose between them. (David Stampe is credited with a joke, well-known to phonologists, which turns on just this premise.) And while linguists might be unhappy over the existence of such cases, there are certainly no grounds for ruling them out in principle; we would simply be forced to say, under such circumstances, that the ideal speaker might incorporate either P_1 or P_2 into his grammar.

To see why the two kinds of indeterminacy are not parallel, we must pursue the analogy. We have been assuming, standardly, that the grammar of the ideal speaker-hearer constitutes a part of his theory about the linguistic practices

of other members of the community; that it is an answer to the question, "What system of rules might I most rationally adopt which would enable me to generate the corpus of utterances with which I am acquainted?" (We are continuing to ignore, of course, the problems posed by a heterogeneous speech-community.) In which case any system that satisfies criteria of economy and descriptive adequacy will do as well. But suppose we put the question differently, asking instead, "What system of rules would other ideal speakers in the community adopt etc.?" Then we run into a problem, for if the speaker knows that either P_1 or P_2 could equally well be incorporated into his grammar, then he knows that other speakers will have reached the same conclusion, and he will have no way of knowing which system they will have (arbitrarily) adopted. In that case, the indeterminacy cannot be resolved by fiat; the best that any speaker will be able to say is that the grammar contains either P_1 or P_2 .

There may seem to be no reason for putting the question in this second way, especially as it leads to just these undesirable complications. So long as a speaker has internalized a system of rules that enables him to produce utterances that will be understood by other members of the community, and to understand their utterances, why should he care if the systems used by other speakers--even other ideal speakers--are isomorphic to his own? Where phonological and syntactic rules are concerned, I think, there is no reason

to care at all, but meaning is a very different sort of affair. And we can see this if we consider the consequences of the idealization under which we have heretofore been operating.

In constructing the grammar of a language, even in a speech-community in which there is absolute uniformity of linguistic practices, the ideal speaker must always take into account the non-linguistic characteristics of members of the community. To take a trivial example, our determination that the unacceptability of multiply self-embedded constructions is not due to a grammatical constraint is made on the basis of assumptions about human memory limitation; if English were spoken among machines with a memory capacity several orders of magnitude greater than our own, we would expect linguistic practices to be different, with no corresponding change in the rules of grammar. But if such devices avoided multiply self-embedded constructions with the same regularity that human speakers do, then we must assume that their grammar was different from ours, and included a restriction on the degree of self-embedding that is allowed. To construct a similar argument in phonology, we need only consider a community in which there was radical systematic variation in the anatomy of the vocal tract.

Faced with a heterogeneity in the linguistic practices of a community, then, the ideal speaker may attribute it to one of two factors. Either the entire community conforms to

the same set of linguistic conventions, and the heterogeneity is due to some difference in the non-linguistic faculties of its members, or the heterogeneity is due to the existence of two competing systems of linguistic rules. But where phonological and syntactic variation are concerned, this observation is pretty much vacuous. Where a speaker is confronted with some systematic variation in the pronunciation of a given form, he will automatically attribute it to a difference in linguistic rules, rather than to a systematic difference in the anatomy of the vocal tracts of members of different sub-communities. (Unsystematic variation, such as might be encountered among persons with speech impediments, or chronic drunkards, is another story.) By the same token, it is unthinkable that the difference between American and British English in the syntactic behavior of main-verb have could be linked to any difference in the perceptual mechanisms of British and American speakers; clearly, two different linguistic rules are involved. It is not presently of interest to us how the speaker incorporates knowledge of variation in such rules into his grammar or grammars, or whether and when he may choose to do so; our point is simply that systematic phonological and syntactic heterogeneities are almost always the consequences of the existence of competing systems of linguistic rules. For this reason the idealization to the "homogeneous speech community" in phonology and syntax is not overly severe, since we will always

presume, in any event, the relative homogeneity of the community with respect to the relevant factors that influence performance.

But use is different. We have seen that the uses to which a word may normally be put, and the determination on the basis of those uses of what the term is assumed to designate, depend on the beliefs of speakers about the normal referents of the term. Speakers who have differing beliefs about newspapers, for example, may not use newspaper to refer to a kind of company, or may analyze the word as designating the company, rather than the publication. And while the idealization to a community of speakers whose mouths and memories are uniform is not severe, the idealization to a community all of whose members have uniform beliefs about all things is so extreme as to be almost untenable. So where a speaker is confronted with variation in the use of newspaper, he must consider the very real possibility that the use-differences are the result of differences among groups of speakers in their encyclopedic beliefs about newspapers, rather than in a difference in the analysis of what it is that newspaper designates. And by the same token, it is entirely plausible that speakers should use newspaper in a way that is entirely uniform, yet that they should analyze its meaning differently, according to variations in their extralinguistic beliefs.

Where use is concerned, then, a speaker can calculate the rationale for speaker practices only against the beliefs of the community that uses it, even where these are at variance with his own. For the speaker is obviously bound by the beliefs of the collectivity; he can never rationally assume that beliefs that he knows he holds idiosyncratically could form a part of the background of presuppositions against which his own utterances will be interpreted, nor is it rational for him to interpret the utterances of others against beliefs that another speaker would not rationally attribute to him. In a community in which there is heterogeneity of belief, the ideal speaker cannot construct the lexicon solely against his own beliefs; he must couch the question we asked earlier as, "What analysis of the use of w would other ideal speakers in this community make?"

We can contrast the consequences of the Chomskian idealization to phonology and syntax and to semantics as follows. In the first two realms, the idealization requires that we make the fairly innocuous assumption that speakers are basically alike with respect to phonetic, perceptual, and memorial considerations; and that speakers are in some way aware of these non-linguistic homogeneities; then we may assume that they can reason that whatever system will best generate an appropriate output for them will work as well for other speakers. In short, they can construct a wholly private grammar. But in semantics, we can make the same assumption

only against the assumption that speakers are alike in their non-linguistic beliefs, and that they are aware of this homogeneity. That, as we have seen, is a very severe idealization; it is hard even to imagine what the members of such a community would have call to use language for, except perhaps collective prayer.¹ Only in such a community is it possible to talk about a "private lexicon," which a speaker may consult without regard for the characteristics of other speakers, and into which he can arbitrarily incorporate any one of a disjunction of possible designata for each entry. But an account of linguistic competence that is applicable only under such circumstances could hardly be said to be an account of human language at all.

We conclude, then, that a rational speaker, even in a linguistically homogeneous speech-community, may never be able to determine the form of the conventions governing the uses of words; the best he will be able to say is that, say, "w is used as if it designated either a or b." This conclusion is not without independent interest. For one thing, it entails that we cannot at once take the grammar to be a characterization of linguistic knowledge, while at the same time assuming that its semantic component pairs sentences with specifications of truth-conditions. So long as the meaning of every term were determinate, we could still cling to the notion of a truth-functional semantics for natural language; we could say, for example, that sentences like The newspaper

fired John were literally false, but that they could be used to implicate true propositions. But we cannot talk about the literal truth or falsehood of the sentence John broke the window, because we cannot say what the word window designates. And this in turn, has consequences for philosophy of language and related areas; for one thing, it suggests an empirical argument that proponents of the "functionalist" account of meaning--Strawson's "theorists of communication-intention"--may be able to bring to bear against the "theorists of formal semantics." It should be noted, however, that the application of these conclusions to artificial and technical languages and sub-languages is arguable, and that many of the issues that concern these philosophers are in large measure independent of any empirical assumptions about natural language use. At the same time, a theory of natural language is concerned only with accounting for how it is that speakers are enabled to assign unique truth-conditions in sentence-tokens, which is all, after all, that its users must be able to do. It is no relevance whether they should be able to make such assignments to sentence-types, because types themselves are never uttered.

The indeterminacy of meaning does have further empirical relevance, and we will come to talk about it again. But the arguments that we raised on its behalf have much graver and more immediate empirical consequences, which must occupy our immediate attention. For if the theory of the lexicon is a complement of the theory of communal beliefs, which together

provide the theory of linguistic practices, then the speaker's awareness of heterogeneity of beliefs must invariably be accompanied by some heterogeneity in his linguistic system. And this may extend, not only to the determination of which category is designated by a given item, but also the way in which designated categories are defined. And if a word may designate different things for different speakers at different times, then we must consider the relative costs of the different idealizations under which we may talk about the "meaning of a word" at all.

3.2 Meanings and Categories

At the beginning of our discussion of ostensive reference in Chapter Two, we said that the reference of a use of a term was determined by the simultaneous application of a theory of meaning, which provided a function from a term to the argument of the referring function--in ostension, the demonstratum; in description, the designatum--and a theory of reference, which provided a function from the demonstratum or designatum to the intended referent of the term. The theory of reference is perforce a theory about word-tokens; we determine the RF relative to circumstances of utterance, given the designatum or demonstratum picked out by the term. But the theory of meaning, it is commonly assumed, must be a theory of word types; a speaker could not know what is designated by a particular token of tiger except in virtue of some prior knowledge of what the word-type tiger means in

English. (cf., of course, the exchange between Alice and Humpty Dumpty over glory.) Accordingly, it is customary to talk about "the meaning of w in L" as a homogeneous object or relation, and to assume that the problems raised by heterogeneities are no more severe than analogous heterogeneities in phonological and syntactic rules.

But we have just seen that the speaker can analyze the meaning of a word only against his theory of the collective beliefs and practices of the community, which must be perceived, in the real world, as heterogeneous. One speaker's theory of a category a may differ from another's in any number of ways, and the speaker who is aware of such discrepancies must accommodate his linguistic behavior to changes in the background of beliefs against which the utterance is to be interpreted. Then at the extreme the speaker's theory of meaning should itself be subject to revision at every conversational turn. A word may have to be used as if it designated a category a on one occasion, and b on another, and so on. In which case we may ask in what degree we are licensed to talk about fixed "meanings" at all, as properties of word-types.

At first glance, the indeterminacies we have discussed do not seem to force such a radical rejection of the idea of type-meaning. For one thing, not all words are indeterminate in meaning; we can confidently ascribe a particular designatum to many proper names like Charles Dickens and

Columbus Circle, for example, (Common names are always susceptible to the type/token indeterminacy, and some proper names may raise other problems: France may be taken to designate either a geographical or political entity; Bulova, a company or a brand of watch.) Moreover, it is one thing to say that "w is always used as if it designated a or b or c," and quite another to say that w has no constant meaning; its meaning may simply have to be given disjunctively.

But there's another problem. If speakers have different theories of the categories designated by natural-language terms, even with respect to the criteria that define categories, what grounds do we have for saying that the categories themselves are "the same" from speaker to speaker? And since speakers must accommodate to one another's beliefs, why should we not expect that they will be constantly revising their representations of the systems of categories that words are used to refer to? Then a speaker might use w as if it designated a on one occasion of use, and a on another, and q on a third, and so on, for a potentially infinite set of categories, as background beliefs change. In which case meaning, like reference, becomes wholly a property of tokens, and we will be hard-put to know how to talk about type-meanings at all.

One could try to avoid this rejection of type-meanings in any of several ways; in the first part of this section, I'll review some of the moves that have been suggested.

Before proceeding, however, a cautionary note. We have already noted that in the philosophical literature, the distinction between designation and reference has either not been made at all, or has been drawn on purely metatheoretical, which is to say, empirically irrelevant, lines (as with Wittgenstein's discussion of red in the Philosophical Investigations.) A good part of the literature on meaning that I will be citing, accordingly, is couched in terms of problems of truth that can have no interest for us, since we have already seen that the meaning of a term does not by itself determine the classes of things to which it can either normally or "correctly" be used to refer. Nothing that we say about the meaning of bachelor, for example, can secure the truth-value of a particular token of the sentence A bachelor is an unmarried man; that must depend on what this token of bachelor is being used to refer to. What is more, the indeterminacy of meaning makes it impossible for us to talk even about a technical notion of sentence-truth in which the semantics specifies only the truth-conditions that attach to the "literal" meaning of the sentence--the meaning it would have if every term were construed as referring to its designatum. Our interest in the controversy over meaning, then, will be in some respects beside its point; we want to know how meanings are fixed, but we have no metaphysical axe to grind.

Uniformity could be imposed upon the linguistic categorization of the world on the basis of metaphysical, phenomenological, or cultural criteria; we will visit each position in turn. The first has recently become popular as a result of the work of Putnam 1975a, 1975b, and Kripke 1972; it has been more recently endorsed by some linguists (see e.g. McCawley [ms.]). In a nutshell, their account runs as follows: just as proper names designate individuals, so common names like tiger and lemon designate "natural kinds" of things. Such categories are metaphysical givens, and their membership is determined by wholly natural considerations, independent of what we may believe about them; they are simply "there." It is not necessary that speakers should know what makes tigers tigers, or that they should all have identical theories of tigers, for us to say that tiger has the same meaning for all speakers, so long as they all analyze the word as designating the same class, where identity is determined metaphysically. To "give the meaning" of a word like tiger, Putnam suggests, we need only identify by ostension or description a member of the species, and say that "tiger designates animals like that one in their essential properties."

There are several virtues to this account of meaning, which we may call "realist name-theory" (the more common "referential semantics" is unavailable to us, for obvious reasons.) In treating common names on the model of indexical

terms (cf. Putnam's formula for "giving the meaning" of a word), we relieve ourselves of having to tie the categorization of the world of which language makes use to the theories of the world that individual speakers may come up with; words indicate categories, but do not describe them.² The individual speaker, Putnam suggests, leaves to specialists the task of determining the criteria by which categories are defined, and agrees to abide by their decisions: in the "factory" of our speech-community, "some people have the 'job' of wearing gold wedding rings, other people have the 'job' of selling gold wedding rings, still other people have the 'job' of telling whether or not something is really gold." (1975b, p. 227). Meaning is thus collectively determined, in virtue of the "division of linguistic labor," a point important to us, as the individual speaker can then be assured that other speakers will conform to the uniform categorization prescribed by authority.

The trouble with this view is that it is unrealistic, so far as the majority of natural-language terms are concerned. It is true that most of the words that Putnam and Kripke talk about--gold, tiger, lemon, and heat, for example--would normally be analyzed as designating natural kinds, and that speakers would leave to specialists the job of determining what is essentially true of those kinds of things. And in specialized sub-languages, such as those used in the physical sciences, the universe of discourse may be restricted so

as to allow only talk of such categories, in which case we may embrace the realist account. But realism has nothing to do with the way other words are defined, as Chomsky 1976 points out. He allows that words like tiger may be indexicals, and agrees with Putnam and Kripke that the designata of words--what he calls the "categories of common-sense understanding"--may be givens, and that speakers need not know what properties a thing must have in order to be a member of a given class. But he balks at saying that these properties are determined by metaphysical necessity, suggesting, rather, that the criterial features for membership in a given class are determined by some (unspecified) "interaction of the language faculties and other faculties of mind." In natural language, he argues, the determination that a thing is or is not "nameable" involves:

assumptions about the nature of the thing named, some conceptual and some factual. In our system of common-sense understanding, natural kinds are defined by internal structure, constitution, origin, function (for artifacts), and other properties. (pp. 44-45)

Chomsky's criticism of Putnam and Kripke is surely justified; the criteria for nameability, and the criterial properties for membership in a nameable kind, are doubtless determined in some part by some interaction of the phenomenal world with the cognitive system that is brought to bear on it, rather than by metaphysical considerations. It may be that a botanist, or learning that the things that are commonly called "vegetables" have no set of essential properties

in common that distinguishes them from other plants, would have to discard the term, or restrict its application to the smaller set of legumes. But speakers of English have proved unreceptive to such re-definitions; tomatoes are vegetables, whether or not they have the same organic structure as beans. The criterial properties for vegetable membership are not simply botanical, but functional and superficial as well. It is true that a scientific result may occasionally persuade speakers to mend their ways, but only because of contingent factors. The discovery that the whale bears its young alive would not be impressive to speakers of English if they didn't know that this property was likely to entail other differences between whales and fishes in behavior, intelligence, savor, and rate of spoilage, all of which are far more relevant to the interests that speakers have in mind when they distinguish among classes of animals.

But Chomsky also seems to say in this chapter that the criteria around which natural kinds are constructed are entirely determined by phenomenological factors; it is on this basis that he says that the truth of the sentence "Nixon is an animate object," (pace our own refusal to countenance the whole idea that sentences can have truth-values) follows, if not de dicto, then from some "necessary connection between the categories of common-sense understanding." In effect, Chomsky would replace metaphysical necessity with epistemic necessity. And this view fails as surely as Putnam's,

because it ignores the contribution of convention and social practice.

We classify tomatoes as vegetables, for example, because they are generally served with the main course. (At least, this is the criterion that the Supreme Court took as central in a 19th century decision.) It would want only a widespread change in eating habits, then, to cause us to re-classify them, without any change in their essential nature, or in the way we perceive them (nor would we say, I think, that the meaning of tomato had changed as a result.) Nor is it defensible to make the weaker claim that the kinds of criteria in terms of which designata shall be defined are in any way necessarily determined; the difference between the criteria which define vegetables, weeds, and conifers follows only from the very contingent differences in the uses we make of them, or the kinds of trouble they cause us.

There are many words, certainly, such that we can be assured that perceptual uniformities among speakers will assure uniformity of the designated categories for all speakers; particularly those which designate "basic" or "abstract" objects. Thus it is reasonable to assume that verbs like come, give, and see or adjectives like bent and rough will be analyzed as designating the same things on all occasions by speakers, as a consequence of purely phenomenological considerations. But we cannot rely on such considerations even to provide the uniform definiens for words like

tomato, and once we come to words whose designata are social institutions--not only Congress, but baseball and jazz--it should be clear that there is no purely phenomenological basis for the structure of the system of categories. The truth of "baseball is a game" or "baseball is a physical activity" could follow only from convention, either linguistic or non-linguistic. This is not to say that the kinds of things named aren't defined on a system of psychological primitives, whatever their nature--and origin--may be. In a later chapter, Chomsky presents a somewhat different picture of the way in which the organization of mind contributes to the construction of the "categories of common-sense understanding":

The notion "physical object" seems to be quite complex. At the very least, some notion of spatio-temporal contiguity seems to be involved. We do not regard a herd of cattle as a physical object, but rather as a collection, though there would be no logical incoherence in the notion of a scattered object, as Quine, Goodman and others have made clear. But even spatio-temporal contiguity does not suffice as a general condition. One wing of an airplane is an object, but its left half, though equally continuous, is not. Clearly some Gestalt process, or notion of function is playing a role. Furthermore, scattered entities can be taken under some condition to be single physical objects under some conditions: consider a picket fence with breaks, or a Calder mobile. The latter is a "thing," while a collection of leaves on a tree is not. The reason, apparently, is that the mobile is created by an act of human will. If this is correct, then beliefs about human will and action and intention play a crucial role in determining even the most simply and elementary of concepts. Whether such factors are involved at early levels of maturation I do not know, but it is clearly an empirical issue and dogmatic assumptions are out of place . . . (p. 203)

We can easily accept most of this. It's reasonable that there should be a universal repertory of basic concepts--properties, relations, and so forth--and that the criteria that define natural kinds are generated over the basic system. (We don't have to worry, fortunately, about where the system comes from, though it is hard to share Chomsky's faith that the problem will some day be empirically resolved to everyone's satisfaction.) But the basic system doesn't determine its application, as Chomsky shows here. Whether or not we decide that something is a physical object, for example, depends on factors quite independent both of the cognitive system and the phenomenal world. But even given universal agreement that something is, or isn't a physical object, we may or may not choose to think that it is criterially so. The truth of the statement "baseball is a physical activity" doesn't follow from the fact that all speakers would perceive all actual examples of it as a physical activity, if they chose to ignore this property in constructing the class. And similarly for "baseball is a human activity," "Baseball is an activity of finite length," etc. (It may be convenient to think of the final categorization as a description of the world, and the basic system as the grammar of the language in which the description is written. Following this analogy, structural linguists have made a great deal of progress in characterizing meanings relationally, an enterprise that is strictly orthogonal to our present concerns. See note.)³

If we can rely on neither mind nor nature to supply the criteria around which categories are constructed, we are forced to turn to culture, which could operate to fix the categories designated by terms in any of several ways. In some cases, instances of an artifact or institution will be so widely uniform throughout the community that speakers must perceive it uniformly. Thus we may assume that there will be little disagreement about the criteria by which we assign membership to the categories of forks, books, or photographs, since it is commonly--we could almost say universally--believed that these things are used and perceived in the same way by virtually all members of the community. (At the very least, there will not be enough variation in beliefs to trouble any idealization to the meaning of the word-types.) In still other cases, there may be a common agreement that the determination of what does or does not constitute a member of a category will be left to the judgment of certain specialists. Just as we may leave the determination of the criterial properties of gold to metallurgists, we may leave to the determination of the criterial properties of the categories of Fords or misdemeanors, to experts in other fields. We may then assume common consent that a word like Ford is an indexical for a kind of car, whose criterial properties are whatever the recognized experts--in this case, the people at Ford--say they are. Again, we can assume that the word designates the same class

of things on all occasions of utterance. Finally, it may be that the definiens of some categories are fixed by explicit stipulation; though this is really only a variant on the appeal to authority. Thus the designatum of strikeout is defined by criteria prescribed by universally recognized authority; the only difference between strikeout and Ford (the car) is that the criteria for correct application are generally available to members of the community.

Yet taken together, these considerations do not assure the homogeneity of word-meanings for all cases. Consider the problems that arise when we try to say to identify the categories designated by such ordinary words as joke, vanity, gossip, tabloid, smoothie, amateur, colonial, hippie, stew, or Canadian. In each case, there is considerable variation among speakers in beliefs about what does or does not constitute a member of the category, nor is there any recognized authority to whom we can appeal for resolution. Take jazz. I may believe that the category includes ragtime, but not blues; you may believe exactly the opposite. After all, we will have been exposed to very different sets of exemplars. And absent a commonly accepted authority, we must construct our own theories of the categories, most probably in the light of varying degrees of musical sophistication. At what cost, then, can we talk about "the meaning of jazz" in the language? Two obvious idealizations suggest themselves, but neither leads to an understanding of the speaker's knowledge of meanings.

First, we could talk about jazz as having a single meaning in langue, where it designates a discrete cultural phenomenon; we could then take the "culture" in which langue abides to be a fictitious abstraction over individual contracts and practices, following Bentham and Hume. Or we could use langue as Saussure (is usually assumed to have) intended, to refer to a Durkheimian conscience collective, which has a primary existence over and above the beliefs and behavior of the members of the community.

But such a solution will take us no closer to describing the linguistic knowledge of the speaker, because the collective awareness does not abide in any individual. In effect, we are once again idealizing to a community in which there are homogeneous beliefs, and that idealization cannot be of help in uncovering the knowledge of the speaker who is aware of communal heterogeneities. It will not help, even, to say that jazz designates a "fuzzy category" in langue; that ragtime is closer to the border of the category than, say, bop. In a conversation between speakers both of whom believe that ragtime is unequivocally not jazz, and are aware of one another's beliefs, such an equivocation comes no closer to describing the meaning of jazz for them; the statement "ragtime is sort of a kind of jazz" would be simply false, so far as they were concerned.

Suppose we took the opposite tack, starting with individual speakers, and said that jazz has a different

meaning for two speakers when they disagree about the criterial properties of its designatum, or when they identify jazz by means of stereotypes that select radically different classes of exemplars. We don't have to say that jazz has the same meaning for two different speakers only when they have exactly the same theories of jazz, or that the word has to designate exactly the same thing for both of them--some measure of rough functional equivalence would suffice us. Moreover, we could allow that speakers are fairly lax in application of criteria and stereotypes, so that the designatum of jazz has fuzzy boundaries even for individuals. Then we could continue to subdivide the speech-community into different dialects, until the idealization of langue was not intolerably severe; for each dialect, we could say that the langue was simply a statistical generalization over parole. In which case, we could argue that knowing the meaning of jazz in langue enables us to determine, with a high degree of probability, what any token of the word was being used to designate. This is roughly the approach that has lately been invoked to deal with variation in phonology and syntax, and it is worth following it up to see why it comes to grief. I'll deal with the problems in order of increasing severity.

First, how do we manage to understand uses of jazz when the speaker's dialect is different from our own? Suppose, for example, that I am talking to an older speaker who

uses jazz so generously as to include all Black-influenced, highly rhythmic, American popular music, including rock, disco, R & B, and so forth. A notion of dialect-borrowing would help here; I understand his usage on the basis of my prior familiarity with his belief-system. And the same principles will enable us to explain how I can come to use the word consistently with his theory of jazz, when I am speaking to him in his dialect. This case has a straightforward analogue in phonology and syntax, as when I understand an Englishman's /haf/ to be equivalent to my/haef/, or when I use his pronunciation to facilitate communication.

Some more complicated cases admit of the same kind of resolution. The different "dialects" in which jazz has a different meaning would not correspond very well with "dialects" defined phonologically or syntactically, so that it would be difficult to determine, in normal conversational contexts, just what a given speaker or hearer believes jazz to be. Where the speaker is ignorant of his hearer's beliefs, and may plausibly assume that the hearer is ignorant of his own, it is reasonable that he should use the word in accordance with some "neutral" belief-system that is most likely to be accessible to both parties. And the reasonable hearer should be able to follow this strategy, and identify the designatum accordingly. Here again, the parallel to other areas of dialectology is clear: it is as when two Frenchmen abandon their respective local patios to converse

in standard Parisian.

I'm using "dialect borrowing" here in a fairly loose way, to refer to any situation in which a speaker may make use of knowledge of more than one homogeneous linguistic system. For cases where variation is fairly systematic throughout a community, however, we might prefer to render all this in terms of the model of a single heterogeneous "variable competence," as introduced by Weinreich, Labov and Herzog 1968, and developed by Labov and others in subsequent work. So we could say that the designatum of jazz is defined for a given speaker on a given occasion of utterance by the application of criteria whose weighting is context-sensitive, just as we say, for example, that the choice of the vocalic nucleus in any New York speaker's pronunciation of bad or ball is determined by variables sensitive to contextual factors. But while it would be an interesting exercise to construct a model of "variable semantic competence" to deal with cases like the above, it would also be a futile one, because of a fundamental difference between semantics and phonology or syntax. Whether we treat formal variation in terms of dialect borrowing, or in terms of variable rules, we presume a specified set of variants from among which a particular token is selected; a speaker can have only two (or three, or six) variant pronunciations for a given word-type in his repertory. (or analogously, can be familiar only with a certain number of other dialects.) He cannot create

new alternatives de novo. Up to here, we have been talking about meaning as if the same situation held, so that a speaker using a term like jazz had merely to choose from among a determinate, even small number of possible categorizations. But there is in fact no determinate limit to the number and variety of belief-systems and resultant systems of categorization that can co-exist within a given community, nor is there a fixed "standard dialect"--a single set of beliefs that is equally available to all speakers. (This is just the problem of the inaccessibility of langue that we mentioned above.) A speaker can never be sure that he knows how a hearer categorizes things, and vice versa. So where categorizations vary, new systems must often be created ad hoc, according to the interlocutors' apprehensions of one another's experience, and the exigencies of context.

Suppose we teach two speakers the meaning of jazz according to Putnam's formula, shorn of its realism: we present them with a set of exemplars, and say "jazz designates this sort of music." (And suppose that we follow a similar procedure for the words rock, blues, boogies, swing and so forth.) Even though the exemplars are the same, the speakers may come up with very different theories of what jazz is, since the question of which properties are criterial (orchestration, tonality, chromaticity, rhythm, etc.) can only be answered relative to a particular set of interests. So we can never be assured that they will agree on the

classification of new cases. And in practice, of course, we are much worse off, because speakers don't construct such categories over the same exemplars, and their knowledge of the categories with which jazz is in opposition is equally variable. How can they ever come to terms?

An example will help here. Suppose that my English uncle stops by to visit me in New York--a man thirty years my senior, who manufactures umbrellas in Salford. I give him a drink, and then move to the record player, asking, "Do you like jazz?" How in the world does he know what sort of music I am talking about? Suppose, for the sake of argument, that he is much more familiar with music than I would reasonably take him to be. Then he could reason as follows:

1. The categories into which my nephew chops up music are not isomorphic to my own, and he knows this, and knows that I know it.
2. My nephew has only imperfect knowledge of what jazz designates in my conceptual system, and vice versa. And we know that we know this, etc.
3. There is no single belief-system such that "jazz" is clearly defined in it, and such that each of us could be assured that the system was accessible to the other. And we both know this, etc.
4. My nephew cannot be using jazz to designate a category in his own conceptual system, since he knows that I do not have access to this system; if he were so using the term, then no purpose would be served by his polite inquiry, since my answer would not enable him to accomodate my desires--he would still not know what I wanted to hear.

5. Nor can he be using the term to designate a category in my system; ignorant of it, he would not know what record to play.
6. Therefore, he must be using the word to designate a category of music whose definiens he thinks I would be able to construct on the basis of contextual assumptions.
7. It is reasonable to assume that the category is to be constructed around a set of exemplars E such that he may expect that I will be able to determine which members of the set of exemplars of jazz that I am familiar with are members of E, and so on. In other words, I must determine which of the things I call jazz are most like that subset of the things that he calls jazz that he would take me to be most likely to be familiar with. Ideally, we will have to be able to find some exemplars that we have in common (thank goodness he didn't ask me about musical comedies.) Bix is too old for him; Chick Corea is too young for me. He does not know that Art Tatum ever played in Europe, and has probably never heard of Django Reinhardt. Say Ellington.
8. It is also reasonable that he should expect me to construct the category along the fairly gross criteria that are available to a musical novice, since he has no reason to credit me with any special expertise. What properties would an unsophisticated listener take as criterial of the kind of music exemplified by Ellington?
9. He cannot expect that I should have a rich or sophisticated theory of popular music types; moreover, he can expect that theory to be different from his own. (He may know about swing, but he cannot know whether contemporary listeners considered it a subordinate of jazz. He cannot know whether I have heard sufficient blues to have considered its place at all.) So the category in question must be restricted to exemplars that are like the ones around which we are assuming that the class is being constructed, in those properties that an unsophisticated listener would find most salient.

By making such assumptions about each other's beliefs, about each's beliefs about the other, and about our conversational purposes, my hearer and I may come to approximate a category of things that corresponds to this token of the word jazz. (I don't know if you could ever say that we had arrived at a meeting of minds.)⁴ But now we can ask some other questions. Why does it matter what I and my hearer think jazz (really) is? My theory of jazz is relevant only insofar as it happens to coincide with the beliefs that somebody else might reasonably attribute to me; I can never use the word to refer to a wholly private category. What is the point, then, of looking for "meanings" in the minds of speakers?

At this point, we are once again arrived at the arguments with which we began this chapter. We saw there that in a heterogeneous speech-community, the determination by a particular speaker that a word designated this or that category must be made on the basis of his apprehension of the beliefs of other members of the community; his knowledge of linguistic conventions is necessarily the complement of his knowledge of non-linguistic beliefs and practices. In this section, we have simply extended the argument to the determination of nature of the designated categories themselves; as the extralinguistic presuppositions that form the background for the utterance vary, so must the speaker's determination of the form of the linguistic conventions that

are being appealed to.

Now it could be objected that in picking examples like jazz, we have unduly clouded the issue, and the objection has some validity. There is less variation in systems of categorization where the categories named are conceptually more "basic" (such as with whole, bring, or object); or where the uniformity of exemplars is guaranteed either by natural law (tiger, water) or by the uniformity of certain cultural institutions (bachelor, doctor); or where a single small set of exemplars is widely known throughout the community (assassinate, movie star). Yet these cases are only different from jazz in the degree to which speakers can assume some uniformity in the cultural belief-system. Even for the most basic words, after all, like cause or line, we can assume that the speaker internalizes a uniform meaning, not simply because phenomenological considerations force on him a single characterization of the designated category, but because he assumes that these same phenomenological considerations affect other speakers just as they do him. If we could make the same assumptions about all of the considerations that lead to category construction, as we can about the performance considerations that enter into the construction of the phonological and syntactic components of grammars, then the idealization to a private linguistic system, realized in competence, would not be pernicious. But we cannot do so, and a "private" theory of the speaker's

knowledge of the meaning of tiger, water, or line that cannot also be a theory of his knowledge of the meaning of jazz, vanity, or Canadian cannot be much use to us.

Yet even if meaning is strictly speaking a property only of word-tokens, there must obviously be some idealization against which we can talk about the meanings of word-types. It may be that a word like jazz is used as if it designated a different category on every occasion of utterance, but there are limits; if someone uses jazz as if it designated a kind of painting, or even an inappropriate musical category, we must protest that he does not know what the word means. And even if he does use jazz as if it designated an appropriate category, but fails to understand other of its normal uses--as when we talk about "one of the best bassists in jazz"--we must protest that he still doesn't know how the word is used. (Note that the two formulations are not equivalent, on our view.) In either case, what we are really saying is that he has not internalized to a sufficient degree the constraints that the community imposes on the kinds of beliefs that can be brought to bear on the interpretation of utterances in a "normal" speech-situation.⁵

I'll begin Chapter Four by setting up an idealization for talking about this sort of knowledge, taking a leaf from some game-theoretic notions; in the course of doing so, I'll try to show how such an idealization can be brought to bear to resolve several problems that have recently occupied the attention of theoretical linguists.

NOTES--CHAPTER THREE

1. I am grateful to D. Terence Langendoen for reminding me of the importance of collective prayer.
2. Putnam suggests, in fact, that words like tiger should be analyzed as containing "hidden indexicals."
3. Since Saussure, most work in linguistic semantics has been directed at the characterization of the structure of the lexicon, rather than at problems of reference. Following the usage of Ricoeur 1968, we would say that the study of semantics has been subordinated to the study of semiotics. The two enterprises are strictly independent; for example, it may be reasonable to say that the "meaning" of bachelor can be represented as a function of the meanings of male, married, adult, and so forth, whether or not we can explicate the meanings of the latter terms. In practice, however, structural studies of the lexicon are going to be interesting only where we can assume a fairly uniform analysis of the meanings of the systematic primitives. It is certainly no accident that the most progress has been made in characterizing the meaning-relations that hold among words like kinship terms, or verbs designating "basic" actions. In the first case, the repertory of semantic primes is constructed over concepts, like "male" and "parent of," which mind and nature conspire to make universal and uniform to a tolerable degree; moreover, the "semantic field" is reasonably well-defined, so that the vocabulary forms a natural paradigm. And finally, the system of categorization that speakers ultimately come up with may be highly informative about some general features of the culture, though on the last point, of course, there has been some dispute. Similarly, verbs like come and bring, or buy and sell, designate categories defined by criteria that appear to be either phenomenologically uniform ("move," "goal," and so on) or culturally so ("money," for example). Structural studies of the lexicon must inevitably fare best with such "one-criterion" words.

By contrast, structural semantics has done less well when analyzing lexical domains in which the world does not supply a discrete field, or in which the definition of the primitive categories may be subject to more variation from speaker to speaker. Lehrer's 1969 analysis of cooking terms, for example, winds up mired in a wash of uncertain oppositions, and "distinguishers" whose interpretation is not likely to be uniform from speaker to speaker. That is inevitable; not only do

practices vary from place to place, but also from speaker to speaker, even within the same household, and there is no reason to expect that, say, parboil or poach should occupy a fixed place in the uniform "semantic system" or English. (Much the same kind of criticism could be leveled at Katz's 1977 suggestions as to how the set of verbs designating speech-acts might be lexically represented;)

There has been some discussion in the literature over how structural analyses of the lexicon ought to be interpreted; in particular, over whether these are analyses of the meanings of words or the criteria around which the categories designated by words are to be defined (or, to put it in other terms, over whether lexical decompositions should be represented in the lexicon, or by means of meaning postulates.) Obviously, the second position is more congenial here; we would be most consistent in saying that bring, for example, designates a category defined by criteria that are best represented as the value of a function that we may call CAUSE on the criteria that define the category designated by come. But again, our position on reference makes this argument in some measure irrelevant; whether or not the bring/come relation is represented linguistically, it has no bearing on the determination of the truth-conditions attached to utterance of sentences containing the terms.

4. Speakers vary enormously, of course, in their ability to make calculations like these, which is why some of them can teach, and others only write.
5. The argument that I have been making in this chapter has its analogue, not only in philosophical debates over language, but even more precisely in current literary theory, where relativism has driven some writers towards semantic nihilism. Obviously, we would agree with those who say that interpretation varies as the reader's belief, all the more because the author is powerless to accomodate to the reader's situation. But that doesn't entail that we can't talk about "the meaning of a text," against an appropriate idealization.

CHAPTER FOUR

KINDS OF WORD-USES

4.1 Normal Uses4.1.1 The Problem of Acceptability
Judgments

A speaker's knowledge of the conventions governing word-use, we have seen, must be a part of his knowledge of the beliefs and practices of his community; the two kinds of knowledge cannot be distinguished, even in theory. In this section, I'll try to set up an idealization which allows us to talk about such general knowledge, reserving for the following section the problem of describing knowledge of linguistic conventions. But although we will not be talking here about linguistic knowledge, we will be addressing a purely linguistic problem: how shall we interpret the speaker's judgments that tokens of sentences presented in the absence of context are acceptable or deviant? The connection between the two problems will emerge presently.

Recall that in section 2.2.2, we adumbrated a taxonomy of word-uses, which we can now flesh out. We can do this by setting up a kind of scale. At one end will be uses that are entirely conventional (though we cannot always

tell which these are); let us say that such uses can be interpreted only in the light of knowledge of purely linguistic conventions whereby certain forms are used to refer to certain things. (The use of Charles Dickens to refer to the man who wrote Bleak House would be an example.) Closely related to these are word-uses that are "conventionalized," or "partially motivated," such as when we use nylon to refer to a stocking, or country to refer to a kind of music. I won't try to validate this distinction until the next section; for the moment, take it on faith that some encyclopedic information must play a part in the interpretation of the latter. Opposed to these are what we called free uses, which can be interpreted solely on the basis of non-linguistic information, given knowledge of the conventions governing some other use of the word; thus if I know that it is conventional to use Charles Dickens to refer to a certain writer, I can interpret solely on the basis of encyclopedic information the use of the word to refer to an oeuvre. Among free uses, we can further distinguish three rough classes. Under their most obvious interpretations, 1-9 exemplify "normal uses"; we would expect that any speaker would accept the uses without reservations as perfectly ordinary English.

1. We had chicken for dinner last night.
2. I haven't read Dickens.
3. John painted a calendar watch.

4. John lives in the South.
5. The turnout astounded the mayor.
6. Game is 21 points.
7. The theater is in a sorry state.
8. The alto arrived late.
9. He is one of the biggest names in radio.

By contrast 10-12 exemplify "local" uses. It is easy to imagine a context, or set of contexts, in which they could be rationally used. But their meaning out of context is less sure, and we would probably judge them to be only marginally acceptable:

10. The ham sandwich is waiting for his check.
11. The Times hasn't arrived at the press conference yet.
12. That's a lot of speaker for a small room.

Finally, 13-15 exemplify metaphorical uses, which speakers would probably judge "odd" or "deviant," but which may be interpretable relative to the aesthetic or affective purposes of a given utterance:

13. The squirrel has gone to his granary.
14. The crag frowns on the river.
15. That snake (John) would say anything to get elected.

I have introduced these categories only for the sake of convenience. The boundaries between them are rough, and I expect that others would disagree over the classification of given examples. Still, the scale, if not this particular

set of divisions, seems to have a sound basis in intuition, which might be measured in any of several ways. As we proceed from normal to metaphorical uses, we become progressively less willing to say that the use is "good" or "standard," and less inclined to say that a speaker who could not understand the use was not "natively competent" in English. In familiar terms, this is a squish.

Some background, now. The procedure of relying on speaker judgments to sort out a corpus over which a grammar shall be constructed is so widely accepted, and so convenient, that even those linguists who disavow it in principle have adopted it in practice, and attempts to supplement it by other means have so far had only a cautionary effect. And where the object of inquiry is the individual competence of the native speaker, as it has been in most recent work on syntax, it is reasonable to expect that a speaker's judgments can be used to provide an entree into his internalized system of linguistic rules, with appropriate reservations about performance, dialect mixture, correction, and so on.

It was natural, moreover, that these same procedures should have been applied to the study of "semantic competence," once generative-transformational linguists turned their attention to problems of meaning, and that the same battery of devices--asterisks, question marks, and the like--should be used to delineate the corpus of standard

word-uses that the semantic component should be asked to generate. The assumption that semantics should be approached on the model of syntax was basic to Fodor and Katz, and the subsequent controversies about degrees of acceptability, and so forth, were simply the analogues of current arguments over the nature of syntactic rules.

But we have rejected out of hand the assumptions that underlie this way of interpreting judgments about well-formedness. If a speaker's judgment that a given use is "good" is not evidence that the use is determined by the rules of language alone, then it follows that the difference between it and a marginal or deviant use can only be laid to pragmatic factors.

What does it signify that speakers should judge some kinds of word-uses "standard" or "acceptable"? Let's start with some uncontroversial examples, and let the pragmatics decide about the others. 16 would have to be judged a workaday use of the item oak:

16. Oak is a kind of wood.

(There is not even the temptation here to say that oak is here "elliptical" for oak wood. "Oak wood is a kind of wood." would be silly.) How shall we interpret this judgment? Recall that Katz and Fodor would say that the fact that an interpretation is available in "the null context" shows that it arises independent of contextual considerations, and must consequently be assigned by purely linguistic rules.

Now "the null context" is an interesting expression to have used here. If the null set is a set, then the fact that a word bears a certain interpretation in the null context scarcely entails that it bears that interpretation in the absence of context. (This is not simply a matter of semantics; the difference will be important to us below.) But let's first take up the claim that Katz and Fodor intended to make: that standard interpretations are assigned by linguistic rules, and are hence independent of context.

We can translate the claim as follows: in any context of utterance, the word oak can be used to refer to a kind of wood (provided, Katz would add, that the selection restrictions of the predicate allow such an interpretation). If this were true, it would follow (in our terms) that certain given functions here--the function from sources to products--always allow derivation of a good RF on a given argument. But this runs counter to the discussion of reference in 2.1, where we showed how the referring function is derived when the range of a given function is restricted to its intersection with the contextually-determined range of reference, and how the determination that a particular given function allowed derivation of a "good" RF could be made only relative to the speaker and hearer's awareness of one another's beliefs.

Whether a speaker can point at a niblick to identify successfully "the game of golf," for example, depends on

what other things the speaker would be likely to refer to in the context, and on whether he can count on the hearer to share his theories of golf and niblicks, and to be aware that he shares the speaker's theories, etc. (It also depends on what other possible demonstrata are available in the context; we have a higher probability of identifying "golf" by pointing at a picture of Lee Trevino if it is not hanging next to a golf club. But the determination of the domain of the RF is more sensitive to contextual factors in ostension than in description, where the set of available arguments--i.e., the set of categories designated by lexical items--is invariant to a much greater degree from utterance to utterance.)²

Strictly speaking, then, we should be able to say only that a particular given function allows the derivation of a good RF only relative to a particular context of utterance. By way of example, consider 17 and 18:

17. Water buffalo goes well with duck sauce.

18. Water buffalo makes good handbags.

17 and 18 involve the use of different RF's: from animals to their meat in 17, and from animals to their hides in 18. But both RF's are derived from the same given function from kinds of things to the substances they are made of. In either case, the RF is derived when the range of the given function is restricted to its intersection with the range of reference: in 17, kinds of meat; in 18, kinds of material.

In these examples the range is suggested, but not determined, by the linguistic context. (It would only be pragmatically odd if water buffalo were being used to refer to a kind of hide in 17, or a kind of meat in 18.) In other contexts, such as when we say,

19. Water buffalo is tough.

the determination that we are talking about one or the other kind of thing, if it can be made at all, must depend entirely on the extralinguistic context.

In a context in which a speaker might equally well be expected to refer to either hides or meat, then, the function of "X is made of Y" would not allow derivation of a good RF. But in that case, the uses in 17 and 18, which we would have few qualms about labeling "standard" general uses, are sensitive to contextual considerations in the same way as the uses we labeled "local," such as in "The ham sandwich wants his check." What then explains the difference in the judgments that speakers make about examples 1-15?

An obvious answer suggests itself. It is frequently the case that the range of reference of a word-token is just the set of kinds of hides or kinds of meat; it is very rarely the case that the range is just the union of those two sets. In consequence, words like water buffalo are often used to refer to hides or meat, and we could say that the speaker's judgment that these uses are "standard"

reflects his awareness that the use is frequent, or has a high probability of occurrence. But at best this doesn't buy any explanation (the standard criticism of this kind of "account"). It is undeniable that we constantly estimate likelihoods (if not calculate probabilities) when we talk, as when we do anything. But when we say that a use is "standard" we are not simply saying that someone is likely to produce it--no more than we are saying that it is generated by a system of formal rules. It is surely just as likely that we will hear persons referred to as "pigs" or "angels" as that we will hear them referred to as "altos" and "captains," but our intuitions, in consultation with considerations independent of sheer probability, pick out only the latter uses as standard.

4.1.2 Normal Beliefs

For Katz and Fodor, the presentation of a sentence in the null context was intended to ensure that the speaker could interpret it solely on the basis of his knowledge of linguistic rules. We have seen that this conclusion is not justified; in the absence of contextual specification, speakers still have access to the body of beliefs that they take to constitute the background against which all utterances in a community are rationally made. In this section, we will be talking about how a speaker arrives at this set of beliefs, and how he brings them to bear to interpret

tokens of sentences presented in the null context.

I will try to keep my usage consistent here, saying that what speakers do is to "judge" that a sentence is "normal," where "normal" is to be understood on the model of social, rather than statistical norms. (Most of the words that we use to describe speaker reactions to sentences show a similar pattern of polysemy. Verbs like judge, evaluate, and accept, and adjectives like standard, normal, acceptable--and above all good--all oscillate between prescriptive and descriptive uses.) We can also use "normal" to describe the system of beliefs that a speaker invokes to interpret utterances; we will assume that a speaker will judge that a particular word-use is "normal" just in case it is licensed in the null-context by the system of normal beliefs alone. But this formulation will be useful only if we elaborate the notion of "normal beliefs" in an interesting way.

By the system of normal beliefs, I do not mean the "collective awareness" of the community, though the notions are connected. Whether or not it makes sense to hypostasize such an entity, we have seen that the individual speaker cannot have direct access to it, so that it could not provide the basis on which he offers judgments about word-uses. But in locating the system of normal beliefs in the individual, we must also avoid the temptation to identify it with the speaker's own system of beliefs, since he will frequently be aware of the discrepancy between what he believes, and

what beliefs most people take as a basis for conversation. Nor need the system of normal beliefs necessarily be attributed to most other speakers, since the informant will also assume that they are aware of a similar discrepancy.

We can most easily describe the system of normal beliefs on the model of game theory, following the line developed by Lewis 1969, in treating conversation as a "game of cooperation," a game in which the players' interests coincide. Consider, for example, one of the "coordination problems" that Lewis presents as an example:

Suppose several of us have been invited to a party. It matters little to anyone how he dresses. But he would be embarrassed if the others dressed alike and he dressed differently, since he knows that some discreditable explanation for that difference can be produced by whoever is so inclined. So each must dress according to his expectations about the way the others will dress; in a tuxedo if the others will wear tuxedos, in a clown suit if the others will wear clown suits (and in what way he pleases if the others will dress in diverse ways).

In this situation, it is obviously in the interests of any one of the participants to behave as the others will; he can succeed in doing so to the extent that he can determine what expectations they will have about his and one another's conduct. It doesn't necessarily matter what he himself believes would be the correct way to dress, nor what the others believe is correct; even if each of the participants believes that it would be absurd to wear a clown suit, he will do so if he thinks the others will.

In trying to solve coordination problems of this

sort, a participant runs into an obvious regressus. Let us reduce the problem to two players A and B, and suppose that either can win only if both dress in the same way. A is then justified in wearing a clown suit only if he expects B to wear a clown suit, and believes that B believes that A expects B to wear a clown suit, and believes that B believes that A believes that B believes that A expects B to wear a clown suit, and so on. This is just the "Gricean regressus" which we encounter in stating the conditions under which a given conversational implicature will go through ("S believes that p is irrelevant, and believes that H believes that p is irrelevant, and believes that H believes that S believes . . ."). But as Lewis shows, there are winning strategies, given a sufficient number of "nested expectations," and assuming some "system of concordant mutual expectations." In the broader social setting, this system of expectations will be determined either by convention, or by the "neutral" system of beliefs that the greatest number of individuals could most plausibly arrive at. Leaving convention aside for the moment, an individual will be most successful in communicating if he can apprehend this system of beliefs, and behave as if he accepted them. Let us say accordingly, that a given speaker will attribute a proposition p to his system of normal beliefs just in case he believes that it is rational that all speakers should behave as if they believed that p, in the absence of

circumstantial evidence to the contrary.

In deciding whether to assign a proposition to the system of normal beliefs, a speaker must take into account both availability and attribution. If a belief is not familiar to most members of the community, then it cannot be assumed as part of the conversational background. Where a speaker believes that only one of a set of mutually contradictory beliefs is generally available, it must be assigned to the system of normal beliefs. For example, if someone says, "You have taken more time to paint the garage than it took God to create the world," we must assume that he intends that we interpret the utterance against the normal belief "God created the world in six days," since of all propositions of the form "God created the world in n days," this is the only one with any general currency. In this case it doesn't matter whether anyone actually believes the proposition. (By analogy, the participants in our party game might all have sufficient reason to wear clown suits if they have collectively heard one of them say, even in jest, "We should all wear clown suits to Mary's party," and if no other form of dress has been collectively proposed.) Note how the regressus pops up here. It is not sufficient that a speaker believe that a proposition is generally available: he must also believe that other speakers believe it is available, and so on.¹

But where a speaker believes that all members of a

set P of mutually contradictory propositions are equally available to members of a community (etc.), attribution must play a role; he will assign to the set of normal beliefs that member of P such that it is most reasonable that members of the community should attribute p to one another, and believe that they attribute p to one another (etc. again). So if the participants in our party game have collectively heard it suggested both that they should all wear clown suits to the party, and that they should all wear jacket and tie, they will most rationally wear the latter, since that behavior is more consistent with those beliefs about correct party dress that it would be most reasonable to attribute to any of them. (Here again, note that it does not matter that anyone should actually hold the belief, but simply that it should be rational to suppose that he holds it, and that he should know that it is rational to suppose that he holds it, etc.).

Some of these considerations will not become important until the discussion in Section 4.2. For present purposes, we will idealize a bit. Just as we can say that a particular word-use is "standard" if a large number of speakers judge it to be normal, so we can assume a large measure of agreement as to which propositions are likely to be assigned to the system of normal beliefs of any member of the community; accordingly, we can talk about "our" system of normal beliefs, to facilitate discussion.

4.1.3 Interpretation in the Normal Context

In assigning an interpretation to an utterance presented in the null context, an informant has available to him both the system of normal beliefs, and the linguistic information contained in the sentence itself. Taken together, these enable him to construct one or more "normal contexts" for the sentence; that is, a context in which the sentence could rationally be used to express some proposition, assuming as background only the system of normal beliefs. A sentence is judged normal, we will assume, only if an informant can construct at least one normal context for it. When we come to talk of a particular word-use in a particular sentence, we modify the formulation as follows: a word w is used normally in a sentence S presented in the null context, only if its linguistic environment, together with the system of normal beliefs, determine a range of reference such that the word would be rationally used to identify some individual in the range. We will say that all such ranges are the null-context normal ranges of w in S.

Let's say that a range of reference can be any set of all those things that are functionally equivalent with respect to a given interest. It may be simply a set of things in the world (the set of all automobiles, the set of all opera companies), or a set of things localized to some restricted universe of discourse (the set of all cars in this lot; the set of singers in such-and-such opera company).

The determination that it is one or the other depends on how the values of "purely" indexical elements in the utterance are assigned, and the fact that a sentence is presented in a null context doesn't rule out either possibility, since the linguistic context may contain indexicals (as in "He's the best _____ here.") that force the localization of the range.

Virtually any group of things may be equivalent for some purpose, of course, but we can make some progress if we distinguish among the kinds of interests that determine the range, according to whether they could be attributed to someone solely on the basis of normal beliefs. Members of the set of automobiles, for example, are normally regarded as functionally equivalent for any number of social purposes: as means of personal transport, in terms of their manner of sales and distribution, or as items for which a substantial part of a household budget must be set aside. By contrast, there is no collectively determined interest such that it is satisfied just by the set of all automobiles that aren't Buicks, or just by the union of the set of all automobiles and the set of all bass-reflex speakers. While a particular context might determine such a range (say when a Buick dealer is talking about adding another line), we would have to know something about the peculiarities of that context in order to restrict the range to such a set.

The null-context normal range of a word is determined

in part by its linguistic environment. Given the sentence "He parked his w in the garage," we would allow that the use of w is normal only if there is a good function from it to a member of a range of things such that they are equivalent for the purpose of being parked in the garage. We do not require that the referent be something that is normally parked in a garage, however; w could be mountain or mother here, and the use would still be normal. We will not say, therefore, that the predication itself limits the normal range beyond the sets of things of which it could rationally be predicated. (For present purposes, we can treat this class as equivalent to the set of things that are "selected" by the predicate, but see section 4.3 for a discussion of this question.) But considerations of what things are normally equivalent under an interest do play a role in determining acceptability, as we will see.

The normal range of a word may constitute a basic conceptual or cultural category of things, but it need not. For example, the set of substances of which handbags can be made is a normal range, but the properties that characterize its members do not define a basic category. Moreover, the normal range of a given word-use is often determined as the intersection of the ranges determined by several predicates, as when we say "w is tasty, and makes good handbags." Clearly, there is no basic category such that all and only its members are possible referents for w here.

Given the referring term itself, several further conditions determine whether a use is normal in the null context. First, there must be a "best" function from its designation to some member of the range determined by the predicate, where cue-validity is calculated solely on the basis of the system of normal beliefs. This condition can be formulated as i:

- i. Given a use of a word w in a linguistic context " w is ϕ " such that w designates a and " ϕ " can be rationally predicated only of members of R , the use is normal only if there is some member b of R such that b is best identified as being the value of some function f at a , given the system of normal beliefs.

Condition i serves to rule out some fairly obvious anomalies, such as "the idea slept late" or "water buffalo has elapsed." Given no further contextual specification, there is no good function from ideas to late-sleepers, or from water buffaloes to periods of time. But the condition does not tell us what to do when there is more than one good function from the designatum of a referring term to a member of the range determined by the predicate, and these are more interesting cases.

Consider 1 and 2, for example:

1. He parked his Ford in the garage.
2. He parked his stock certificates in the garage.

Both of these examples have normal interpretations in the null context, of course. But why doesn't 1 have two normal interpretations, in one of which Ford is used to refer to a

kind of stock? (We could tease out the second interpretation, of course, by further specifying the context, but then we suspend the relevant normal beliefs. See below.)

Obviously, it is important here what things normally satisfy the interest determined by the predicate. 3-4 and 5-6 provide other examples:

3. He likes to eat steel.
4. He likes to eat chicken ($\neq$ "chicken feathers")
5. Frenchmen are wet.
6. Hawaii is wet. ($\neq$ "the population of Hawaii")

Here again, it is highly unlikely that 4 and 6 would be recognized as ambiguous, except given further contextual specification. We must add to i condition ii, then, which states that uses of words to refer to things that do not normally satisfy the interest determined by the predicate are normal only if there is no good function from their designata to things that do normally satisfy that interest.

- ii. Given a use of a word w in a linguistic context " w is $\emptyset$ " such that w designates a , and " w is $\emptyset$ " can be rationally predicated only of members of R , and is normally predicated only of some subset R' of R , then if there is some member of R' such that b is best identified as being the value of some function f at a , w cannot normally be used to refer to any member of R that is not also a member of R' .

Finally, consider a third sort of case. A predication like "is tasty" can be both rationally and normally made both of chicken meat and chicken liver meat. So a sentence like 7 should be normally ambiguous in the null

context by conditions i and ii:

7. Chicken is tasty.

But 7 has only one normal interpretation, in which chicken refers to meat. The reason is that the function from chickens to chicken meat is necessarily better than the function to chicken liver meat, since the second is necessarily a non-trivial composite that includes the first as one of its elements (whether we express it as "liver of meat of" or "meat of liver of"). A third condition, then, would be necessary to stipulate that given two possible referents b and c in the natural range of a word w that designates a, such that there is a better function from a to b than from a to c, we can only interpret w in the null context as referring to b. (I will not bother to formulate this further, since it follows immediately from the conditions of reference as we gave them in Section 2.1 above.)

These conditions conspire to limit the number of normal interpretations that a given word-use can have in the null context. Where we have two equally good functions to different objects in the range, neither of which can be ruled out by any of the above criteria, then both interpretations are available in the null context; accordingly, we judge that a sentence like Water buffalo is tough is "ambiguous." (Though the case is rather more complicated than that. Note that the analogous "plankton is tough" has only a single interpretation. We cannot interpret water

buffalo here as referring to a single kind of substance, since our normal theories of water buffalos, unlike plankton, do not allow that it is made of a single substance which can be differentiated with respect to texture from other things. By contrast, "Water buffalo has 32 chromosomes per cell" allows a single interpretation, since the difference between meat and hides is not relevant under the suggested interest.)

I noted above, in discussing the unavailability of the "stock" interpretation for "He parked his Ford in the garage," that it would be fairly easy to jog an informant into supplying the desired interpretation. All we would have to do, in fact, would be to change the linguistic context a bit, say to:

8. He was worried that his stock certificates would be destroyed if his house caught fire, so he parked his Ford in the garage.

So that considerations of textual coherence could be brought to bear in fixing the normal range of Ford. More easily, of course, we could change the example to:

9. He sold his Ford at 71-3/4.

and the desired interpretation would be judged entirely normal in the null context. Similarly with 7, which we could amend to 10:

10. I like some kinds of liver; chicken is tasty.

All of this is familiar, and none of it is troubling; we have been arguing all along that a judgment that a use is "normal" in the null context is made on the basis of the

particular linguistic environment in which the use is presented, as well as on the system of normal beliefs. Normality is a property of tokens of word-uses; depending as we specify the range, we can get speakers to judge as normal an indeterminately large number of uses of a given word:

11. The blood plasma of animals is thin; chicken is the most viscous.

But very little follows from this, except that the "null context" is not a theoretically interesting construct after all. It doesn't matter whether a speaker judges a use solely against the contextual background established by its linguistic setting, or against some background specified extralinguistically, so long as, given the background, he believes that the use would be licensed by the belief-system that a rational member of the community would bring to bear in interpreting it. This is the important factor that distinguishes normal uses from local ones, which are licensed by systems of normal beliefs that are available only to a sub-group of the community; or metaphors, which are not licensed by any system of normal belief at all.

4.1.4 Local Uses

In Chapter Three, we argued that one could not idealize the "meanings" of words, even in langue, without suppressing the importance of the social processes wherein meaning was created. Since then, however, we have been assuming that idealization in order to deal with other matters.

It is true that we have shifted the emphasis somewhat, from the collective awareness that determines the pattern of usage of langue, to its instantiation in individual speakers--our "system of normal beliefs"--which can only determine parole. But there is no less distortion in talking about a speaker's perception of his community as a homogeneous object; we saw in Chapter Three how speakers could create belief-systems to suit the communicative requirements of a particular interaction. In this section, we will slowly relinquish the idealization of the system of normal beliefs, in order to set up the discussion of convention that follows.

When a speaker judges that a word-use is normal, we will want to ask, "normal for whom?" Over what community, real or imagined, does he construct the belief-system that licenses the use? The answer, as always, is "it depends." If the speaker is presented with a forced choice--every example must be labeled "in" or "out"--then his responses will be notoriously subject to manipulation, since the context of presentation (there are no true null contexts) will invariably lead him to assume something about the experimenter's interests and intentions, and the speaker will adjust his system of normal beliefs accordingly. But it doesn't really matter whether a speaker is judging that a use would be normal for anyone in the English-speaking community, or just for Americans, or Upper West Side Basque-

Americans, because it is not our object to isolate one group of users as "normal English." Even replicable, universal agreement that a use is normal, after all, doesn't mean that it is justified on linguistic grounds alone.

Such judgments are more interesting when they can be scaled (though here too the temptation to think of "normal" as picking out a range on a single axis may get us in trouble). The relative degree to which a use is judged normal, however we coax out the appropriate responses, will reflect the degree to which the speaker feels certain that it is justified by beliefs that must form a part of the system of beliefs that any member of his community would rationally bring to the neutral context.

A speaker can be most certain that all other speakers would best behave as if they believed that p when he believes that there could be no rational alternative to p. This does not mean that he must believe that all other speakers do behave as if they believed that p. For one thing, no belief would qualify under this requirement--not even the beliefs that license the use of the identity, probably, if we included the youngest speakers. It suffices that the speaker should be able to say, "well, if he isn't talking as if he believed that p, then who knows what he has in mind?" (Suppose someone says, "It took you longer to paint the garage than it took God to make the world." Either he means "longer than six days," or we could have no

way of telling what he means.)

A speaker can be most confident in his assignment of a proposition to his system of normal beliefs if he believes that no rational speaker, whatever his experience, could rationally believe other than *p*. (The "general" functions that we described above, such as the type/token function, the use/mention function, and so forth, are normal in virtue of this sort of assumption.) He can assign a belief to the system of normal beliefs with slightly less certainty if he believes that the experience of the members of the community is sufficiently uniform that no rational member, given that experience, could fail to believe that *p*. (The belief that licenses the use of chicken to refer to a kind of meat is like that.) Other propositions will be such that some rational members of the community may not accept them, but all must behave as if they believed them, and so on down the line.

At a certain point, however, a speaker must judge that others might incorporate either of two mutually contradictory propositions *p* and *q* into their systems of normal belief. That is, he must arrive at an apprehension of what Parsons, in his 1960 critique of Durkheim, called "differentiated norms." The existence of such discrepancies does not automatically lead to communicative problems, nor are the uses licensed only by (perceived) differentiated norms necessarily counted less than normal. Suppose, for example,

that there is a general awareness (etc.) that both p and q are available to the community at large, and suppose that one group of speakers--say, men--believes that it would be most rational for everyone to behave as if they believed that p, while another group--say, women--believes that it would be most rational for everyone to behave as if they believed that q. And finally, suppose that there is a general awareness (etc.) of this discrepancy. It is true that there would then be two systems of belief that could be brought to bear in interpreting utterances produced in mixed company, and that we wouldn't always know for certain which system was being invoked, since members of either sex might accommodate themselves to the other's belief-system in order to facilitate communication. In practice, however, such discrepancies can usually be resolved, according, for example, as participants are taking a conciliatory or mulish line. The situation is just the analogue of one in which speakers of different dialects accommodate to one another's usage; we can generally tell in context what an Englishman talking to an American intends to refer to by a given use of "public school."

A speaker may judge that a use is normal, then, even if he believes that the belief-system that licenses it would be considered normal only by a portion of the speech-community. Where members of a set of mutually contradictory propositions {p, q...} vary in accessibility or attributability according to the experience of speakers, and where

the relevant differences in experience are roughly predictable on the basis of perceived differences in social role or place, then speakers may be able to approximate to one another's belief-systems, as we showed in our discussion of words like jazz in Chapter Three.

Earlier we said that a speaker would assign a proposition p to his system of normal beliefs when he believed that it would be rational if all members of his community behaved as if they believed that p, absent any specification to the contrary. We can now amend this, saying that a speaker will count p a normal belief if he believes that there are some circumstances in which it would be rational for all speakers to behave as if they believed that p, again absent specification. That is, a belief is normal if a speaker believes that others can (normally) tell when it constitutes part of the unspecified background against which an utterance is to be interpreted.

Given a belief p that a speaker associates with a certain group G, it becomes progressively less likely that a speaker will call p normal (or better "judge that uses licensed by p are normal") as he is less certain that p is generally familiar, or that it is associable with G, or that members of G are themselves discriminable under ordinary discourse conditions. (This doesn't mean that the speaker himself doesn't ever behave as if he believed that p. His decisions about normality still depend on what he

thinks it is rational for all others to do.) So there will be word-uses which he believes to be rational for some members of the community, yet which he does not judge normal.

Linguists concerned with heterogeneity have tended to focus on certain kinds of social groups--as defined by sex, geography, social class, and race and ethnicity, for example--as being of greatest linguistic interest. This is certainly true so far as phonological and syntactic features are concerned, but variation in lexical use is more easily addressed as it reflects other sorts of community subdivisions. This is not to say that groups like social classes do not vary significantly in the beliefs they bring to the null context, or that these differences do not affect the way they use language. But these differences are invariably tied to differences in evaluative and prescriptive norms, and hence reflect, not simply that we think it rational to behave as if such-and-such, but that we think it honorable to do so. It follows that a speaker's judgments about a word-use licensed only by the beliefs associated with a particular social class will be much affected by prescriptive norms. (Explicit linguistic prescription is only an accomplice; we can judge that a use is vulgar without having heard it declared so, if it is licensed only by vulgar beliefs.)

We can't pursue here the topic of prescriptive norms, though I think it leads in an interesting direction.² (At

the end of 4.3, we will touch on the phenomena called "use-conditions," such as "obscene" and "written standard," which have up to now been regarded, like polysemy, as being only of minor importance to theories of language.) For our purposes, it will be easier to talk about beliefs associated with groups whose norms are not subject to prescriptive evaluations, such as what sociologists have called "collectivities," which we can take here, broadly, to include all groups whose members share common specialized goals or interests; e.g. classes of linguists, teenagers, sailing buffs, lawyers, or intellectuals. These divisions are obviously irrelevant to the interests of the dialectologist interested in phonological or syntactic variation. But they provide a much more convenient model for talking about variation in word-use.

Take a simple example. Suppose one sailor says to another,

1. His (wind) shadow was hurtful.

This use of shadow to refer to the area to the leeward side of a boat is rational only against a system of beliefs that is particular to a certain collectivity; it is licensed only when members of that collectivity are talking to one another. And they know this as well as anyone; they would use another expression in talking to a non-sailor, or to someone about whose interests they had no information at all. At the same time, a non-sailor may be able to apprehend that

1 would be a rational use, given the system of beliefs particular to that collectivity, even if he could not identify the referent of shadow, or produce the use spontaneously. The sailor and the non-sailor, then, might agree as to the relevant properties of the use in 1: that the entire community could perceive it as rational against a background of beliefs that is available only to some of its members.

Let's say that the "local uses" of a word are those that a speaker believes are generally perceived as rational against a system of beliefs that is available only to a subsection of the community. (It doesn't matter whether the speaker himself has access to those beliefs.) Now we could go on sub-categorizing such uses. 2 and 3, for example, are licensed by beliefs that can be assigned only to persons who are acting in certain roles within a collectivity:

2. The bartender said that the steak special walked out on his check.
3. The Times asked if he could have an exclusive interview.

Obviously, a number of considerations determined how a given use will be evaluated--what kind of roles and collectivities are involved, the degree to which the interests of the sub-group are roughly known to the community at large, and so forth. But none of these considerations is particularly interesting to us, nor do we care how many speakers say that a particular use is normal (or "normal for us/them.") There are simply too many kinds of heterogeneities

that a given speaker might be taking into account in making such judgments.

Uses like those in 1-3 are most interesting for the way in which they contrast with uses like those in 4-6:

4. He said he'd give it to me at less than the book price.
5. The team is controlled by a holding company.
6. I like Dick Powerll better as a juvenile.

These uses are not different from those in 1-3 with respect to the kind of justification that they are accorded. In either case, it would be generally agreed that the use is entirely rational against a background of local beliefs. Whether or not we know which book it is that determines the book price of a car, we assume that there is such a book, and that speakers who are familiar with car merchandising would generate "book price" spontaneously. (Note that it also doesn't matter whether we are correct in such assumptions, though it is obviously preferable to stick to examples in which we are. The fact that there didn't actually happen to be a group of speakers for whom "book price" was entirely rational might make no more difference to the way it was used than the actual non-existence of a winged horse would make to the way that a Greek used Pegasus.)

There is an obvious difference between 1-3 and 4-6 however. In general, only speakers who have access to the beliefs that license wind shadow use the expression, but a much greater proportion of speakers use book price, even in

talking to others whom they have no reason to believe would be party to the beliefs that license it. And this leads us to ask whether the use of "book price" to refer to (let us say) "standard market price" is a linguistic convention for speakers not in the car business. The answer would seem to depend on how we define convention. On the one hand, such speakers could not generate the use spontaneously, so we might want to say that their usage follows from their knowledge of the rules of language. On the other hand, they believe that it is entirely rational that the word book should be used in referring to standard car prices, perhaps even that the usage could not rationally be otherwise. It is not terribly important, of course, whether we define "convention" to include this case or not, but it is important that we should be able to characterize it in such a way as to distinguish it from other kinds of regularities that have also been laid to "convention," because the differences will have important consequences to how the expression can be used.

4.2 Conventional and Conventionalized Uses

4.2.1 Convention and Predictability

In 4.1, we drew an analogy between the processes whereby the speakers decide that given uses of words are normal, and the processes that participants must go through in order to arrive at the best strategy for solving coordination problems in general. We noted in passing that

coordination problems could be solved in two ways: participants could act on the basis of their beliefs about one another's mutual expectations; i.e., on the basis of what they take to be the system of normal beliefs. Or participants could simply agree to behave in a certain arbitrarily chosen way whenever such and such a situation arose; i.e., they could establish a social convention. In 4.1 we talked only about the first sort of use, because we were mainly interested in showing that judgments about word-uses were not based solely on knowledge of linguistic rules. But it is obvious that each such judgment must always be based at least in part on beliefs about what is conventional in a given community. It is only by convention that a given form designates anything at all; given that convention, we can rationally use it to refer to other kinds of things.

Accordingly, we will want a way of defining convention which is consistent with what we have up to now had to say about meaning and designation, and which will allow us to distinguish intermediate cases of "partial motivation." The problem of describing convention is extensively discussed by Lewis 1969, 1976; we may take his definition of convention (1969, p. 78) as a starting point:³

A regularity R in the behavior of members of a population P when they are agents in a recurring situation S is a convention if and only if it is true that, and it is common knowledge in P that, in almost every instance of S among members of P,

- (1) almost everyone conforms to R;

- (2) almost everyone expects almost everyone else to conform to R;
- (3) almost everyone has approximately the same preferences regarding all possible combinations of actions;
- (4) almost everyone prefers that any one more conform to R, on condition that almost everyone conform to R;
- (5) almost everyone would prefer that any one more conform to R', on condition that almost everyone conform to R'.

where R' is some possible regularity in the behavior of members of P such that almost no one in almost any instance of S among members of P could conform both to R' and to R.

(Parts of this definition are unclear out of context; Lewis uses "any one more" instead of "everyone" in 4 and 5, for example, in response to potential objections that needn't concern us here. 3 is intended to indicate that S presents a coordination problem.) For our purposes, this rough paraphrase of Lewis's definition will do; a regularity is conventional when everyone conforms to it in the interest of consistency, rather than to some other regularity which would work as well if everyone conformed to it. Thus, it is in the interest of consistency that we all drive on the same side of the road; and it is by convention that in America, we drive on the right, for consistency would be equally well served if we should all drive on the left.

Lewis offers this definition against the assumption that meaning determines extension; i.e., that the meanings of all statements are determined by linguistic conventions alone;

moreover, he is not much concerned over the problem of heterogeneity within the community (though his "almost everyone" is intended to allow some deviation from regularity by a few members of the community, and he gives as well a quantitative definition of convention, whereby a regularity may be conventional according to the degree to which the entire community conforms to it). Accordingly, we will have to modify this definition in two ways. First, there are regularities which would satisfy Lewis's definition which we would have some hesitation in calling conventional. Take the use of chicken to refer to a kind of meat. Everyone conforms to this use, and expects that everyone else will conform to it. And everyone prefers that chicken meat should be referred to as "chicken" on condition that everyone else refer to it so. Moreover, we could refer to chicken meat by some other term--say, "pullet," just as we call steer-meat "beef," and pig meat "pork," (see Scott 1819) in which case everyone would prefer that everyone else refer to chicken meat in this way.

How shall we distinguish this use, then, from the use of chicken to refer to birds? We would like to say that a regularity R is conventional in a situation S just if people would not conform to it if it were not conventional to do so, or equivalently, if there is no other convention R" for behavior in another situation S" (or set of such conventions) such that knowledge of R" would allow speakers to

apprehend each other's expectations in S, and arrive at a coordination equilibrium accordingly. But in either case, our definition would become circular. To avoid the circularity, we must add to our definition a condition 6, as follows:

(R is conventional among members of P when they are agents in S if and only if . . .)

- (6) Almost everyone expects that, if there were no general conformity to R in S, then there is no other situation S' such that most everyone would conform to R in S'.

In other words, if individuals conform to the same regularity in behavior in several situations, always in the interest of consistency, then we will assume that that behavior is conventional only in those situations for which we could not predict conformity to R on the basis of conformity to R in some other situation.

This is just a paraphrase, to be sure, of what we said about meaning in Chapter Three: the "meaning" of a word is that of its uses that most efficiently predicts the others. But the same pattern of "polysemy" affects non-linguistic regularities as affects linguistic ones; indeed, it would have to, since the processes that generate polysemy are non-linguistic to begin with. Take the behavioral regularity we call handshaking, to which we conform to greet, take leave, congratulate, declare truces, agree to conform, and so forth. It is unlikely that all of these situations involve separate conventions of handshaking; rather, we can

assume that the regularity has a single social "meaning" (but we do not really need the shudder quotes) from which all other uses are derived (though some may be "idiomatized"; there is an element of convention in the conformity to handshaking before prize fights, for example). Moreover, it is quite likely that there are non-linguistic regularities of this sort whose meaning is indeterminate--though I do not know how to tell if handshaking is one of them. Certainly the meaning of conventions may change, as one or another situation of conformity becomes prior.⁴

Not unexpectedly, we run into the same difficulties in trying to accommodate Lewis's definition of convention to a heterogeneous community that we did in trying to get from langue to parole. Obviously a word-use may be conventional for some members of a community, but not for others. The same may be true of non-linguistic regularities; one man may believe that it is entirely a matter of convention that one is married in a morning-coat; another, who wears morning-coats routinely to receptions and inaugurations, might be able to predict that a morning-coat would be worn at weddings. So we will have to amend Lewis's definition to talk about what status an individual accords to a given regularity in a given recurrent situation. As a first pass, we could embed all of the definition in a belief-context, saying "A speaker believes that R is conventional in S among members of P if and only if he believes that"

But in a heterogeneous community, the definition of S must often vary from member to member; the category of recurrent situations that we call "greetings," for example, is no better defined than the category of music we call "jazz." And speakers are aware of such discrepancies, and adjust their theories of the situation to the requirements of a particular context. (So that the determination that a handshake is appropriate in a given context can be made only by reference to the participants' takes on one another. And the meaning of a particular handshake, offered, say, as token of a willingness to cooperate in some collective venture, may be determinable only after litigation.)

It is important that we be able to generalize the notion of convention in this way, but it is more important that we be able to make it more explicit for word-uses in particular. Our revised definition of convention, relativized to the individual speaker, requires still more elaboration before we can deal with the problems that we set out at the beginning of section 4.2. So far we have said that a speaker will judge that a regularity R is conventional in a situation S when he believes in addition to conditions 1-5, that almost everyone expects that general conformity to R in a recurrent situation S is a prerequisite to conformity to R in any other situations. This "almost everyone" will be problematic, however, when we want to talk about perceived heterogeneities; we might better recast this proviso,

using the notion of "normal beliefs" that we introduced in 4.1.⁵ Let us say that a speaker believes that a regularity is conventional in S if he believes that it is normally believed that conformity to R in S is a prerequisite to conformity to R in any other situation.

We can now contrast this with a definition of "free uses"; we will say that a speaker will judge that a regularity is free in S if he believes that it is normally believed that general conformity to R in S is predictable on the basis of general conformity to R in some other situation S' (and believes as well that the other conditions 1-5 in the definition hold). This raises the entire issue of "predictability," which will become extremely important for our subsequent discussion of "partial motivation." We say that kick the bucket is an idiom, for example, because its meaning is not predictable given only knowledge of the meanings of its constituents. But on what basis do we determine that a given use of a given expression is predictable at all? And for whom does a use have to be predictable before we grant that it is normally held to be so?

The questions are obviously related. The strongest criterion for predictability that we could offer would equate the predictability of a use with absolute regularity and productivity of the processes that generate it. Then we would say that a use of w to refer to a in a situation S is predictable for a given speaker only if he would always produce

w in S, without any precedents to guide him. But this is surely too strong. I do not know with what regularity we could count on speakers to generate spontaneously the use of newspaper to refer to a publisher, or of radio to refer to a collective commercial activity. Let us assume, however, that many speakers would not produce these uses, or that they would do so only with prompting, or that they would not do so with the regularity that is in fact collectively observed. But this surely doesn't entail that the use is then not predictable for them. Speakers who would not produce the relevant use of newspaper might nonetheless judge that the use was entirely rational if they were presented with it, again without precedents, and might accordingly adopt it thereafter. (The "ideal speaker-hearer," moreover, could not judge that a use was predictable if he did not believe that other speakers would not judge it as such.)

We can draw a useful analogy here to strategies in other games. Let R be the regularity that we observe in the behavior of baseball players when they are on first with two outs and a three-and-two count on the batter: they run on the pitch. Anyone familiar with baseball would allow that R is entirely predictable, given a knowledge of the rules of the game. And one could observe a huge number of baseball games without ever observing a departure from R; it is likely that violations of actual

rules would be more frequently encountered. Yet it does not follow from this that all players generate R spontaneously: a youngster who is learning the game might fail to go on the pitch even after he has mastered the rules that optimize that strategy. On being corrected, however, he will eventually come to see that the strategy is justified, and will conform to it thereafter. And he will doubtless judge as others do, that R is predictable.

If the requirement that a speaker spontaneously produce a predictable use is too strong, however, the requirement that he merely be able to understand it is too weak. There are a number of expressions that we would call "idiomatic," for example, whose meanings seem to be recoverable in most contexts without prior knowledge of that particular expression. Consider phrasal idioms like give one's word, by chance, or rumor has it that . . ., or the use of round-trip to refer to a kind of travel ticket, or of close to refer to a kind of friendship. It is reasonable to assume that a French speaker knowledgeable in English would have little trouble in understanding these uses, even if he had never heard of them, just as an English speaker knowledgeable in French will have little trouble in arriving at the referent when he hears that a fare is "cent francs retour," or when someone talks about "une amitié profonde." By analogy, we would not want to say that it is entirely predictable that boxers should touch gloves before a fight,

yet someone from a culture in which handshakes function as they do in ours could easily divine the motivation for the practice, even if he had never seen a boxing match before (just as we can determine why Sumo wrestlers bow.)

When we say that a usage is predictable, then, we entail less than that someone might produce it in some instance of that situation, and be understood. But this much follows, in fact from the way in which we have already defined free uses, as those uses that a speaker judges would be normally believed to be predictable. We can take this as equivalent to saying that a speaker will judge that a use is free if it is licensed by the conventions governing the use of the same word in some other situation, together with those assumptions that constitute part of the system of normal beliefs, as we have defined it in 4.1. Then we can distinguish among degrees of conventionalization, according as the beliefs that license a use to which most speakers conform (i.e., which satisfies conditions 1-5 of Lewis's definition) depart from the system of normal beliefs.

We have already spoken at length about free uses like those of chicken to refer to meat, and of newspaper to refer to companies. The beliefs that license these uses are accessible to most speakers, and would be generally attributed to the system of normal beliefs. Moreover, the conventions governing the use of these words in other

situations (i.e., the use of chicken to refer to a kind of bird) are also recognized to form a part of the belief-system that most rational speakers would bring to bear in interpreting utterances. These uses will be free, accordingly, for any speaker who is aware of the distribution of these beliefs--in practice, for most present-day speakers of English.

Uses may depart from free in one of two ways. We have already talked in 4.1 about local uses, such as the use of shadow to refer to the leeward side of a vessel, or of sandwich to refer to a restaurant patron. These uses are not licensed by normal beliefs, but neither are they conformed to by the majority of speakers. Since they are not regularities in the behavior of members of P, there is no question of their being conventionalized at all. (In the case of this use of shadow, in fact, there is no recurrent situation at all such that most members of the population even have the opportunity to behave regularly in it, so the possibility of convention is ruled out on prior grounds.) We may choose to say that these uses are free relative to a given speaker's perception of the norms of the collectivity whose beliefs allow them, or establish a separate category of local uses; the difference is purely terminological.

Regular uses that are not free may be purely conventional, of course, or may assume an intermediate status. Take the uses of book to refer to a kind of car price, or of

juvenile to refer both to a kind of dramatic role and to the actors who are suitable to play it, which we mentioned at the end of 4.1. These uses are only predictable (given the conventions governing the use of book and juvenile to refer to other things) against a background of beliefs about automobile merchandising and play-making that is not generally available, and so which would not be assigned to the system of normal beliefs by most speakers, whether or not they are party to the specialized information. At the same time, many more speakers use the words thus than are members of the collectivities for which the uses are entirely rational, and their beliefs about the use could be assigned to the system of normal beliefs. It is reasonable to assume that it is normally believed that:

- i. The beliefs that license these uses of book and juvenile must include an awareness of the convention whereby book is used to refer to publications, or juvenile to a stage of life.

In other words, it will be generally agreed from which other use of book and juvenile this one is derived. We could say as much, of course, for uses that are far more conventionalized than these; it would be generally agreed that the verb plant, as in "They planted the evidence," is derived from plant "planta" rather than plant "factory," or that the kick of kick the bucket is the same as "strike with the leg," rather than "renounce." Still, agreement over i does distinguish these uses from examples of pure homonymy (bear vs. bare), or from cases of polysemy in which it would be

hard to say which psenses are most closely related, as with the run of "give him a run for his money," the dope of "this is the inside dope," or the stand of "a two-months stand," which might be linked to any of several other uses of the words.

Beyond i, we may assume normal agreement over ii:

- ii. These uses of book and juvenile are most likely licensed by the beliefs "Standard car prices are all listed in a certain book," and "There is a stock part in many plays that is filled by a younger actor or actress."

Thus, there will be general agreement over which beliefs are likely to have motivated the use in the first place. Note that ii does not entail that there is any normal acceptance of these beliefs, or even that anyone in the community might now attribute them to any system of normal or local beliefs. So again, there are uses that are much more conventionalized than those on which we could assume the same kind of general agreement. For example, it is likely that there would be general agreement that the use of book in make book was licensed by the common assumption that bets were entered in books, or that the use of book in "they booked the suspect" was licensed by the common assumption that arrests were entered in a book. But ii serves to distinguish these uses of book and juvenile from another class of uses, which are not sorted out by i. We may be able to link the noun hit "success" with the verb hit "strike," or the noun strike, as in "teacher strike" with the verb strike "hit." But there

would not be general agreement as to which beliefs licensed these uses.

Further general agreement, on iii, serves to distinguish these uses of book and juvenile from still another class:

- iii. The beliefs that license these uses are in fact part of the background against which some members of the community usually act.

That is, the uses are licensed by beliefs that some speakers do in fact assume for conversational purposes, and hence there are presently speakers for whom the uses are entirely free. Note that we would not allow that iii is normally assumed for the use of book in "make book," or for the use of seat in "a seat on the stock exchange," or for the use of thesis in "master's thesis." But there is one more distinction to be made, this one rather fine.

Consider the use of serious and country to refer to kinds of music; of stereo to refer to a kind of appliance; of smoker to refer to a kind of party; or of nylon to refer to a kind of stocking. In each case, we may assume that there is a group of speakers against whose normal beliefs the uses are entirely rational. (That is, there are people who believe that country music is best identified as that music that has a rural origin, or that stockings are best identified as those garments that are made of nylon.) But these groups aren't accorded the same status within the community as those to whom the relevant beliefs about these

uses of book and juvenile are attributed. Accordingly, we can add iv to the set of assumptions that would normally be made about these uses:

- iv. The groups for whom these uses of book and juvenile are licensed are those who would normally be assigned authority with respect to the categories of things to which the words are used to refer.

We have had occasion to appeal to this same notion of "authority" before, in our discussion of meaning in Chapter Three. We said there, following up on Putnam's notion of the "division of linguistic labor," that the determination of the criteria that defined some categories--such as gold, prime cuts, Coupes de Villes, and doctors--was by common consent left to specialists, with other speakers agreeing to conform to the usage prescribed by authority. (With other words, like jazz and newspaper, no single group exists to whom authority is generally assigned, while with still others, like vanity and vegetable, authority is vested in the community at large; there are no experts.)⁶

Where we entrust the definition of categories to certain authorities, we entrust them as well with the task of determining how those categories will be best identified. And we copy their usage, not because it is conventional to do so, but because we believe that we would agree that their use was entirely rational if we were party to their specialized knowledge, and that other speakers would agree as well. An analogy from chess will help here. Most players who have

any familiarity with the literature at all accept the often-noted assumption that 1. P-K4, P-Q4 cannot lead to equality for Black. So Black avoids the move with a regularity that increases as the sophistication of the player. But for many players, including myself, acceptance of this assumption rests entirely on faith in the authorities who tell us that 1 . . . , P-Q4 cannot be played. We believe that the regularity is entirely rational, and hence predictable; but we cannot give its rationale.

In the course of showing that these uses of book and juvenile are not really conventional, we have set up the framework that will allow us to talk about degrees of conventionalization. We have said that a regular word-use is predictable for a given speaker if he believes that it is rational, given his apprehension of the normal conventions governing its use, and of normal beliefs. Now we can make this definition relative, saying that a word-use is predictable for a given speaker to the degree that he believes that normal beliefs provide its rationale. Then we can distinguish several stages of conventionalization (though not without some idealization) according to the kind of rationale that the speaker believes that normal belief would accord to a use. Given a speaker S, who is aware both of a regular use R of a word w to refer to a, and of general conformity to another regular use R' of w to refer to b, we will say that R is free for S if either I or II holds:

Stage I S believes that it is normally believed that, given R', R is licensed by normal beliefs. (newspaper, etc.)

Stage II S believes that it is normally believed that, given R', R is licensed by beliefs to which members of the community would normally conform. (book price, juvenile)

A use becomes progressively more conventionalized for S as we move from 3 to 6:

Stage III S believes that it is normally believed that given R', R is licensed by beliefs that may be accounted normal by some members of the community, but to which members of the community would not normally conform. (nylon, country)

Stage IV S believes that it is normally believed that, given R', the beliefs that license R are not conformed to by any members of the community, though these beliefs might still be normally accessible. (make book, etc.)

Stage V S believes that it is normally believed that, given R', there exist beliefs such that R would be licensed by them. (kick the bucket, storm the fort)

Stage VI S believes that it is normally believed that there is no regularity R' in the use of w such that R might be predicted. (bear, inside dope. This is "p'rre" convention, and is equivalent to our earlier definition 1.)

In a way, 1-6 give no more than a taxonomy of the kinds of processes that give rise to what have been called folk etymologies, in which speakers guess wrong about the original rationale for a use, and reanalyze, say, asparagus as sparrow grass. As such, these processes have been of interest only as the source of a few diachronic curiosities. But speakers guess right much more often than they guess

wrong. (Perhaps "folk etymology" is an unfortunate phrase; we might better have stuck with Bréal's "false perception," or simply have used "reanalysis" throughout.) And in any event, it is more interesting that speakers should be able to come to some kind of partial agreement about the motivation for a use, than that they should be able to apprehend the actual intention of its original users.

I have no further theoretical points to make; the rest of this section, as well as the largest part of 4.3, will be given over to a discussion of the ways in which we can apply this account of meaning and reference to some specific problems in syntax and semantics. (If I were to revise this thesis, I would have included 4.1 and 4.2.1 in Chapter Three.) In the following section, I'll be talking about the syntactic behavior of phrasal idioms; in the next, about some of the syntactic problems raised by the English do so construction, and about how we may interpret the speaker judgment that linguists indicate with the mark "?." In 4.2.4, I'll talk about how we can describe change in lexical meaning. In 4.3 I'll show how we can account for "metaphorical" word-uses without further apparatus than we already have available; at the end of that section, I'll be talking about some of the problems that arise when we try to attach usage-conditions like "slang" and "formal" to lexical entries. I am not under the delusion that there is much overall coherence to what follows, although it will occasionally

be the case that one topic leads naturally to the discussion of another. The generous reader will accordingly take the rest of this thesis as a series of exercises, and if he is smart, he will bother to look only at those sections that deal with topics in which he has some immediate interest.

4.2.2 The Syntax of Idioms

The degree to which a use has become conventionalized may affect its distribution in several ways. Let me use the term "idiom" to refer to any regular use of a word or expression that is generally analyzed as in Stages III-V above. This is not quite the way the word is normally used in the literature; most writers take an idiom to be any expression whose meaning is not predictable on the basis of its structure and the meanings of its parts, so that technically every lexical item is an idiom, as Hockett 1958 pointed out. Under our definition, however, only uses that are analyzed as in some measure licensed by some other use will qualify as idioms. Then we can distinguish lexical idioms, such as the use of plant in plant the evidence, or of nylon and country to refer to stockings and musical genres; and phrasal idioms, such as kick the bucket, pay lip service, and I should hope to kiss a pig.

Lexical idioms have interested linguists primarily because of the problems they pose for morphology. For example we mentioned above Aronoff's example of the verb stand, which has the same past tense form stood in all of

its uses, which seem otherwise to have no salient meaning-relation. But such patterns are most striking only when uses are reached the level of convention of stages V and VI above. It is interesting that 1 and 2 should have the same past tense, but it is unremarkable that 3 should have the same past tense as 4:

1. John stood the umbrella by the door.
2. John stood her obstinacy as long as he could.
3. John stood in the corner for several hours.
4. It stood for many years as the classic paper on the problem.

That is, we are not surprised to find a formal relation between uses where we can perceive a relation between their rationales, even if the latter has become conventionalized to some degree. Accordingly, the difference between stages I-VI does not correlate interestingly with any difference in purely formal regularities, since formal identity (i.e., the fact that there exists some other use R' of w such that R can be predicted on the basis of R') is presumed at all stages. Moreover, such regularities often arise as a result of paradigmatic pressures which cannot be dealt with in our terms. (Though we might note that it is reasonable to assume that no formal differentiation of uses can take place if they are related to R and R' in stages I and II.)

More interesting to us is the relation between conventionalization and the anaphoric processes that we talked about in Chapter Two, which can be used only to relate free

uses of words. But we cannot begin to discuss this problem here; because we would need a full account of the apparatus of reference--relative clauses, pronouns, determiners, and so forth--to do it justice. And here as well the intermediate stages of conventionalization are not interestingly differentiated. There is no apparent difference in acceptability, for example, between examples 5 and 6, though the relation between the controlling and controlled uses is quite different in the two cases:⁷

5. *John stood the umbrella by the door, and her obstinacy as long as he could.
6. *The book stood on the shelf for two weeks, and as an example of careful historical reconstruction for over thirty years.

The behavior of phrasal idioms raises a number of problems for syntax and semantics. From a purely semantic point of view, it has always seemed as if idioms like kick the bucket and shoot the breeze should be treated as single unanalyzable lexical items, which may be optionally inserted in place of their monomorphemic "synonyms" die and schmooze. At the same time, various morphological and syntactic regularities associated with these expressions seem to require that they sometimes be treated as analyzable strings. The shoot of shoot the breeze shows the same past tense shot as the non-idiomatic shoot of "shoot the sheriff." Similarly, as Binnick 1971 noted, a number of idiomatic expressions containing the lexical item come show parallel causative expressions with bring; thus we get the relations

shown in the a and b pairs below:

- 7a. John came to the party.
- b. Alice brought John to the party.
- 8a. John came to when he was given smelling salts.
- b. We brought John to with smelling salts.
- 9a. The topic came up in discussion.
- b. John brought up the topic in discussion.
- 10a. The party came off without a hitch.
- b. John brought the party off without a hitch.

This has been taken as evidence that the come in the idiomatic 8a, 9a, and 10a must be identified with the come of the non-idiomatic 7a. And most significantly, we have Weinreich's 1966 observation that some idioms must be regarded as structurally transparent to the operation of some syntactic rules; thus we get 11 and 12 alongside of 13 and 14:

- 11. The ice was broken by John's timely remark.
- 12. Sooner or later, the piper must be paid.
- 13. *The bucket was kicked by John at 11 last night.
- 14. *The breeze was shot by the boys for two hours.

Problems like these have led linguists to propose a number of formal devices which would predict the semantic and syntactic regularities associated with idioms (see, e.g. Weinreich 1966, Fraser 1970, Dong 1971, Newmeyer 1972, Katz 1973). None of these is entirely satisfactory, however, for reasons that have nothing to do with their relative elegance or descriptive adequacy. For example, consider the ways in which a grammar might be asked to account for the fact

that kick the bucket does not undergo passive, while pay the piper does. (For simplicity's sake, we will assume a passive transformation; nothing hangs in the balance.)

Let us grant the usual assumption that phrasal idioms must be lexically entered as substrings whose internal structure is fully specified, which may be inserted into any appropriate phrase-marker configuration or subconfiguration; i.e., that kick the bucket is entered with a single semantic representation, and is lexically specified as having the structure V + NP. Then we might follow the suggestion of Weinreich, marking the entire expression with a feature [-Passive]. Further economies could be realized by following Fraser's proposal, marking the idiom with a single feature that assigns it to a class of idioms, such as dance up a storm, which behave similarly with respect to the operation of syntactic rules. We may even be able to get by with Katz's suggestion that some of the constituents of some idioms be marked with the single syntactic feature [+Idiom], which blocks the application of a transformation to strings of which the marked constituent is a component if they otherwise satisfy its structural description. (Thus the presence of the [+Idiom] feature on the NP in kick the bucket will ensure that passive cannot apply to strings of which it constitutes the VP.)⁸ But however we state such exceptions, it is hard to see how speakers could learn them. How do we determine, for example, that the bucket of kick

the bucket is marked with the feature [+Idiom] (or whatever) while the piper of pay the piper is not? Not, surely, by observing that we have never heard "The bucket was kicked"--we do not have to have heard "the piper must be paid" in order to produce it. And not because there is any structural difference between the two expressions.

We must assume, then, that the syntactic distinction between kick the bucket and pay the piper is related in some way to the meanings of the expressions. In this connection, Newmeyer has made the ingenious suggestion that the syntactic behavior is predictable on the basis of meaning alone. He proposes that idioms like kick the bucket, sit on pins and needles, shoot the bull, and make the scene do not passivize because their non-idiomatic equivalents die, wait, talk and arrive are intransitive. Unfortunately, this proposal doesn't work, for all its prettiness. Give up the ghost, for example, does passivize; at least 16 is much more acceptable than 15:

15. *Once the bucket has been kicked, there's nothing medical science can do.

16. Once the ghost has been given up, there is nothing medical science can do.

Similarly, idioms like throw in the sponge and pop the question can be passivized, although their non-idiomatic equivalents resign and propose are intransitive. Or consider the idiom have a shot at, which doesn't passivize, as opposed to its synonym try:

17. Bowling was tried (*had a shot at) by John.
 This failure is surely due to the presence of have . . .
at in the surface form of the idiom, rather than to any
 consideration of meaning.

If the behavior of idioms can be predicted neither by
 form or meaning alone, then it must be due to the relation
 between the two. But we cannot address this relation if
 we say that phrasal idioms are simply given a single seman-
 tic representation, which is unrelated to the meanings of
 their parts. Rather, we have to talk about idioms as par-
 tially analyzable. It has been observed that there is often
 an obvious diachronic relation between the "literal" and
 idiomatic uses of expressions like these; thus Sadock 1974
 p. 98 writes, "It is not surprising that the idiom down in
the dumps means "depressed" and that we do not use in sev-
enth heaven or on cloud nine in this sense. . . . Yet from
 a synchronic point of view, there is no reason why these
 idioms should have exactly the sense they do, and not, say,
 the opposite senses." But it is one thing to say that there
 is no obvious reason why down in the dumps should have
 exactly the use that it does, and another to say that, given
 this use, it cannot be accorded a post-hoc motivation.
 After all, if the use is "not surprising," it is because of
 a perceived synchronic relation to the "literal meaning."
 (I assume that most speakers now analyze the dumps here as
 related to dumps "refuse heap," rather than to the much

older dumps "dazed or puzzled state," which is the actual source of the expression. CF. the OED citation from Swift's letters, 1714: "he tells me that he left you horridly in the dumps." The "refuse heap" use is first cited from 1871.) And there is no reason why this perceived relation should not affect the way in which the form is used.

Take the idiom kick the bucket, which is standardly glossed as "die." J. R. Ross has noted (viva voce potshot, after a CLS talk) that this expression is diachronically motivated at least to the extent that it contains an active verb; it would be surprising to find wait for the train used for the same purposes. But this observation has a synchronic relevance as well; kick the bucket is not used to refer to a protracted death:

18. John is slowly dying (going to meet his maker, giving up the ghost, ?kicking the bucket).

Similarly, contrast the following pairs:

19. John abruptly kicked the bucket (?shuffled off his mortal coil).
20. I am slowly becoming panicked (?hitting the panic button).
21. John proposed (?popped the question) by degrees, in a series of increasingly amorous letters.
22. He has been increasingly taking a stand (?nailing his colors to the mast) over the past few years.
23. They kill (?snuff) their prisoners by starving them.

We can assume then, that speakers must make some assumptions

about the rationale whereby kick the bucket is used to refer to dying; in particular, that kick is the same word that is used to refer to the act of striking with the foot, and that the best function from this takes us to an abrupt act, rather than to a gradual act, or to a simply change of state. (For this reason, we do not use kick the bucket to refer to the "death" of inanimates; we can say that it has been ten years since Life magazine "gave up the ghost," "passed on," or even "went to meet its maker," but not since it "kicked the bucket.") Moreover, we can assume that speakers generally attribute awareness of this rationale to the system of normal beliefs, which is why there is a large measure of agreement over the use of the expression. It does not matter whether the normal assumptions about the relation of "literal" and idiomatic uses are historically justified, (cf. the use of dumps above), so long as they are grossly uniform for the community.

In fact, it should be clear by now that we have reason to suspect the standard notion that we can even state the "literal meaning" of an idiom without reference to its particular use. Consider several phrasal idioms of the form hit NP. Intuitively, we would say that the hit of hit the panic button is equivalent to strike, as in "He hit the wall with his fist." But we would identify the hit of hit the deck, hit the sack with the "collide with violently" use, as in "He hit the guardrail going sixty-five." And the hit

of hit the highpoints, hit bottom, seems to be equivalent to the "make contact with, reach" use, as in "We hit the Cross-Bronx at rush hour." But we could not make these determinations if we could not agree as to what assumptions speakers are likely to make about the contribution that one or another use of hit was making to the use of the expression. In fact, we must also assume that they are capable of reckoning the cue-validity of functions from several uses of hit to its idiomatic referent, which entails that they must in every case have fairly uniform assumptions about the beliefs that license the idiomatic use. Or consider a wholly hypothetical example. We would identify the lay down of lay down the law with the "set before someone" use, as in "The tribal council laid down six conditions for membership." But suppose the same expression were idiomatically used to mean "die." Then given "Poor John laid down the law at six last night, after a long battle with cancer," we would probably assume that this lay down was derived from the "doff" use, as in "John laid down his burden." Here again we can have intuitions for the "literal meaning" only by making hypotheses about how the idiomatic use is compositionally derived.

Given that normal assumptions about the rationale that licenses a phrasal idiom can affect the way in which it is used (cf. kick the bucket/give up the ghost examples above) we should not be surprised to find that they affect

syntactic behavior as well. We have already seen that we cannot account for the transformational defects of idioms solely by introducing unmotivated syntactic or syntactic-semantic restrictions into the grammar; now we may have the say of doing without such conditions entirely. Suppose that we can let the grammar freely generate such strings as The bucket was kicked by John, and then explain the unavailability of the idiomatic interpretation on functional grounds. This would clearly be a desirable result from the syntactician's point of view.

We could only give such explanations on a case-by-case basis, to be sure. For every instance in which a given rule does not yield an acceptable output when applied to a structurally appropriate phrasal idiom, we will have to be able to say why it would be odd to apply that rule to that idiom, given both the conversational function of the rule, and the assumptions that speakers are likely to make about the way in which components of the idiom contribute to its interpretation. Obviously, then, we cannot hope to be able to give a systematic explanation for the syntactic behavior of idioms; the best we can do is to give a plausible account of a few cases, and let the argument follow inductively. (For this reason, such accounts are always vulnerable to "counterexamples"; the apparatus of explanation has to be cranked up anew for each case. But pragmatic explanations are not subject to the same canons of economy that

govern syntactic descriptions. It doesn't matter if we have to invoke twenty different principles to explain twenty different regularities; so long as each of the principles is independently motivated, there is no increase in cost.)

Let us restrict ourselves to looking at the application of passive to idiomatic transitive VP's, such as kick the bucket and pull up stakes. Admittedly, the function of passive is particularly hard to characterize (as opposed, say, to that of restrictive relative clauses). It is universally agreed that speakers use passives when they want to "focus" on the logical object of the verb (or "foreground" it, or make it the "topic," "theme," or "old information") usually de-focussing (etc.) the logical subject in the process. But the notion of focus is hardly clear, and critics are justified in pointing out that there is little agreement over what regularities ought to follow once we have focussed on something. Sometimes, it is true, the reason for the oddness of a particular example may be intuitively so apparent that a more explicit account is not required. Thus, we do not have to inquire deeply into the nature of focus to explain why it is strange to say "A book was read by me." But other regularities associated with passive sentences are not so easily explained. There is no obvious connection, for example, between focus and the scope of quantifiers, which is why even linguists who readily admit the rhetorical function of passive may insist that scope

relations are determined by independent semantic rules. And the burden of proof here clearly rests on the functionalists; we cannot be asked to accept simply as a matter of faith that focus has a certain idiosyncratic effect of scope, or that passive has some other less obvious function that we could appeal to to explain the regularity.

Keeping this in mind, let us begin with "focus" and grind out what we can. Consider first the class of idioms given in 24, all of which passivize freely:

- 24a. Keep tabs on, lay down the law, make headway,
break the ice, pay heed.
- b. Pop the question, spout nonsense, pick a fight,
shave points, spin a yarn.

It could be argued (see, e.g., Newmeyer 1972) that these are not phrasal idioms at all, since in each case, only one of the components is used abnormally. Thus keep tabs on corresponds to the non-idiomatic keep a close watch on; make headway to make progress; lay down the law to lay down the rules, and pay heed to pay attention. In the b examples, only the verb seems to be idiomatically used; thus pop the question is (roughly) equivalent to ask the question, and spout nonsense to speak nonsense. (I have even heard it suggested that the headway of make headway should be treated as simply a positional variant of progress, which can co-occur only with make; and heed as a variant of attention, and so on. This will not do, because the expressions are never quite synonymous. To pay heed is to comply, while to pay

attention is only to listen; we can keep tabs on someone's performance without keeping a close watch on it. Even headway and progress are different; the former is used to refer only to the initial stages of advancement towards a predetermined goal, so it would be odd to say, for example, "We have made a lot of headway in chemistry since 1965." And the use of shave in shave points has not even a near synonym; we would certainly not want to say that it was a positional variant of refrain from making, in the interest of keeping the winning margin smaller than the spread. It is not possible, then, to handle these cases simply with morphological rules.)

Formally, of course, there is little difference between saying that headway is a lexical idiom, and saying that make headway is a phrasal idiom; that is, between saying that headway has a certain idiomatic interpretation just when it co-occurs with make, and saying that make headway has a certain idiomatic interpretation assigned to it as a collocation. Moreover, it is rare that the interpretation of the non-idiomatic constituent is itself entirely unconventionalized. There is only one question that can be "popped," and the keep of keep tabs is obviously a different use from the keep of "keep the ticket stubs," or "keep chickens." What is interesting about these idioms is that they are particularly easy to decompose: we can discern just how each of the components contributes to the

interpretation of the whole, and we assume that other speakers will break it down as we do.

Let us say that verb phrases "refer" to states and activities, and that transitive verb phrases normally refer to states and activities that are best identified as "open relations" of the form Rxb , where "R" stands for the relation referred to by the verb, "x" is a variable for the referent of the sentence subject, and "b" stands for the referent of the object NP.¹⁰ Then we will say that an idiomatic transitive VP is decomposable just in case it is used to refer to a state or activity such that it would be normally believed that that activity could be identified as an open relation Rxb , such that the object NP of the idiom refers to b, and the verb to R. Take pop the question. We would normally identify its referent as a relation between a person and a proposal; we can easily see how the verb pop refers to that relation, and the NP the question to the proposal. The beliefs that license these uses are not entirely normal, to be sure. The question can normally be used to refer to a marriage proposal only in contexts in which no other question is salient, but the idiom is used in other contexts. And pop is normally used transitively to mean "cause to emerge suddenly"--the use that is relevant here--only with an explicit adverbial, like out or up.

It is not surprising that we should be able to use pop the question in passives, then, whenever we want to focus

on the referent of its NP object--i.e., the proposal. And the same holds for the object NP's of the other idioms in 24. The NP in break the ice is clearly used to refer to a mood of reserve; the law of lay down the law to the rules governing conduct in certain situations. We can passivize it when it is appropriate to focus on that mood or those rules.

It is not necessary that one constituent of an idiom should be used normally, or almost normally, for an idiom to be decomposable, though the determination that the idiom is decomposable is most clear in those cases. Take expressions like spill the beans, paint a pretty picture, which are idiomatically used to refer to speech acts in which a certain kind of information is conveyed in a certain way. In each case, it is reasonable to assume that the verb refers to the mode of transmission, and that the object must refer to the material that is transmitted. (Though the rationale for using beans is obscure, and we arrive at its referent, presumably, by subtraction.)

Note also that it is not necessary that the referent should be "normally" referred to with a transitive construction in order that we should be able to break it down into components to which the constituents can refer. We normally refer to dying with die as if it were a change in the state of a single individual. But we can also identify death as a two-place relation, which holds between persons

and their spirits, or between souls and bodies. Then given give up the ghost, we will be able to determine that there is a function from ghosts to spirits, and from the act of giving up something to the act of relinquishing one's hold on the spirit. (Originally, of course, give up was "render," and ghost designated a class like that now designated by spirit. The use is obviously no longer licensed by normal beliefs about death, but there is normal awareness of the kinds of beliefs that speakers used to have, or are supposed to have had, about dying. Cf. the convention in cartoon drawings of showing a ghost rising from a dying man.) By the same token, we analyze shuffle off this (one's) mortal coil as referring to the act of discarding one's body. (Note that if we say, "John has given up the ghost," John must refer to a corporeal person, while with "John has shuffled off his mortal coil," John refers to a personality. Accordingly, there is a certain oddness to 25, but not to 26:

25. ?John has given up the ghost; he's with the angels now.

26. John has shuffled off his mortal coil; he's with the angels now.

Thus the way in which we decompose the idiom will affect the way in which we interpret the referent of the subject NP as well.)

By contrast, we cannot suppose that speakers will be able to agree to analyze dying into a relation and

entity such that there is a function to each of them from the components of kick the bucket. We have already seen that kick must be construed as referring to an abrupt act. But there is no normal way of analyzing death as an abrupt act performed on something else, such that we can perceive that there is any good function from buckets to it. (If the idiom were kick this vale, we might interpret kick on the basis of its "abandon, depart from" use; even kick this bucket might allow that interpretation, with bucket construed as referring to the world. But actual idioms are hard enough to deal with.) Or contrast take a powder with hit the road, both of which refer to kinds of departures. The latter is decomposable (though it is not easily passivized--see below); the road is construed as referring to the route taken in departure, and hit to the manner of reaching it. But of the entities that might be perceived as components of an act of departure--the commitment to the place departed from, (cf. pull up stakes), the permission to leave (cf. take leave), the goal, the route (cf. hit the road), or the marks of motion (cf. make tracks)-- there is none such that we have any normal rationale for referring to it with powder, nor is take restricted to referring to any one relation that could hold between any one of these entities and the subject. We might, in other circumstances, use take to refer to the relation between leaver and permission to leave, as in take one's leave, or to the route

"taken," or reflexively, to the leaver himself, as in take (oneself) off. With take a powder, however, there is no way of telling what relation take is being used to refer to.

We can assume then, that speakers will not passivize idiomatic VP's that are not decomposable. If there is no assurance that hearers will be able to identify the referent of the focussed NP, then it makes no sense to focus on it. Examples of apparently non-decomposable phrasal idioms are given in 27:

27. Kick the bucket, take a powder, shoot the breeze, give a damn, chew the fat, hit the spot, play hooky.

The fact that an idiom is decomposable, however, does not ensure that it will show up in passives. First, it may be that the referent of the object NP is not the sort of thing that a speaker might reasonably be expected to focus on. Consider the idioms make tracks and hit the panic button, which are odd in passive sentences:

28. ?The panic button was hit by John when he heard the news about Mary.
29. ?Tracks were made by the robbers as soon as they saw the patrol car.

But these idioms are decomposable. We have already seen that we could identify a leave-taking by referring to a relation between a leaver and the path he makes in departing. Similarly, we can identify someone's becoming panicked by reference to the act they perform in initiating certain

kinds of behavior. If these idioms are different from give up the ghost or pop the question, it is in the kinds of assumptions involved in identifying referents in this way.

We said above that an idiom was decomposable if its referents could be identified as a relation of the form "Rxb" such that, etc. By this criterion, there is no difference between idioms that refer to the act of dying as a kind of relation between a person and his spirit (give up the ghost), or between a soul and a body (shuffle off his mortal coil), or between a person and the world (pass away, with an implicit point of departure), or between a person and either the place he goes to after dying or the people he meets there (go to the happy hunting ground, join one's ancestors). But these means of identification are not all equally rational, given our normal beliefs about death. An idiom like go to meet one's maker really involves two different kinds of conventions: one whereby the word maker can be used to refer to God, and the other whereby we can refer to death metaphorically by invoking a relation between the subject and God. (Though each has a certain rationale.) Let us say that an idiom is normally decomposable when it is analyzed as involving only conventions whereby each of its constituents can be used to refer to the constituents of its referent in certain situations. Thus, the convention governing the use of pop the question could be stated as "The NP the question is used to refer to a marriage proposal

when the verb pop is used to refer to the act of uttering it." Examples of normally decomposable idioms are given in 30:

30. Pop the question, pass the buck, draw the line, foot the bill, lay a ghost, blaze a trail, pick holes in, steal the limelight, stir up mud, stretch a point, take note of, keep the pot boiling, find fault, break the ice, pull strings.

By contrast, consider an idiom like go to heaven. The convention governing its use does not have to mention word-use at all. (Note that we could substitute synonyms salve referentia; thus "He has proceeded to (gone on to, traveled to) the happy hunting ground (St. Peter's gate, the great quantifier in the sky)". Rather, there is a "conventional metaphor" (see 4.3) whereby we may refer either normally or non-normally to one relation in order to identify another. This same sort of convention enables us to use phrases like show someone to the door, or throw someone out on his ass (cf. "kick someone out on his butt," etc.).

It is not clear whether expressions like go to heaven, which involve only conventional metaphors of this second type, are idioms at all. But there is another class of expressions that involve both sorts of conventions; these include the idioms listed in 31:

- 31a. Line one's pockets, dip into one's purse, raise the roof, hit the ceiling, throw in the sponge, (go to) meet one's maker.
- b. Make tracks, ring a bell, hit (push) the panic button, give the gate to, run circles around, shake the dust from one's feet, take the fifth, carry the torch.

Let us say that these idioms are abnormally decomposable; we can identify the referents of their constituent terms, but it is only in virtue of our knowledge of conventional metaphors that we know what that relation is invoked to identify.

It is not hard now to explain the behavior of these two classes of idioms with passive. Since the object NP of abnormally decomposable idioms does not itself refer to some component of the idiomatic referent, but only to a component of the relation by means of which that referent is conventionally identified, it is irrational to produce that NP in a sentence position which is reserved for terms whose referents are "in focus" in the discourse. Thus, passives are odd with abnormally decomposable idioms for the same reasons that they are odd with idioms that are not decomposable at all; there need be nothing in the idiomatic referent to which the object NP refers. Only (but not all--see below) idioms that are normally decomposable can be passivized.

At this point, it could be argued that we have introduced too many distinctions. The criterion relevant for predicting the behavior of passive with idioms, it turns out, is whether the components of the idioms refer separately to components of the referent. Why should we distinguish, then, between idioms like make tracks and idioms like kick the bucket, since in either case this criterion is not satisfied? We can find a justification only by looking at other kinds of regularities, where the distinction between these two

classes is relevant. It has been noted (as by Sadock 1974) that idiomatic VP's may vary according to the degree to which the NP can be freely modified. Thus we can modify the trail of blaze the trail, or the strings of pull strings in any number of ways:

32. John blazed a pioneering (much-followed, important, valuable) trail in microscopy.
33. John had to pull some powerful (hidden, uniquely available) strings to get the job.

The object NP's of other idioms can appear only with a much more restricted group of modifiers. We can say "make deep tracks" or "ring a familiar (loud) bell" but not much else. And finally, the object NP's of idioms like kick the bucket and do the trick admit only the modifiers ol' and proverbial.

There is no mystery in the freedom with the NP's of normally decomposable idioms can be modified. Since they refer to components of the idiomatic referent, any modifier that occurs with them is simply construed as further identifying their actual referents. Thus with pull strings, we normally interpret strings as referring to connections, so that hidden strings is interpreted as "hidden connections" and so forth. It is equally clear why nothing other than ol' or proverbial can be used with kick the bucket; the only thing that anyone could know about the referent of bucket is that it exists in some unavailable proverbial scenario. But examples like make deep tracks and ring a loud bell are more interesting. Since the referents of tracks and bell are

not present in the idiomatic referent, the modifiers can't be construed as directly modifying any component of the referent. But the fact that such modifiers can be used with the idiom, and that they affect its interpretation, indicates that the decompositions of the idiom are relevant to its use.

We can assume that speakers will generally agree as to what beliefs would license the use of make tracks to departures, and in particular, that the expression would be most rationally used to refer to departures in which the departure was especially vigorous and abrupt; i.e., where make tracks could best refer metaphorically to the relevant action (see the following section). In just those circumstances in which make tracks could be used metaphorically to refer to a rapid departure, make deep tracks could be used to much the same purpose, to a more emphatic effect. Similarly, we can adduce the circumstances in which ring a bell could rationally be used to refer to the act of triggering a certain kind of mental event, and we know that the use of a loud bell would be licensed by the same beliefs to refer to an instance of this event that was particularly intense. In each case, then, we interpret the modified NP in the light of the rationale we accord to the convention whereby one relation can be used to refer metaphorically to another in the first place. We could give the same sort of explanation for uses like those in 34-36; where the

idiomatic sense is preserved under modification of the NP:

- 34. You will have to pay that implacable piper sooner or later.
- 35. He lined his already-bulging pockets with graft from the new construction project.
- 36. He has gone to meet his welcoming maker.

We can assume normal agreement that pay the piper involves a convention whereby piper is used to refer to an animate agent who exacts the consequences of folly, and a conventional metaphor whereby we invoke a relation between persons and the agents who exact such consequences to refer to the relation between persons and the consequences themselves. (Cf. pay the fiddler, give the devil his due, "You'll have to make your peace with God.") Moreover, we can assume normal agreement that if there were such an agent, and he were implacable, then the consequences of folly would be unavoidable. Again, assumptions about the way in which the idiom is decomposed enable us to modify its constituents.

As we noted, there is considerable disagreement about the acceptability of these idioms in the passive. We can assume that this is due in part to disagreements about how idioms would normally be decomposed, and about what beliefs are normal in the first place. But we should also note that there is an inherent grading effect that reflects the degree to which it is conventional that a certain metaphor can be used to refer to the idiomatic referent. For example, the beliefs that license give up the ghost are not entirely

normal, but we also assume that some speakers normally behave as if they were. The beliefs that license pay the piper are presumably not accounted normal by any speakers. But in either case, we also know that all speakers may occasionally act against these beliefs, in the service of some aesthetic or affective purpose. Thus we often talk as if we believed that there was someone who exacted the consequences of folly, or that a spirit departed from the body when someone died. By contrast, it is only in uttering idioms like make tracks that we ever behave as if we believed that departures were best identified as relations between the departed and the marks of their motion, or in uttering hit the panic button, as if we believed that the initiation of panic was best identified as a relation between a person and a mechanism that initiates panicky behavior. Accordingly, we can distinguish among the two classes of abnormally decomposable idioms given as 31a and b. The a examples are licensed by non-normal beliefs against which speakers may sometimes act, and so might be licensed in the passive under certain circumstances. (Examples like give up the ghost however, which are licensed by local beliefs, are better counted as normally decomposable.) At the least, we can roughly grade the acceptability of some relevant examples as follows:

37. The piper must be paid.

38. ?His pockets were lined with the graft from the new dam.

- 39. *The crap was cut by John.
- 40. *Tracks were made by all of the burglars when they heard the patrol car.
- 41. *Circles were run around Mary by John.

But these judgments are too much subject to confounds to be of any real value.

Let me touch briefly on one other factor that may contribute to the acceptability of passivized idiomatic VP's. Consider 42 and 43, which give examples of passivized idioms of the form hit NP and take NP:

- 42a. The jackpot was hit by a little old lady.
- b. Once the highpoints have been hit, you may as well go home.
- c. *Once the sack was hit, they all slept soundly.
- d. *The deck was hit by all of the recruits simultaneously.
- e. *When John is depressed, the sauce is hit hard.
- f. *The headlines were hit by the story the very next day.
- 43a. Wagers are taken on every conceivable event.
- b. No notice was taken of our departure.
- c. Turns were taken by all the participants.
- d. ?The wind was taken out of his sails by John's criticisms.
- e. ?The rap (fall) was taken by Wilmer.
- f. *Pot luck was taken by all of the dinner guests.
- g. *A back seat (in the proceedings) was taken by John.
- h. *His medicine was not taken gladly.

All of these examples, we can assume, involve idioms that are normally decomposable (the object NP's can be freely modified, for one thing). Where they are unacceptable in the passive, then, we cannot deal with them as we would with take a powder or hit the spot. Rather, the difference seems to correlate with normal intuitions for which use of

hit and take is involved in the composition of a particular idiom. For example, we would say that hit the jackpot and hit the highpoints involve the same use of hit as in hit the answer, hit the lucky number, hit all the reasons for not going. And each of these would passivize freely. On the other hand, the hit of hit the sack, hit the deck seems to be the same as the hit of The paper I dropped hit the floor, which is odd in the passive. And the hit of The story hit the headlines seems to be like We hit the Cross-Bronx at rush-hour, which also doesn't passivize.

Similarly with take. The use in take wagers seems to be the same "accept voluntarily" use as in take Canadian money; the take of take turns can be identified with the use in take a good shot at the turkey or take the job. In each case, the analogous normal use can be passivized. But the take of take the fall, take pot luck, take one's medicine, take a back seat seems to be more like the "accept involuntarily" use, as in Mary took his abuse because she was afraid to leave, or The tree took a terrible beating in the storm, which does not passivize easily.

Strictly speaking, we need to say little else to make the relevant point about idioms. We have already seen that the determination that one or another use of a form is involved in the derivation of an idiom depends on assumptions about how the parts of the idiom contribute to the interpretation of the whole. Where we can show that the

verb of an idiomatic VP is analyzed as identical to a use that does not show up in the passive in non-idiomatic uses, we have satisfactorily demonstrated that the failure of the idiom to show up in the passive has nothing to do with any idiosyncratic syntactic marking assigned to it. But we can push a little further here, because the behavior of these idioms can shed some light, indirectly, on the nature of lexical exceptions.

Lees 1960 labeled as "middle verbs" a class of verbs including cost, weigh, meet "encounter," marry "épouser" which appeared not to occur in the passive (or to take manner adverbials, though that issue is now hopelessly clouded). And there have been several attempts to show how the behavior of such items could be predicted from the interaction of various syntactic rules and proposals as to how the behavior of such items might be listed in the lexicon. More recently, however, it has been noted that the relevant judgments are not easily predicted. For example, Lakoff and Peters 1969 argued that their account of the ordering of conjunct movement and passive could predict the unacceptability of sentences like 44:

44. *The table is touched by the chair a foot above the ground.

But Dougherty 1970 pointed out (in the service of another point) that their argument equally well predicts the unacceptability of 45:

45. France is touched by Germany on one side, and Spain on another.

Similarly, Chomsky 1965 argued that live in, but not die in should be accorded distinct representation in the lexicon, in order to account for the apparent pattern of acceptability shown in 46 and 47:

46. England was lived in by many great men.

47. *England was died in by many great men.

But as Ivan Sag has noted, a combination of stress and contrast may make die in more acceptable in the passive; his example is:

48. This England, which has been lived in, and died in, by so many great men.

Most of the other relevant judgments about lexical exceptions to passive are similarly blurrable, it seems.

It is more plausible, then, that the explanation for such exceptions can be laid to pragmatic factors; we would like to be able to say that the unacceptability of a given passive of this type is due to the fact that the nature of the particular predication is inconsistent with the function associated with the passive construction. Bolinger 1973 has argued along these lines, suggesting that a passive is acceptable only when the action referred to by a verb is notionally transitive. He notes a contrast between 49 and 50, for example:

49. *The bridge was walked under by the dog.

50. The bridge has been walked under by generations of lovers.

The fact that generations of lovers have walked under the bridge, he suggests, "affects" it; the fact that the dog walked under it does not. (We might argue similarly with respect to the difference between 44 and 45.)

If the difference between 49 and 50 is easy to appreciate, however, it is very hard to describe. Bolinger's account of the meaning of passive suffers from all the problems we noted earlier. This does not mean that we must reject it, but merely that we cannot formulate it satisfactorily given the present state of the metalanguage. In fact, some other observations about the behavior of the different psenses of verbs like hit and take suggests that we have no choice but to assume that all exceptions to passive have pragmatic explanations.

First note that sentences like 51-53 are entirely acceptable, although VP deletion has applied to delete one use of a verb under identity with another:

- 51. John touched the wall right above where the table does.
- 52. The paper I dropped hit the ground near where the meteor had.
- 53. The IRS took more of her money than the yacht ever did.¹²

Note, however, that one of the uses can be passivized, while the other (generally) cannot:

- 54a. *The wall was touched by the chair.
- b. The wall was touched by John.
- 55a. *The ground was hit by the piece of paper I dropped.

- b. The ground was hit hard in several places by the meteor.

56a. *Lots of money was taken by her boat.

- b. Lots of money was taken by the IRS.

Recall that in Chapter Two, we argued that such behavior with respect to anaphoric processes was evidence for the fact that the two uses must differ only pragmatically; semantically, there can be only one item take, hit or touch involved in each of these sentences. But in that case, we have no hope of being able to characterize linguistically the class of lexical exceptions to passive, since the determination that a verb is to be given an "active" or "middle" interpretation cannot be made on linguistic grounds alone.

The two observations we have made do not suffice to predict the behavior of all idioms with respect to passive. In other cases, passive seems to be odd as a result of factors like those that Keenan has described in terms of "relative independence of reference"; this would explain the oddness of *John's mind was made up by him, or "John's tongue was held. Still other exceptions may be due to factors that are even less well understood than these. (I suspect that the relative oddness of *Time was marked by John is due to the same factors that make it odd to say ?Water was spouted by the fountain. These are likely particularly salient cases of the problems that arise when an indefinite is placed in a focus position at the expense of a definite; cf. ?A book was read by him.) But the point should be

sufficiently clear. We can only explain regularities like these by an appeal to the notion of degrees of convention, and we can only explain that notion by talking about meaning in terms of the ways in which speakers are likely to apprehend one another's beliefs.

4.2.3 "Do so" and "?"

Before concluding this discussion of idioms, let me sketch out a way in which the same kind of analysis can be brought to bear on a problem of a very different sort.

Consider the anomalous behavior of the do so construction, which in some ways resembles "surface anaphoric processes," such as deletion (see Hankamer and Sag 1976 and Sag 1976), and in other ways resembles deep anaphoric devices such as do it. Like VP deletion, but unlike do it anaphora, do so cannot function indexically; one cannot point at a man trying to climb a tree and say,

- 57a. ?Do you think he can do so?
- b. ?Do you think he can?

Though one can of course say,

- 58. Do you think he can do it?

Similarly, do so patterns with VP deletion with respect to the "missing antecedents" phenomenon, though judgments tend to be more labile here:

- 59. I've never ridden a camel, but Jack { did
and it stank horribly. { did so
 *did it

With respect to other tests, do so seems to fall out somewhere between VP deletion and do it anaphora. For example,

consider the requirement that the anaphoric element be parallel in form to its antecedent:

60. The garbage needed to be taken out, so
 John { *did.
 { ?did so.
 { did it.

At the same time, there is good reason for supposing that the do of do so is to be identified with main verb do, rather than as an Aux. For one thing, it takes Auxes of its own:

61. John shook hands with the mayor, but Mary
 didn't { *do.
 { do so.
 { do it.

Moreover, like main verb do, it is not generally used when its antecedent is a stative verb:

62. They said the book would contain three chapters,
 but it didn't { [∅]
 { *do so.
 { *do it.

In this way do so is unlike other surface anaphoric devices, for it contains lexical material, and its use is constrained in part by semantic conditions.

There is some question, of course, over how the distinction between deep and surface anaphoric processes should be syntactically represented. It could be argued that only the former are base-generated, or that the two processes are distinguished in the kinds of base-generated elements that they contain, or even that they involve the application of two sorts of transformations. But however we account for the distinction, we will want to be able to explain the

peculiar status of do so. Let us focus here on one observation about the construction. Though do so is not comfortable with stative antecedents, it seems less aberrant there than do it, as noted by Ross 1969. Example 63 is from a recent book on the IQ controversy; 64 is given by Ross:

- 63. None of the examples we talk about here will involve twin studies, but later ones will do so.
- 64. ?If any of the samples contain Strontium 90-- and some have been known to do so (*do it)-- throw them out.

Let us grant, for argument's sake, that these judgments are correct. Then the status of the judgment represented here by "?"--that an example is "partially acceptable" or "so-so"--is quite interesting. I would not suggest, as some have done, that we should allow the grammar to determine degrees of grammaticality; even if we were willing to forgive the affront to simplicity, this line doesn't explain anything. At the same time, we don't get off the hook simply by saying that the example is wholly ungrammatical, or wholly grammatical, for the speaker who assigns it an equivocal judgment; we are still at a loss as to how its marginal status can be explained.

Taking the broad view, we might say that do so is in a process of change, and that do is in the process of being reanalyzed in this use, having become "Auxier" than the do of do it. This is true as far as it goes, but how could we

incorporate this observation into the linguistic knowledge of an individual speaker, without attributing some sort of squish to his grammar? We could argue as follows.

Suppose that a given speaker S analyzes the do of do so as main verb do, on the basis of his acquaintance with a corpus of utterances C containing examples of other sorts of uses of main verb do, so and other expressions. And suppose also that S believes that it is plausible that other speakers familiar with corpora like C will analyze the construction as he has, and that they will attribute that analysis to one another etc., so that he feels justified in assigning this analysis to normal belief. But suppose, also, that S is aware that there is some systematic variation in the uses of do and so, so that it is possible that there should be some group of speakers who analyze the construction on the basis of a corpus C' which contains no tokens of some of the use-types exemplified in C; and that S believes that the analysis of do so as containing Aux do would be plausible on the basis of acquaintance with C'. And suppose, finally, that S would attribute a mutual awareness of this variation and of its consequences to other members of the community, so that he could assign to normal belief an awareness of the possibility that some members of the community will come up with the Aux analysis, and that they will assign it to their systems of normal belief. Then depending on the certainty with which S can make these

assumptions, he may judge 64--with do so used anaphorically for a stative--a possible utterance, since it is consistent with the analysis of do so that would be constructed over a kind of corpus with which it is plausible to assume that some speakers are in fact acquainted. Note that it is not necessary that any speaker should actually have made the Aux analysis of do so, or that S should actually have heard spoken an utterance like 64, in order that S should equivocate in rejecting it. It suffices that S should believe that the Aux analysis would be normally regarded as plausible.

The facts surrounding the use of do so are quite complicated, and we can only make some suggestive remarks here. Let us assume that do so would be analyzed as involving uses of do and so that were entirely free only if main verb do were freely used intransitively as a rough synonym for "act" or "behave," and if so were freely used as a pro-adverbial: then the force of the expression would be roughly, "act in such manner." Do and so did have these uses as late as the nineteenth century, but their patterns of use have since been changing, as a result of more sweeping disruptions in the language, whose source is ultimately quite mysterious. These changes, in turn, have left a residue of synchronic variation behind them.

At an earlier stage of English, so and thus were used as "manner demonstratives," and could function

indexically or anaphorically in any adverbial slot. Thus one could say,

65. Mary was delighted at the news, and when so (thus) happy, she seemed to me the most beautiful creature in the world.

66. While so (thus) engaged, he seemed distracted.

67. Your dinner can be delayed, if you choose so (thus).

68. Hold the tennis ball so (thus).

All of these uses are either obsolete, or on the way, having yielded to forms with this or that. The use in 65 is rarely encountered; we now say "When she was that happy . . ." (so may be used with adjectives as an intensifier, but not anaphorically for a previously stipulated indication of degree). Except in formal style, we would now say, "While engaged in this," and "if you choose that (it)." And the demonstrative uses of so have largely been superseded in conversation by expressions like this way or like this. At the same time, so has acquired new uses which might not be licensed by the assumption that it was a pro-adverbial, such as in

69. I am so a member in good standing.

These changes are surely tied to changes in other uses of so, as an introductory particle, ("And so to bed") and as a conjunction; we might want to lay to the loss of this pro-adverbial function the reanalysis of so that, as in:

70. He chose his arguments so that none could refute him.

We would now parse 70 with a clause boundary after arguments, but 200 years ago the boundary would more plausibly have been placed after so, as suggested by the obsolete alternative "He so chose his arguments . . ."; so is no longer analyzed as anaphoric for the content of the result clause.

The reasons for this shift are beyond us; so is the proper analysis of the data. We may conclude, however, that the analysis of so as a pro-adverbial depends on familiarity with at least some of the uses exemplified in 65-68--the more the merrier--and that it is normally assumed that there is a good possibility that not all speakers are acquainted with these uses.

Main verb do has also changed in use. Two hundred years ago, one could use it with any adverbial to mean "behave in such-and-such a manner." The uses in 71 and 72 are now obsolete; those in 73 are at best highly formal:

- 71. How shall I do to answer your letter?
- 72. How did they do to get money?
- 73. You did imprudently (wisely) to select him as your companion.

Do is still used with as-clauses, and with adverbials like likewise, but these uses are also somewhat formal, and are not part of the idiolects of all speakers.

- 74. You should do as your scoutmaster tells you.
- 75. John shot a pigeon, and Mary did likewise.

(Do like S is entirely normal, however, and the formality

of 74 may be due to the formality of as in these contexts. It may be, accordingly, that intransitive do is simply coming to be subcategorized only for sentential complements.) Do is now normally used in all contexts only with adverbs like well and better, and these uses are in some measure conventionalized. Note that 76 is somewhat archaic:

76. You did well in coming here to see me (to come here to see me.)

Do in do well is used only to refer to the performance of a certain routinized task or activity.

It is reasonable to assume, then, that do so is normally analyzed as in some measure conventionalized by all speakers, and that the degree to which it is conventionalized is a function of the speaker's familiarity with other uses of do and so. The speaker familiar with expressions like do likewise, and with the uses of so in 71-73, may analyze the expression into main verb do + pro-adverb so, provided he believes that other speakers will be familiar with these uses, and that they will assume a common familiarity, etc. In that case, we may assume that he will not use do so anaphorically for statives. But the speaker who is not familiar with these uses will come up with another analysis, say in which do is an Aux, and so a flag that indicates that VP deletion has taken place. Note that even a speaker who is familiar with these uses may come up with the Aux analysis, if he believes that they would not

be sufficiently familiar to other speakers that they would assume that the main-verb analysis was normal, or that they would assume that other speakers would be familiar with the uses, etc. Thus the mere fact that a speaker is acquainted with the uses that license a certain analysis of a form does not guarantee that he will himself adopt that analysis; he must also reckon what other speakers are likely to assume about the relevant uses.

Mark Aronoff (personal communication) has suggested that the anomalous behavior of do so is due simply to the fact that it is by way of becoming a learned expression, and as such is subject to reanalysis and hypercorrection. This is quite consistent with what we have argued here; we have gone further only in trying to explain why do so is becoming more formal. Since the uses of do and so that license do so are normal only for earlier stages of the language, we can be most sure that another speaker will be acquainted with them when we believe that he is familiar with a large corpus of formal written English, in part because he is then more likely to have encountered them in older texts, and in part because of the deliberately archaic character of formal prose. Accordingly, do so has come increasingly to be restricted to communicative contexts in which knowledge of the uses that license it may be presupposed; that is, to formal style. (One more case of a use creating its own use-conditions.) This may

explain some observations made about do so by Bolinger 1970, who suggests that do so has an "adverse connotation" ("It is neutral," he says, "So is somber."). Some of his examples are given as 87-80; (the judgments are his; others might only think the starred examples a bit odd):

- 77. *If you were invited to go to a cocktail party at your friend's house, wouldn't you do so?
- 78. If you were ordered to attend a strategy session at the Pentagon, wouldn't you do so?
- 79. *If kissing a girl seems like fun to you, why don't you do so?
- 80. If calling me a liar will relieve your mind, then do so.

Without getting involved in the details of the argument, we may speculate that it is the association of do so with formal style that gives it this character; formal style, after all, is associated with the exercise of authority, and such expressions may be used in conversation to invoke authority, especially when referring to activity that is unpleasant or coerced. In this regard, it is interesting that most of Bolinger's examples involve second person subjects for do so, because we would then be tempted to construct a spoken context around them, and the slightly formal do so becomes more conspicuous.

This account is obviously sketchy, and it would be hard, in any case, to devise a methodology for getting at the assumptions that speakers make about one another's assumptions about how constructions like these are analyzed.

It may be that we will never be able to free ourselves from the need for speculative argument. At the same time, the general line seems promising, and might be turned to explain the anomalous behavior associated with other phrasal idioms of this type, such as each other, himself, anymore, and the for-to complementizer, with resulting economies for syntactic description.

4.2.4 Theory-Change, Use-Change, and Reanalysis

New conventions arise, of course, out of historical processes, and it is the speaker's perception of these processes--whether accurate or not--that determines the degree to which he will assume that a given use has become conventionalized, as we saw in the last section. In a sense, the difficulty that a purely semantic account of meaning would encounter in trying to deal with conventionalization reflects the more general difficulties of trying to accommodate any such account to the facts of change itself. If we draw a sharp line between meaning and use, then we are always faced with the problem of how to determine at what point a change in use results in a change in the linguistic system. We alluded to this problem in Chapter Two, when we asked on what grounds we distinguished metaphorical word-uses, which are supposed to be generated by the theory of semi-sentences, from standard uses, which are supposed to be determined by the grammar; without a

good account of the difference, we can scarcely hope to explain how word-uses may change from one class to the other.

This problem has concerned linguists since the distinction of synchrony and diachrony was first introduced. (See, e.g., Schuchardt 1912.) If it has more recently been ignored by investigators of linguistic semantics, that is more likely because of its intractability than its irrelevance. The current attitude seems to be that few useful generalizations can be made about meaning-change; Bréal's attempts to adduce "laws" of semantic change on the model of sound-laws are now the objects of universal obloquy. Yet we have seen that the attempt to systematize synchronic meanings on the same model are equally subject to criticism. If meanings abide in culture, then historical semantics is no less an impossible undertaking than cultural history, just as synchronic semantics is no less impossible than anthropology. In this section, I will try to sketch out very briefly some of the things that we can hope to say about the phenomena that have been labeled "semantic changes."

Writers on semantic change have usually agreed that the term can be used to cover either of two processes, which we may call "redefinition" and "transfer." In redefinition (or "change in sense" or "substitution") the category designated by a term is redefined, as a result either

of a change in the category, or in the normal theory of the category. The Modern English development of words like parliament, felony, artillery, atom, or Germany instances this process. In transfer (or "addition/loss of senses," or "displacement") a word comes to be normally used to refer to an entirely new category (or ceases to be used to refer to a category, though this aspect of the process has been largely ignored). Thus OE bēdu "prayer" has become Modern English bead, and ME grain "seed" has come to be used, by stages, to refer to a longitudinal arrangement of wood fibers. Redefinition has been of primary interest to philosophers, who have asked, for example, on what grounds we say that Rutherford was talking about "the same thing" as a modern physicist when he spoke about "atoms"; and to literary historians, concerned to know, say, how the Renaissance "concept" of the things designated by words like nature and sincere was different from our own. Linguists and psychologists have been primarily interested in the processes of transfer, for the reasons that we mentioned in the discussion of polysemy in Chapter Two.

We will talk about both of these topics, but not as "semantic change," which implies changes in the linguistic system itself. Rather, let us talk about changes in the relation between a term and its extension, where "extension" is taken to be the class or classes of things that the term

is normally used to refer to. This relation can change in two ways. The extension can change, either through redefinition of the classes that comprise it, or by the addition or loss of categories of possible referents. In that case, a term will change in use. Or the nature of the relation itself may change, either through a reanalysis of the array of functions that determine the use of the term, or through the conventionalization of various uses. In that case, the use of a term need not be further affected.

Several factors may contribute to use-change, and the apparent unpredictability of developments, particularly in transfer, has led most linguists to conclude that historical semantics cannot be systematized. With sound-change, we can rule out certain changes as impossible; thus we assume that there could not be a change in any language in which /t/ > /m/, without a certain number of necessary intermediate stages. But the shift of bead from "prayer" to "small perforated ball" seems entirely capricious, and if the French use of grain or the German use of Perle are less spectacular, they seem to be equally undetermined. But this is only one aspect of the problem raised by change, which corresponds to that Weinreich, Labov, and Herzog 1968 (henceforward WLH) discussed as the problems of "constraints" and "transition": what language states can immediately succeed one another, and what intervening stages must be posited between any two forms of a language for a

given community?¹³

For some of the other problems that WLH mentioned, however, we may be able to provide a more interesting solution for use-change than for form-change. In particular, consider the problems of "actuation"--why did a given change happen when and where it did?--and "spread"--how was the innovation transmitted to the community at large? Actuation and spread can't be treated independently of one another, though the problems are not identical. (Cf. WLH's critique of Hoenigswald 1960, who treated all sound change as borrowing.) It makes no sense to put the actuation problem as "who was the first speaker to use the new form, and why did he do it?" For one thing, we know that a new pronunciation is likely to have been used countless times by different speakers throughout the community before it becomes a normal variant for even a tiny subgroup of the community, and only then can it be called an inceptive change. And when the change is phonetic or syntactic, we cannot say anything interesting about why that sub-group should have adopted that change then, since the determination will depend solely on the relative prestige of individual members. There is nothing that we can point to in the experience of Lower East Side Jews that could explain why they should have adopted a raised vowel in bad before Italians did. The problems become linguistically significant only after the linguistic change has already begun; i.e.,

"when one of the many features characteristic of speech variation spreads throughout a specific subgroup of the community . . . [and] assumes a certain social significance." (WLH, p. 186) Then we can invoke more generalized models of social interaction, population movement, and so forth to explain how it may spread throughout the community at large, though the problem is always complicated by continuing change within the community, and by the possibility that the innovation may be reevaluated socially, or reinterpreted linguistically, at any point along the way.

With use-change, we have other opportunities. Let us begin by assuming a homogeneous speech-community, so that the problems of actuation and spread can be collapsed; we will ask, "Why does a new use come to be adopted by a community when it does?" And let us further confine ourselves to cases in which a use-change is motivated entirely by the requirements of communicative efficiency--we won't talk at all about the social factors that lead to the adoption of loan-words, or to dialect-borrowing, and we will wait until the end of the following section to talk about the role of metaphor in creating slang words whose introduction is strictly unnecessary from a communicative point of view. In this way, we can arrive at a schematic explanation for all use-changes and reanalyses, though we will have to be chary in applying it to actual cases.

Both use-change and reanalysis, we can assume, result

from changes in beliefs. Let's start with use-change. Suppose we have a word w, which is normally used to refer to members of categories A, B, and C. Then a change in the normal theories of any of these categories may lead to a change in the use of w. If the category is redefined --if the criteria for class membership are revised--then w will be used to refer to new kinds of things, which did not qualify under the old criteria. So, for example, France is now used to refer to an area that includes Alsace, and cub is no longer used to refer to the young of apes and whales, which are no longer included in its super-ordinate category of "wild beasts." (We are skipping over an important problem, of course, when we say that the category has been redefined; on what grounds can we connect the old and new categories at all, if the criteria for membership are different? We will have to make do with observing that the categories are constructed around like exemplars, and that the members of the new category bear a family resemblance to the members of the old.) Note that it does not matter which of {A,B,C} w is taken to designate; the use of artillery will change if we revise the criteria that either define artillery pieces or the military units charged with their use.

Change in normal beliefs may affect use in another way, because such changes may affect the cue-validity of the functions that connect the normal uses of a word, so

that new free uses are licensed, or old ones rendered less than rational. Suppose, for example, that w is normally analyzed as designating A, and that our theory of A should change in such a way that there is no longer a best function from A to B, and B would be better identified as the value of a function that takes some other class D as its argument. Then the use of w to refer to B would no longer be licensed by normal beliefs. And of course, a change in the normal theory of B, or of D, could also have this effect. In fact, suppose there was another category E, such that E and B were both members of the same superordinate "natural" category, and a change in the theory of E resulted in there being a better function from A to E than from A to B. Then the network of "best" functions from A would again be disrupted.

Such changes need not involve the criterial properties of the relevant categories; beliefs about the contingent properties of things and kinds of things may be equally relevant to the determination of cue-validity. (For example, it is not criterial of newspapers that they are published by just one organization, or of the organizations that publish them that they publish only one newspaper, but a change in either situation would reduce the cue-validity of the function from one to the other.) We might also note that changes in theories which lead to use-changes need not involve only revision of factual beliefs;

changes in evaluative judgments may also affect cue-validity. So, for example, it was the addition to the normal theory of plastic of the proposition "Plastic is used to make cheap imitations of more expensive things" that licensed the use of plastic, during the 1960s, to refer to a vulgar, artificial lifestyle; the subsequent loss of that proposition ("But it's real plastic," children say now) probably contributed to the extinction of that use. Finally theories may become either more or less accessible in the system of normal beliefs, but that question is better saved for a discussion of belief-heterogeneities.

Consider a fairly simple example of the ways in which theory-change can lead to use-change (though we will always be faced with enormous difficulties when we try to isolate the theory changes that have led to an actual change in use). The toys that children make of folded paper to sail across the room are referred to by the word that designates airplanes in a number of languages; in English, they used to be called "darts." In this case, we can assume that the change in use arose when a new category was introduced such that there was a better function from it to the toys than there was from darts. But of course, the etiolation of the theory of darts, as the game became less common, may also have contributed; so may a general uplift among schoolboys, which led them to put the toys to more benign

purposes. (I don't know whether the toys looked then as they do now, but it is conceivable that a design change, perhaps due to a change in the composition of copybook paper, also played a role.) In any event, a series of theory changes led to a recalculation of the cue-validity of various functions, and a change in use resulted.

The same recalculations that lead to change in use may lead to change in meaning. Suppose, in our schematic example, that the theories of A and B should change so as to affect the relative cue-validity of the function from A to B and its inverse, so that the function from B to A becomes better than the function from A to B. Then speakers must re-analyze t as designating B; if they continue to use t to refer to A it will be by some function other than the identity. (Where there are several normal uses, things become even more complicated; a change in the theory of A and B could lead to a reanalysis of t as designating C, or even some new category that it had not previously been used to refer to.)

The history of English provides a number of celebrated examples of such meaning-shifts. Bead is one; others include the development of Mod E deer (which designated simply beasts in OE), dainty (from Latin dignus), meat (originally just "food"), hackney (from "mare") and silly (originally "holy": cf. German selig). In each of these cases, however, it is difficult to contrive access to

the relevant theory-changes, and a number of extraneous factors appear to have played a role as well. The best that we can do, given the philological evidence available, is to chronicle changes in the use of the term; we may then be able to tell when, if not why, the meaning was reanalyzed. Suppose, for example, that w is used at some time t_1 to refer only to A, and comes to be used at some later time t_2 to refer to both A and B. It will be hard to tell, in most cases, whether w is still analyzed as designating A. But if w should acquire at t_3 a new use to refer to C, such that there is a good function from B to C, but not from A to C, then we must assume that reanalysis took place at some time between t_2 and t_3 , since if w continued to designate A, its use to refer to C would not be licensed.

Bead, for example, is used in Old English only to refer to prayers. (We should really say "bedu is used . . .," but we will ignore the problems of form-change.) This use continues into the sixteenth century, but after 1550 it is used only to refer to the prayers of the rosary, a use that continues until the eighteenth century. Its use to refer to rosary beads is first cited in 1377. (The rosary itself was introduced in England, along with the virgin cult, in the thirteenth or fourteenth century.) We assume, accordingly, that the reanalysis of bead as designating rosary beads took place at some time between 1400 and 1800.

But we can fix the point of shift more accurately than this: the use of bead to refer to beads in general became common in the early seventeenth century. (The OED gives one such citation for 1400, but none thereafter until 1599.) Since there is no good function from prayers to beads in general, we assume that this use could have arisen only after the meaning of bead had been reanalyzed. We have further evidence for this in the use of bead from c1600 to refer to drops of sweat and blood. Moreover, the occurrence of forms like prayer-bead (1630), which suggests a need for differentiation, leads us to assume that by that time bead must have been taken as designating beads in general. We conclude that the shift in the meaning of bead from "prayer" to "rosary bead" must have taken place between 1400 and 1600, and that the further extension to beads in general must have followed almost immediately.

But we cannot begin to unravel the changes in the theories of these and other categories that led to these developments. Doctrinal and liturgical problems aside, we can't talk about bead in isolation; there is also the "semantic field" to consider. Beads were also called grains in the fourteenth century--the French influence may have played a role here--and the history of grain is of almost Balkan complexity. Moreover, bead "prayer" was in competition with French orison from the twelfth century, and with French prayer from the fourteenth, so we can

assume that there were other social factors that made bead superfluous in that use. In any event, we can be assured that heterogeneities of which we are now ignorant played an important part in the change, and we have not yet talked about how heterogeneity affects the picture.

Thus far, we've managed to avoid having to talk about the "collective awareness" of the community, with all its attendant problems, by addressing ourselves to the way in which an individual speaker perceives his community. When we come to deal with change in a heterogeneous community, however, this becomes a little awkward, since the beliefs of a community can change, say as its members are replaced, without any change in the beliefs of individual speakers. So we had best talk here about generations of speakers. Changes in the community must invariably be accompanied by changes in the systems of belief against which younger speakers act, with effects both on use and the analysis of use.

It may be, for example, that there is a change in the distribution of a certain set of beliefs among the members of a community. Suppose, for example, that in 1930, 90% of all Americans believe that p and 10% that -p, and that in 1950 these figures are reversed. All things being equal, we would assume that p would rationally be assigned to the system of normal beliefs of a typical speaker in 1930, and -p in 1950. Then if p licenses a certain use R of a

word w, R will be entirely rational in 1930, but not in 1950. In that case, one of two things may happen: R may be replaced by some other word-use, as (paper) plane replaced paper dart. Or R may be reanalyzed; what was a free use may become in some way conventionalized. (We would explain the change in the status of the use of nylon to refer to stockings in this way.) This may lead to other changes in use; speakers may no longer ignore the difference between R and other uses of w in using anaphoric devices, for example. Not only change in the beliefs of speakers may have this effect. Suppose that in 1950, 90% of the members of p are aware of the fact that some people used to believe that p, while in 1970 only 10% are so aware. Then if p licenses R, R will be generally accounted conventionalized as in stage III above in 1950, and as in stage IV in 1970. And this may also affect use. If R involves the use of an idiom, for example, speakers may no longer passivize it.

Other kinds of change in the community may have similar effects on use. Suppose the use of w to refer to A is licensed by a belief p, such that it is normally believed that members of the subgroup S entrusted with collective authority over A behave as if they believed that p. Then this use will be accounted free, as we saw above. But a change in the relative authority assigned to S may lead to a change in the evaluation of the use. Authority over A

might devolve on the community at large--say, as instances of A become more common. Or S might be supplanted in the public esteem by another subgroup S', whose beliefs are known to differ from those of S. In either case, we could expect either that use will change, or that its status will be reanalyzed.

There are other variations on this theme. We can explain the spread of a new use, for example, as a result of the process whereby the beliefs of collectivity become accessible to the system of normal beliefs, or pass into it. But this sort of schematic treatment has no predictive value; it is interesting only insofar as it allows us to organize the facts surrounding developments that have actually taken place. The details of such cases are usually quite complex, and hard to document; the task has to be left to social or cultural historians. If there is any value at all in our talking about a representative example, it is because the historians, while recognizing the importance of use-change, are often unable to evaluate the process.

We had best stick close to home. Over the past forty years, swing has come to be used to refer to a musical genre, and the age associated with it. The process began with the introduction of a new local use of the verb in the twenties. As Marshall Stearns tells it (1956:157):

During the 'twenties,' the expressions 'sweet' and 'hot' were in use among musicians to distinguish between the

music, for example, of Guy Lombardo and Duke Ellington. Both were generally thought of as jazz. The word 'swing' was only a verb used to describe the basic jazz quality: good jazz should 'swing.'

We don't know anything, of course, about the circumstances that led to the introduction of this use, but we should note two things here. First, this use would have been accounted slang, and had an obviously metaphorical source. We will further explore the relation between metaphor and "use-conditions" at the end of the next section, but in connection with less obscure examples. Second, this use was local, and remains so. It is still encountered, for example, in liner notes, but like (wind) shadow, it has no general currency, and is reserved for situations in which the speaker can assume that the hearer is a member of the same collectivity as himself. (It did, however, lead to the use of the participial swinging, which has been re-analyzed as it passed into normal use.)

From this local use another arose, among a somewhat larger group. In the late twenties and early thirties, swing was used to non-evaluatively refer to the activity of playing in a "loose" jazz fashion. We can trace the changes that led to this change in use: around this time, as "hot" dance orchestras became popular, bandleaders began to try to induce a swinging sound from larger ensembles, either by hiring a few hot jazzmen to take solos, or by commissioning "swinging" arrangements. Once the responsibility for a swinging sound passed to the arranger, one could say

that a band "swung" without entailing that its members were particularly good musicians. This use, too, remains local. (For example, Henry Pleasants 1969 writes, "Nobody has yet been able to make a string section swing.")

These changes in musical practice led directly to the nominal uses of swing, to refer to the quality that music had when it swung, whether that quality had its source in the efforts of the performers, or in the arrangement. For example, Benny Goodman 1939 refers to "the art of making an arrangement that a band can play with swing." (And cf. also, "It don't mean a thing if it ain't got that swing.") This use led in turn to other nominal uses. Swing was used to refer to a kind of arrangement, or to a song arranged in a swing fashion, as in Bob Wills':

You can change the words of an old song
Re-arrange it and make it a swing

The appearance of these uses in song-lyrics suggests that they had some general currency in the thirties, but they are entirely local today. (For example, Guthrie Schuller 1968 writes, "This extra dimension in the rhythmic impulse of a jazz phrase is what we call 'swing,'" using the term to refer to the musical criterion whereby jazz and non-jazz are to be distinguished.) Inasmuch as these uses were normally available in the thirties, we can assume that the introduction of swing to refer to the music played by the big white dance bands like Goodman and the Casa Loma Band was perceived as a free use; these were "swing bands"

that played "swing music"--i.e., music that had swing. (In this connection, Stearns mentions that there is a story that the expression "swing music" was introduced by the B.B.C. as a euphemism, to avoid the sexual connotations that they perceived in "hot jazz." "They were probably right," Stearns says, "but they were fifty years too late." Even if this is true, it would not explain why swing caught on throughout the community; we must assume that the use was perceived as motivated.) But the general public was largely ignorant of the jazz sources of Goodman's music, and the category of swing music came to be normally defined by a variety of criteria, such as tempo (since it was closely associated with certain dances) orchestration, timbre, as well as rhythmic accent. Thus swing music was normally defined as a certain variety of big (white) band jazz. (Cf. Aaron Copland, writing in 1941: "The revival of jazz of the so-called 'hot' variety, which came into vogue around 1935 under the name of swing. . . .")

As jazz continued to develop, however, and as it came to influence all styles of popular music, the theory of the category designated by swing changed; the music of Goodman and Miller had to be normally differentiated both from other jazz-influenced popular music, and other styles of jazz. (The late thirties and early forties saw an enormously popular Dixieland revival, for one thing, and the earliest developments of the jazz style called bop, which

enjoyed considerable public attention. Then too, the war, and a tax on dance floors, made the big bands impractical, and the big swing bandleaders began to employ smaller groups, and changed their orchestrations accordingly.) Speakers could have continued to use swing, of course, to refer to all jazz-influenced popular music, but it would not then have been clear what swing was opposed to. (Henry Pleasants writes 1969, p. 171, "Swing did not die. It simply became so pervasive a part of the musical vernacular that the jazz musician, in his quest of distinction . . . disowned it.") Rather, they continued to use the word as if it designated the kind of music of which Goodman et al. were the most salient examples, but that category was now analyzed as a genre, rather than a style.¹⁴

We have been able to give only the outline of the changes in musical practice and trends that led to the changes in the use of swing, and this is most likely wrong. At best, we have shown what sorts of considerations must be taken into account in explaining how a word came to be used in a certain way. Meillet's apothegm "Chaque mot a son histoire" is apposite here; we couldn't hope to be able to adduce any general laws of use-change. But at the same time, we have seen how use-change can be used as evidence for theory-change. The nominal uses of swing, for example, could not have developed until the theory of what made jazz sound like jazz had changed.

4.3.1 Metaphorical Uses

The literatures on "metaphor" are so extensive that it would be futile even to make a pretense of responsibility to them. For the present purposes, I will simply take "metaphor" as a technical term, to refer to tokens of uses of words that satisfy conditions 1 and 2:¹⁴

1. A speaker uses a word or expression e to refer to a in a context C, even though there is another expression e' available to the speaker such that the likelihood of successful reference to a would be increased by use of e', given the background of beliefs that may be rationally assumed to be governing use in the context.

2. The speaker knows that his use of e is not rational, and expects his hearer to recognize that he knows, etc. And the use is intended to further the purpose of the talk-exchange; i.e., the speaker is not "flouting a maxim," as Grice would say. In other words, we will say that metaphors are word-uses in which the communication of propositional content is subordinated to some other conversational purpose; they form a proper subset of the class of language-uses that Grice calls "conversational implicatures."

Let me clarify one or two points, before going on. As formulated, 1 excludes "dead metaphors"--uses that are licensed by the linguistic practices of the community, such as when we call someone a "golden boy" or a "Don Juan." To the extent that these uses are conventional, they are

justified on purely communicative grounds, and do not involve violation of the rules of communication. For the moment, we can also avoid discussion of "faded metaphors" --either as when we call a person "a cold fish," or life "a journey," or as when we use a phrasal idiom like make tracks. (But we will return to these in the next section.)¹⁵ Let me add also, that I do not mean for condition 2 to be taken as anything more than an allusion to Grice's characterization of the conditions for conversational implicature, perhaps amended according to the suggestions for various commentators.

I won't have anything in particular to say about the "problem" that most of the people interested in metaphor have accorded a central position: how the name of one thing (or idea) can be made to stand for another; we have been talking about little else, after all, over the last three chapters. As we have defined metaphors, they do not differ in quality from other uses of words--no more than sentences used "literally" differ from sentences used to imply. Nor is there any difference in the inferential strategies whereby we arrive at the referents of words used metaphorically; only our apprehension of the speaker's intention is different, and the question of how that enters into our calculations is beyond the ability of any analytic theory of language to deal with.

Basically then, we already have all the "theory of metaphor" that we can hope for, and we will be able to "account for" these uses ("how," not "why") without introducing any new notions. Let me show this with two simple examples, and then give over the rest of the section to a discussion of a literary metaphor of the sort that has generally figured in the literature, with an occasional side-trip to visit a related problem, or to take a swipe at semantic theories of word-meaning, whose inadequacies here are particularly obvious and embarrassing.

Suppose I say to the mother of an industrious child, "How's the little beaver?" Let's take it for granted that the best function from beavers to persons takes us first to a set of properties that beavers are normally supposed to have, and from these to an equivalence class of persons that manifest those properties, or other properties that are expressible as the values of good functions from these. (There are any number of such equivalence classes, of course; in another context we could use beaver to refer to a person with big teeth. We could construct a less equivocal example by choosing an argument for the RF such that only one property or set of properties was normally attributed to it; say a fictional character like Mr. Pecksniff. But let's assume that beavers are normally attributed only with the property of industry, and press on.) This use would be entirely rational against the background

beliefs, "The quality of industry is best identified as that property exemplified by beavers," and, more important here, "Your son is best identified by reference to his work habits, rather than to his family role, sex, or age; that is, his industry is his most enduring and salient characteristic." This last belief is not normal; if we are interested solely in distinguishing one person from others, we will do better to choose some property that is more obvious, and less evaluative. What aesthetic force the metaphor carries, then, results from our behaving as if it were assumed that things were other than we believe them to be, as opposed to asserting that they are so; and it is this, I think, that lies behind all implicatures and ironies. (See the end of this section, though as I said, we are interested in the how and not the why.) In interpreting metaphors, then, we have only to recover the beliefs that license the use; they differ from other uses only in that such beliefs would not rationally be invoked for the straightforward communication of propositional content.

We will have no reason here to introduce the rhetorician's distinction between metaphor and metonymy. When we use wheels to refer to a car, for example, or ear to refer to attention, the RF is not a composite involving the hypostasizing function from tokens to types; otherwise, the process is the same. The first use is licensed by the belief that cars are best differentiated from other

possessions by the presence of wheels; the second, that dispositions are best differentiated according to the body-part that each calls into play. The difference between the classes of figures that rhetoricians have introduced (part/whole, whole/part, attribute/possessor, and so forth) reflect solely the differences between the kinds of RF that each involves.

Let's turn now to a metaphor of the classic type, and examine it in greater detail. In "Childe Harold" (III, LV) we find:

3. The castled crag of Drachenfels
Frowns o'er the wide and winding
Rhine.

Frown is used here in an obviously metaphorical way; in fact, this is a very ordinary, even trite, Romantic figure. (Eliot 1937 wrote that one could say of Byron "as of no other English poet of his eminence, that he added nothing to the language, that he discovered nothing in the sounds, and discovered nothing in the meaning, of individual words.") Moreover, the "meaning" of the line is not hard to recover, and we can talk with some certainty about what the referent of frown must be. (No matter, for our purposes, what sorts of things verbs "refer to.") I would not try to give an exact non-metaphorical paraphrase, but let us say that frown is used here to refer to the way in which one thing is disposed relative to another.

It will be useful to consider how semantic theories

would try to deal with uses like this one. Why is it "odd" (or "aberrant," or "deviant") to use frown in this way?

Katz would say that frown is entered in the lexicon together with a requirement that its underlying subject must bear the feature [+human], and that the sentence is for this reason ill-formed. This position has been attacked by McCawley, Lakoff, and Ross, among others, who have insisted that any condition on the use of a verb like frown can mention only the referent of the subject NP. Their examples have sometimes been marginal (and occasionally gory); the point is more convincingly made if we consider cases where a term that designates a non-human NP is used to refer to a human or group of humans, by some RF other than the identity. All of the following examples must either involve violations of selection restrictions, on Katz's view, or require that we accord multiple entries to polysemous items. Yet none of them are aberrant or literary:

4. The school frowns on blue jeans.
5. The tenor frowned at the soprano.
6. Building number 20 is betting on the Red Sox to win in six.
7. There was a tattered Chomsky and Halle on his desk.
8. The newspaper fired John for lateness.
9. The radio is too loud.

Obviously, we cannot adopt anything like Katz's selection restrictions unless we say that the meaning of a term does

not by itself determine what things it is ordinarily used to refer to.

The view has also been suggested that co-occurrence restrictions of this sort should be stated as linguistic conditions on the referents of terms, rather than on the terms themselves. Examples 4-9 would not present any problem for such an approach. Ultimately, however, this view turns out to create roughly the same problems that Katz's selection restrictions do. Suppose that we said that the lexical entry for shatter (intrans.) required that its (underlying) subject must refer to a physical object, or that the subject of give must be animate. Then how would we account for the uses of these verbs to refer to activities other than those they designate, which can be reasonably predicated of other sorts of things:

10. Her hopes shattered.

11. The sky gave its color to the lake.

and so forth? Or consider the havoc that would be wrought in the derivational morphology. We could plausibly analyze the agentive suffix -er as an explicit marker of a function from classes of acts to associated classes of characteristic agents (or instruments, but leave that aside). But then the class of things that could be called writers or caterers would have to be defined as the class of things to which the subject of write or cater can characteristically refer. In which case, the uses of these words to refer to other kinds

of things would also create problems:

12. All the Russian writers on the seventh tier of the stacks have been ripped up and shredded.
13. The caterer was looted and burned during the blackout.

Actually, there are ways around these problems, at no cost to the derivational morphology, which preserves the approach via conditions of reference. And there are classes of terms for which it is indistinguishable from a pragmatic account of co-occurrence, as we'll see below. In the end, however, it fails too in important ways.

Putnam raises the crucial problem in his discussion of natural-kind words like lemon. He asks on what grounds we would say that a machine-lemon was not a "real" lemon. Perhaps we would want to argue that machine lemons do not really "grow." But, he counters, grow is a natural-kind verb, and we determine the criterial properties for "growing" just as we determine the criterial properties for lemons. For Putnam, it would be in consequence of our knowledge of biology that we can predicate grow of lemons (and more recently, of viruses) but not of rocks (or humors).

Putnam's realism aside, the point is valid. Let us say that we have among our categories of actions a certain class--including those called murder, sell, and orate--which are partly defined by the presence of a certain will or intention. Then it would be odd to predicate these actions of things to which we do not attribute will--

rocks, lemons, and so forth. But whether or not we attribute will to a thing depends on our beliefs about it. We could discover that rocks are willful, or that (all other) humans are automata, and that could change our use such that 14 would be fine, and 15 peculiar:

14. The rock murdered the pebble.

15. John murdered Mary.

But this discovery would involve no change in the meaning of rock or human, since the presence of will is not, in this faithless age, a criterion for membership in either class. No purely linguistic knowledge could tell us what sorts of things the subject of murder must refer to.

All these arguments are really beside the point, however. There are simpler ways of showing the inadequacies of semantic accounts of metaphor, because we can have metaphor without semantic anomaly:

16. The Yankees annihilated the Red Sox.

17. He bestrides the narrow world like a Colossus.

Annihilate and bestride are no less "aberrant" here than frown in 3, but their deviance is solely due to the pragmatic oddness of saying that one baseball team annihilated another, or that a man can bestride a planet. It is this oddness that forces the metaphorical readings, not anything in the linguistic form of the sentence. So we cannot ask the semantics to tell us when a use is metaphorical; even under the most powerful theories, it doesn't have enough

information. And inasmuch as metaphors are a subset of implicatures, this is just what we would expect.

Let's explain the oddness of the expression "the crag frowned," in the following way. Frown designates a kind of act in which a certain facial expression is presented to view. (Hard to say whether that expression is criterially the manifestation of an inner state of bemusement or disapproval. If it were, then we could talk about "insincere frowns," which sounds funny.) Now whether or not only humans can frown (for it seems that the act could be predicated of robots or highly intelligent dogs), it is clear that crags are not a member of the class of things over whose activities the category frown is constructed. Then if someone says "the crag frowned," we can assume one of two things:

- 18a. The RF is the identity for both subject and verb; i. e., a "real" crag "really" frowned. In which case the speaker is talking as if his beliefs about crags were very different from what we normally believe about them, such as we might do in telling a fairy tale.
- b. Either crag refers here to something that can frown, or frown refers to something that crags can do, in our system of normal beliefs. Thus in another context it could be that crag was being used metaphorically. (Imagine, say, that the phrase "the castled crag of Drachenfels" should refer to a Lincolnesque German chessplayer.)

We can rule out 18a as unlikely here. Byron does not ask us to believe that the crag of Drachenfels is really animate, not even in the context of the poem; in the

world of Childe Harold, crags do not really feel things, or express their feelings. Among the alternatives in 18b, moreover, we can assume that it is frown, not crag, which is being used to refer to something other than what it designates, since this assumption will give us an interpretation more consistent with the rest of the poem. For one thing, there is nothing else around to which crag could reasonably--or even unreasonably--be used to refer. (Note, however, that this assumption is more comforting when we encounter the line in Byron than in Prévert.)

This takes us to a second question: how do we arrive at an interpretation for the utterance? It has often been suggested (as by Weinreich 1972, Levin 1971) that the interpretations of metaphorical sentences should be derived by the introduction of devices such as "transfer rules," which map the features selected by the verb, say, onto the (term or referent) of the subject NP. In this case, the interpretation of crag would be augmented with the feature [+human] or [+affective], or whatever it is that frown "selects." This proposal fails on the same basic grounds as the semantic proposals for identifying metaphors; it cannot be brought to bear on cases like "The Yankees annihilated the Red Sox," or "He bestrides the world," where no selection restrictions are violated at all. And even in cases like this one, where it could be invoked, it doesn't begin to tell us either what the interpretation is, or how

it arises. Saying that the feature (+human) is assigned to crag is not different in any interesting way from saying that the crag is here "personified," save for being more obscure. We still don't know what the personified crag is supposed to be doing; i.e., what frown is being used to refer to. And finally, this approach makes it impossible to say how it is that metaphorical uses acquire their aesthetic or emotional import. Why should a word-use whose interpretation requires the application of "transfer-rules" be evaluated any differently from a word-use interpreted by means of another kind of formal device? It's true that we have ruled this question beyond our immediate aims or abilities, but we are at least obliged to provide an entree to it by characterizing metaphor in terms of beliefs and intentions, which must be grist for any serious discussion of the aesthetic problems that metaphor raises.

Let me try to give the reasoning that enables us to interpret this use of frown, with occasional excursions. First, we know on the basis of its linguistic environment alone that frown (o'er) refers to a two-place relation; and with the further assumption that the RF is the identity for both the subject and the prepositional object, that the relation holds between a crag and the Rhine. (The crag of Drachenfels, of course, was not an outcropping, but a precipitous rock; the latter use is less common in American English.) On the basis of this information alone, we can

restrict the set of possible normal ranges that the verb could have, though not so far as to be sure that the verb refers to a spatial disposition. If it were "ushers forth," or "given birth to," the referent would be a kind of causal relation; if the verb were "recalls" it would be a relation in the mind of some third party. Moreover, we cannot tell apart from the verb whether the range is the set of relations that an object can assume with respect to a stationary object, as here, or a moving object (as in "The crag blocks the Rhine") since Rhine may refer either to a body of water in motion, or to the path it follows (and it is indeterminate which of these is its designatum).

From its environment, then, we know that frown is being used to refer to some member of the union R of sets of normal ranges, each a set of kinds of relations that can hold between crags and rivers. We will then ask, as with any normal use of a term, which member of R is best identified as being the value of some function that takes the designatum of frown as its argument? In the absence of extended context, it is hard to see how frown could be used to identify a causal relation, or a relation in mind, but we can easily begin to see how frown could be used to refer to a spatial disposition. When persons frown at things, they characteristically (though perhaps not criterially) face them. So if there is a good function from the parts of their bodies that persons present to an object when they

frown on it, to a part of a crag, we can use frown to differentiate one subset of possible physical dispositions.

If the line read "The crag faces the Rhine," of course, it would be assigned a normal interpretation. Face is a member of a large class of verbs, including sit, stand, lie, hold, and touch, reach, which are normally used to refer both to an activity in which an animate assumes or maintains a certain position with respect to an object, and to the disposition that is thus assumed. These verbs are quite clearly "the same" in both uses (cf. "I touched the wall just above where the table does.") but the different uses seem to involve different extensions; we would not want to say that actions form a natural category with their resulting states. When we say "John was touching the table," for example, we may interpret touch in one of two ways; accordingly, we would say that there are (at least) two possible RF's in the normal context. This sort of approach works equally well for verbs like pass and cross (as in "John crossed the highway near where Route 7 does.") which can be used to refer to actions in the course of which one object changes its position relative to another, or to the disposition of a path of potential movement relative to an object.

In passing, it's worth noting that examples like these have often been dealt with by assuming that there is some ambiguity or vagueness in the case signals, such that

the referent of a nominative NP can refer either to an agent, an "experiencer," and so forth. It is more reasonable, I think, to attribute the indeterminacy to the verb, which may be freely used to refer to any of several kinds of relations. This was what the generative semanticists had in mind with their talk of "higher predicates," of course, except that we are not attributing the different interpretations to any markers that are present in the semantics; the various psenses of the verbs are pragmatically derived. (This approach runs into some problems, I should note, when we come to deal with the oblique cases, and I will not pursue it here.)

I have up to now avoided saying what it is that a verb like face designates; I have claimed only that its various uses can be pragmatically derived from a single source. If we want to push the analogy with frown, however, which clearly designates something that humans do (or present), we will want to be able to say that face also designates an anthropomorphically defined category. As we saw, face can be used to refer either to an activity or a disposition ("John faced the wall" may or may not entail that John moved). It doesn't matter, for our purpose, which of these we take to be the relation that face designates; a good RF could be derived in either direction (i.e., from states to activities resulting in them, or vice versa). More interesting to us is the question of

whether the set of dispositions that face can be used to refer to constitutes a single extension, or whether several kinds of dispositions are involved. That is, is the RF "the same" in all of 19-24?

19. John sat facing the wall.
20. The car is facing the street.
21. The room faces the park.
22. The T.V. faces the bed.
23. The chair faces the desk.
24. The cliff faces the river.

In each case, we would say that face is used to refer to a spatial relation whereby an object a is disposed towards an object b such that a particular side of a is presented to b (or so that the side is closest to b). But how do we know which side of the object it is? (Note that it need not be the "front" of the object; a house can face a park in front, as well as face an empty lot in back; a room can face north, south and west at the same time. And in any case, front presents the same problem as face does: how do we know which side of an object is the front?) The "face-side" of a car is that side that leads in normal motion; of a room, the side or sides containing openings to the outside of the structure; of a chair, the side from which the "interior" is most easily accessible; of a cliff or mountain, a broad lateral surface. The "face-side" of a T.V. is the side on which a picture appears; the "face-side" of a

radio may be the side containing the speaker, or the side on which the controls are displayed. (A T.V. cannot face something except with its screen-side; a radio may face something even if it has no speakers on the side presented.) The "face-sides" of cannons, knives, and shoes require other definitions.

On the face of things, these definitions seem to vary idiosyncratically from object to object; we could not stipulate a single set of criteria which would characterize the face-sides of all (eligible) objects. But all of the criteria that we apply to objects to determine their face-sides hold of the face-side of persons--the side we present when we face something; that is the side of us from which information is emitted, the side that ordinarily leads in motion, the side providing visual access to the outside, and so on. If we say that face designates a kind of disposition that people assume relative to objects, then we will have good functions to the disposition of various inanimates, to the degree that we can distinguish parts of them as "sides," and determine which of the functions from people-sides best satisfies the inverse-image condition, relative to the conversational interests. (A book that is "facing the wall" may have either its spine or back cover outwards, depending on whether we define the face-side as the one from which physical access to its contents, or visual access to its title, is most easily attained.)

We could go on at length about how such relations are defined on the basis of anthropomorphic models, but we are already too far afield, and the subject has been extensively discussed by Fillmore 1972, though in other terms. Let us return to Drachenfels. We have determined that frowns here refers to a spatial disposition of crag and river in which the face of a crag is presented to the Rhine. We can assume that the relevant side of the crag is a broad lateral surface, for crags do not have openings, nor do they move, or have any other properties to which there might be a good function from the part of the body presented in frowning. (This is a castled crag, of course, and the author could intend that the disposition be determined relative to the "face" of the castle. We'll ignore this interpretation, however.) But after all this, frown is doing no more for us than face would. Assuming that Byron is not violating the maxim of quantity, then, the referent of this use of frown must be a member of a still narrower subset of relations.

Suppose we render the criteria whereby we distinguish the class of frowns from other classes of things as in A and B:

- A. Frowns are actions, or their resultant states, in which the brows are knitted.
- B. Frowns are characteristically the outward manifestation of an inner state of bemusement, worry, or disapprobation.

It was from A that we derived a function to a kind of

disposition; we can map from the part of the body presented in frowning to the part of the crag presented to the river. Are other parts of A relevant to the derivation of the RF? It follows from A, together with assumptions about how faces "normally" look, that C:

C. Frowning requires an act of will.

But we are not in a position to attribute will to the crag, so cannot distinguish sub-sets of spatial dispositions by its presence or absence. Frowns are distinguishable from other expressions in terms of spatial configurations, however, and there could be a function from the disposition of parts of the face to the disposition of parts of the crag. Against a background of normal beliefs, it is true, there is not a "best function" from the features of one to the features of the other, since the features of a rock face are not arranged with any regularity; moreover, their appearance changes according to perspective. Whether we could identify a particular part of a crag as being the value of a function from noses, chins, or brows would depend on its idiosyncratic properties, as well as where we were standing. Yet there is a fair function; we would look for brows (as in "the brow of the hill") towards the top of the crag, and so forth. And if we presented a group of subjects with a number of pictures of crags, we would probably be able to get fair agreement as to which of the crags had a feature that could be identified as "a brow," such that it

was arranged with respect to the other features of the crag as a brow is arranged in a frown. (The expressive function of frowning may confuse the issue somewhat. Consider the act of pursing one's lips, which has no particular affective significance. Where we have a good function from lips to something else, we can identify their configuration as "pursed," as in "the pursed lips of the cave." I am not saying that this use isn't "metaphorical," only that its interpretation requires less adjustment of the system of normal beliefs than "the crag frowned.")

What is most striking about the line, however, is that a word that designates a way of expressing an inner state has been used to refer to the appearance of something to which no inner state can be attributed. It could be, of course, that the expressive function of frowning is not relevant to the identification of the referent (just as we ignore the fact that the "face-side" of the body contains its eyes when we use face to identify the relation between crag and river). As it happens, there is no single term that designates the class of expressions formed by lowering the eyebrows, independent of what they are characteristically used to express, i.e., no superordinate for frown, scowl, glower, etc. So if we wish to identify the arrangement of features of a crag simply as the value of the function from that particular way of arranging one's facial features, we would have to pick one of the

subordinate terms, as if at random. No group of subjects, after all, could be counted on to distinguish scowling crags, frowning crags, and glowering crags with any measure of agreement.

Obviously, Byron did not choose frown solely because it is euphonious with "winding" and "Rhine"; he intended that we should be able to further identify the intended referent in virtue of some other properties of frowning. ("Obviously" here stands for "Because of what poetry is.") We have a good function from frowning to a kind of disposition in which a certain lateral surface is presented, and a fair function from frowning to a disposition in which the broadest lateral surface is presented with an overhanging ridge at its top. What is the next best function to a still smaller class of possible referents?

There is a function from frowning to the inner state it expresses, but how can we avail ourselves of it? Byron's crag is not animated, and has neither inner states nor the will to express them. But there is also a function--not so good--from frowning to the inner states of others. If frowning is the expression of disapproval (let worry and bemusement slide) then we feel when someone frowns in our direction as we do when someone expresses disapproval of us. (Hence "frown on," in part conventional, for disapproval without outward expression.) So we could identify the

referent further as a disposition of such-and-such type such that it induced in its beholder those feelings that he customarily had when being frowned over (say, apprehension and wariness, though our goal isn't an exact interpretation). Here, by the way, is where the "o'er" becomes important, because it distributes the affect over an area surrounding the river, to include the position of any potential viewer.

Now this is clearly a very poor way of going about enabling someone to identify a certain spatial relation. Our hypothetical subjects are not going to be able to come to much agreement in sorting out the crags that make them apprehensive from those that make them feel happy or secure. They might do better than random, but the probability of successful reference in a given context is quite low, given the system of normal beliefs.

One could say, of course, that it does not really matter that we be able to identify the referent exactly, so long as we are given a rough idea of what it is. It isn't as if there were three castled crags in the area, identical in appearance, and we had been asked to pick out that one of the three which made us feel apprehensive. In which case the functions to the "brow" of the crag, and the feelings that the crag induced, are used nonrestrictively, for coloring or "comment." But then why not be explicit, and write:

The castled crag of Drachenfels faces the
wide and winding Rhine
And with its overhanging ledge, it
made to thump this heart of mine.

It is true that we do not have to be able to identify the referent exactly in order to recover from the line enough of its propositional content to be able to follow the poem. But in the process of figuring out what properties the (exact) referent of frown is supposed to have, we have already recovered much more than that. In determining the RF, we have had to select on the basis of the degree to which the use of a particular function was consistent with our system of normal beliefs. On that basis, we said that functions were "good," "fair," or "poor." But in so doing we have had to construct the belief systems that licensed the rational use of the function, asking, in effect, "How would the world have to be such that we could use the words frown o'er to refer to a certain disposition, with a high probability of successful reference?"

Note that we can use words that designate inner states, or the acts that express or induce them, to identify certain natural phenomena. We say that a stream or brook is "raging" or "tranquil"; that the sky or clouds are "lowering" or "threatening"; that the heat is "oppressive"; that the day is "pleasant," or "miserable"; that the frost is "cruel," and so forth. (Note also the pervasive polysemy of words like gloomy, sunny, and bright.) In every case, however, there is good general agreement as to which states are induced by which kinds of phenomena; we needn't worry here about the source of the cognitive associations.

(Nor does it matter that some of these uses are "conventional"--if there were no agreement at all, the convention could not have arisen.)

Most of the phenomena that we refer to in this way are "active," either in their own actual movement (brooks and storms) or in the potential for motion that inheres in them (clouds) or in their direct effect on persons (heat and cold). Where a natural object does not move or act directly on us, we can identify the effects it induces or simulates only on the basis of its appearance, which will vary according to light and perspective. We could expect to agree on its effects only if we could credit what Wimsatt 1960 called "an animate and plastic nature, not transcending but immanent in all things," to which the poetical sensibility provides access. (Later in the canto [XCII] Childe Harold compares an Alpine storm to the violent explosion of his mind in poetry.) In short, we have only to reconstruct the Romantics' metaphysics and theory of poetry to understand this line; Byron uses frown as if those theories constituted a part of the system of normal beliefs.

For all metaphors, the process is the same: to identify the referent, we construct the most plausible belief-system against which the use would be rational. The process is not always simple, of course, We can speak of the Romantic world-view with any confidence only because

the Romantics were at such pains to explain themselves, and because it is evident in such a large body of accessible work. Where a use is licensed by beliefs that are more idiosyncratic, or "stranger," than these, the meaning may be harder to recover.¹⁶

In fact, a theory of language has little to contribute to the understanding of metaphor. It cannot tell us anything about why some experiences (such as being greeted) should be perceived on the model of others (such as being warm). Nor can it tell us on what basis we attribute something to the system of normal beliefs, or why we are willing to be more generous in interpreting some kinds of discourse than others. Above all, it cannot supply the motive for metaphor. Still, the approach suggested here may have some value, not because it offers any new theory of metaphor, but because it makes unnecessary the old ones, with all their talk of "vehicles" and "tenor," or "implied analogies." In this way, it makes the aesthetic problem a bit easier to tackle.

Earlier on, I suggested that metaphor was simply a particular kind of conversational implicature, involving word-uses rather than sentence-uses. (Equivalently, we could say that with metaphors, we must work to recover the propositional content of a sentence, and with implicatures, of a discourse.) This analogy might be pushed in several useful ways. First, we could say that with implicatures,

as with metaphors, we discover the "conveyed" interpretation of a sentence token S by constructing the belief system against which S could be used to convey a relevant proposition p, given the context. (Drawing the analogy more closely, at some metaphysical peril, we could say that we understand an implicature by determining the beliefs that license the use of S to refer to a proposition p in the normal range of reference established by the application of the maxim of relation--"Be relevant"--in that context.) We may expect, then, that the two phenomena involve the same aesthetic process.

Our talk of "constructing a system of beliefs" to interpret metaphor must also recall linguistic discussion of pragmatic and semantic presuppositions, which is connected to discussions of implicature in an obvious way. While I don't think that "presupposition" is a very useful notion (we have trouble enough with "meaning") it is worth noting that some of the phenomena that might be called "presupposition violations" may be treated as a special class of metaphors in which the abnormally used expressions are grammatical functions or syntactic constructions, rather than lexical items. The technique of jokes makes frequent use of such devices, as when someone says of God, "She's Black" or when someone writes, "Galbraith has predicted nine out of the last three recessions." In either case, it can be argued, we recover the intended meaning by

constructing a system of beliefs against which the use would be rational ("God is a woman," for example) and the aesthetic effect arises out of the juxtaposition of that system with our system of normal beliefs.

4.3.2 Metaphor and Slang

I want to make one other observation about metaphor that is relevant to the aesthetic problems it raises. In most discussions of metaphor, as in this one, there is a presumption that the motive for using language in an abnormal way is primarily aesthetic or affective. (This tradition begins with the earliest Western discussion of metaphor, in the Poetics.) But there is a large class of uses that qualify as metaphors by most definitions--including ours--which reflect primarily the efforts of speakers to be sociable to one another. And these uses should be of special interest to linguists, because it is out of them that new words (such as swing) most frequently arise.

Recall that Weinreich, Labov and Herzog 1968 defined the actuation problem as the problem of determining why a particular change takes place in a given language at a given time. We have already addressed this problem with respect to certain kinds of change in meaning and use, but we also noted that many such changes could not be explained solely in terms of the purely communicative function of language. In 4.2.3, we talked about cases like (paper)

plane, where a use-change results from a change in normal beliefs that affects the cue-validity of the functions that connect the uses of words. But new word-uses may also arise that serve no direct communicative ends, which may even serve to reduce the communicative efficiency of the linguistic system.

Examples of this sort of development are particularly abundant in the classes of word-uses that we label as "slang" or "colloquial." And such uses are quite often metaphorical, by the criteria we gave above. Consider, for example, the words that are currently used to refer to marijuana. (I hope I may be forgiven an excursion into what can only be called pop philology. We will want to be able to appeal, in the course of discussion, to some fairly fine differences in sensibility, which would be inaccessible for a more remote state of the language.) The words grass, weed, herb, smoke, tea, and dope have each of them an obvious metaphorical origin; only pot is obscure and arbitrary.¹⁷ It will be instructive to ask after the purpose of these metaphors, and how this purpose is connected to social evaluations attached to them.

The nature of usage categories like "slang" has been wholly ignored in recent work of lexical semantics. This is perhaps understandable, since it is generally assumed that such evaluations do not affect the determination of truth-conditions. (Though we cannot explain the neglect of

the topic by sociolinguists in this way.) Most linguists would take a line on slang roughly as follows. Some words and expressions--which tend to be particularly "colorful"--are appropriately used in certain kinds of social contexts, and must have various use-conditions attached to their lexical entries. Slang words, for example, are appropriately used only in informal exchanges; other words are used only in writing, or in talking with children, and so forth. But to use such words inappropriately is only to commit a social gaffe; communication is not affected.

This formulation raises several questions. First, talk of "use-conditions" obscures more than it explains. How are such conditions entered? (Consider the use-condition that we would enter with obscene words: "Don't.") More important, how are the conditions learned? That is, how do we know that a word is slang when we hear it? We could say that we know that certain words are slang because we do not encounter them in written texts, or in formal speaking, but this sort of argument has gotten us into trouble before, with metaphors and passivized idioms. (Besides, we do encounter slang words in formal contexts, and recognize their deviance.) We also run into trouble when we try to explain the assignment of the slang evaluation on the basis of the contexts in which the word is used. Sometimes, it is true, a word is associated with a certain subgroup within the speech-community, so that its

use by other speakers is affected by the social evaluation of that subgroup. Terms from Black English, like jive, are certainly evaluated in this way. But many slang expressions are used by all members of the community, such as kick the bucket, or bushed for "tired." In any event, the fact that a word is not appropriate in certain contexts can hardly be learned from observing that it is appropriate in others. Finally, we cannot reasonably suppose that words are slang by stipulation. Occasionally, the prescriptive grammarians may get their hands on a word like ain't, but no rhetoric could even keep up with slang. In most cases, there are no institutionalized interdictions against the formal use of a given expression. Why then is the use of get one's hands on inappropriate in written English, while keep up with is less so?

The answer to this question is tied to the answer to another: why should use-conditions exist at all? The standard view is modeled after examples like pickpockets' jargon, which allows a pickpocket to communicate with his confederate without alarming his victim. In an analogous way, it is argued, the slang of teenagers or jazz buffs provides a semi-private language, to which the uninitiated have no access, and in using slang, a speaker can explicitly affirm the special relationship that holds between him and the hearer. There is certainly a measure of truth in this, at least for some of the categories that would standardly be

handled by means of usage-conditions. (This kind of account can't begin to explain the existence of obscenities, but neither can any other of which I'm aware.) But there are problems here, as well. What will we say about slang expressions that are used by the entire community? We could say that they are used to indicate that an informal mood is presupposed. But this takes us back to our earlier question--how could we come to learn that an expression had just this function?

A third problem is related to the first two. A number of usage categories have been introduced by lexicographers (the only people who continue to concern themselves with the question); if we confine ourselves only to the evaluations that are assigned to words that are universally used (avoiding "jargon" and "cant," for example), we are left with "slang," "colloquial," "informal," "idiomatic," "vulgar," "obscene," "precious," "refined," "casual," and "intimate," among others. But at best, these categories can only be rough indicators of usage. And even using all of them--which no dictionary does--we can't begin to capture the variety of usage-conditions that can be associated with expressions. For example, how shall we describe the difference in the situations in which we use gripe, bitch, or shoot one's mouth off, or get going, move on, split, and take a powder?

We can't address most of these problems here; I

raise them primarily to show that evaluations like "slang" are worthy of much more attention than they have received. But we can at least suggest a line of attack. Rather than talking about the categories to which uses are assigned, let's talk about the uses themselves. We said earlier that the uses of grass, tea, etc. to refer to marijuana had "metaphorical origins." But why shouldn't we just say that they are metaphors? It is true that they are regular; speakers do not coin them anew each time they are used. But regularity doesn't guarantee conventionalization, as we have seen. We can assume that most speakers would be able to interpret these uses on first hearing.

We could divide the slang words used to refer to marijuana into several classes, according to when they were introduced. Tea, pot, reefer, maryjane and possibly weed were in use well before 1960; grass was introduced around 1960, and herb, dope, and smoke were introduced in the late 1960s or early 1970s. But this chronology is a bit misleading, because there has been considerable variation in use from group to group and place to place. A particular term may have been introduced on the West Coast well before it was used in the East, or vice-versa. Some terms were restricted only to certain groups; thus charge, while used in England and among American Blacks, has never had any general currency. Terms have also varied in evaluation; tea, for example, is usually felt to be more precious than grass.

Moreover, the community of speakers who use the terms has changed enormously in size and composition over the last twenty years, so that even those terms that were available in the 1950s to a small collectivity have changed in distribution since then. Things are worse confounded by the absence of citations for most uses.

We can still make some useful generalizations, however. Of the terms used before 1960, only tea is obviously metaphorical. (Though we might regard maryjane as a kind of "formal" metaphor; see below.) A speaker can know what drug pot is used to refer to only by being initiate to the convention that governs that use. The analogy to a code is perfectly apt; someone who has never heard the use could not understand it. Code-words like these are especially characteristic of certain varieties of local slang and argot. We already mentioned the example of pickpockets' argot; we could add tinkers' cant, or the slang of schoolboys and teenagers. In each case, members of the relevant collectivity can assume a common acquaintance (as opposed to just a common set of attitudes and beliefs), and the criteria for membership in the group are clearly defined. Thus it is highly likely that all and only the members of the group will have heard a given word-use. (It might appear odd to say that teenagers and schoolboys are all connected by a common acquaintance. In general, however, communication between members of geographically separated

subgroups does not play a role.)

By contrast, it is not unreasonable that a speaker familiar with the referent should be able to understand a use of tea, grass, weed or herb to refer to marijuana. In using a metaphor, the speaker presumes a common set of beliefs, rather than a common acquaintance. These particular uses would be rational only in contexts in which it was presupposed that marijuana was best identified on the basis of its appearance; in referring to marijuana with grass, a speaker behaves as if the proposition "Marijuana is brownish-green and flaky" could be assumed as part of the conversational background. Note that such metaphors function less to keep the uninitiated from understanding, than to indicate a certain rapport. By the time that grass became current among a large sub-section of American youth, a much larger proportion of the community was familiar with the appearance of marijuana, and a given token of the use might have been understandable to an older listener-on. But older members of the general community could not have used the word among themselves, since there could be no presumption that the appearance of marijuana was a matter of common knowledge. Thus, the use created its own speech-community: just the group of speakers who could identify one another as likely to be familiar with the appearance of the drug.

The uses of tea, grass, weed and herb are not

equivalent, of course. Tea preceded the others by some time, and is now accounted either precious or démodé. (A New York musician has told me that tea was "old-fashioned" as early as 1955.) The words differ in several ways. First, tea, unlike the others, designates a category which might itself be a candidate for reference in many of the contexts in which it could be used to refer to marijuana, and there will rarely be anything in the linguistic context alone which would resolve an indeterminacy of reference. Thus "Would you like some tea," or "Bring some tea along" are not licensed in all contexts in which familiarity with the appearance of marijuana is presupposed; they require as well a set of mutual expectations as to which commodities are likely to be offered, and are hence more likely to appear innocent to the non-initiate listener-on. Hence the furtiveness of this use, and hence its rejection by the community of marijuana users of the 1960s, who looked with disdain on the ritual and paranoia that surrounded its use in other eras. The use of tea presupposed that marijuana use was a guilty secret.¹⁸

We might note that tea seems to be enjoying a renaissance (a recent issue of the trade paper High Times advertises a plastic container as "the perfect cache for your choice Colombian tea"). This follows from the fact that the use is now perceived as so absurdly furtive that a younger speaker can now use it in the assurance that he

will be credited with an ironic intent. If the use of tea is now a little precious, it is because it presupposes something like, "We both know that we are not the sort of people who are furtive about our use of drugs, don't we?" Again, the metaphor creates its own evaluation.

Grass, weed, and herb differ from tea in that their use to refer to marijuana does not invite misconstrual by non-initiates, but they differ among themselves as well. Both grass and weed have been around for some time, but weed has remained a local use, while grass has become, with pot, the normal slang term for marijuana, and appears, for example, in newspaper headlines. (I have seen herb in print, and am told ~~by~~ undergraduates that the word is sometimes used in conversation, but I have not heard it, and will have nothing to say about it here.) Weed differs from grass, herb, and tea in several ways. First, there is not really a good function from the appearance of weeds to the appearance of marijuana; nor do we customarily see weeds chopped up or dried, as we do the other kinds of plants. Depending on speaker beliefs, either of two RF's from weeds to marijuana could be derived. In the thirties, "the weed" was used to refer to marijuana against the belief, "Marijuana is a harmful plant." More recently, "weed" has been used against the common assumption, "Marijuana grows wild everywhere." But this use also appeals to the beliefs that license the earlier one; speakers are obviously not unaware

of the social evaluation attached to weeds (cf. the use of weed for "fool, jerk"). Whatever the motivation for making ironic reference to the harmful properties of the drug, the move is common among users of all intoxicants. Thus a marijuana user may describe himself as being "destroyed," "wasted," "ruined," "wrecked," or "blown away." (Hunter S. Thompson has built an entire journalistic genre around this conceit.) "Wrecked" and "ruined" are used to refer to drunkenness, as well; so are "blasted," "blotto," and "shit-faced."

We might note an interesting difference between the use of weed in the 1930s and its current use. The earlier use was generally with the definite article, and referred to the plant cannabis; the later use is as a mass term, to refer to the substance derived from cannabis. The irony of the later use is relevant here, as well. Note that the mass use involves a composite function, from weeds to cannabis, and from cannabis to its characteristic substance. (The composite could not be computed in the other order, since the belief that a single substance is characteristic of all weeds is both abnormal, and irrelevant to this use.) It thus presupposes more than the use with the; it must be a part of the conversational background both that cannabis is best identified in terms of its source and harmful effects, and that it is reduced to a "mass" substance before consumption. (If the drug were ingested by eating its fruit, we

would not say, "Do you have any weed?" anymore than we say, "Do you have any nut?") So the irony arises in part because the use presupposes both the belief that marijuana is invidious, and the knowledge of how it is prepared, a disjunction that could not be realized in most contexts, given the distribution of beliefs in our culture.

The same sensibility that licenses weed licenses dope, and both words have enjoyed a recrudescant popularity in the seventies. Dope, of course, is the only one of these terms to have been used to refer to other drugs. (The OED Supplement evaluates it as "slang" in its use to refer to opiates; the first citation is American, for 1889. But it was used in standard written English, as the 1922 citation from Public Opinion attests: "The cabarets in Constantinople are a meeting place for all of the world's dope purveyors.") Dope was used in the 1930s to refer to marijuana, under the assumption that marijuana was itself an opiate--what we may call "Donnellan-reference." (It has been becoming obsolete as the use of opium declined, and as the normal category of "dangerous drugs" was redefined to include a variety of biologically unrelated substances, for which drug was used as the superordinate. (Thus "drug addict" replaced "dope addict," etc.)

When a contemporary speaker refers to marijuana as "dope," he invokes as background the beliefs that licensed its use in the 1930s. (This use, and the beliefs that

surrounded it, have become widely known in consequence of the popularity of old anti-marijuana movies and posters, as part of a general nostalgie de Comstock. See, e.g., figure 1 above.) In so doing, the speaker denies the earnest efforts of marijuana users of the 1960s to draw a sharp line between marijuana and "hard" drugs, and thus drives a wedge between the speaker and hearer and the liberal community at large, which has since accepted the distinction. Like tea, moreover, this use invites misconstruction, but aggressively; here, the non-initiate listener-on will understand the reference as to a more dangerous drug. For this reason, the use is not available in contexts in which rejection of the beliefs that license it is not presupposed. A speaker can't use dope in this way unless he is certain that his listener will be able to apprehend his ironic intent. Accordingly, the use is perceived as unavailable to members of sub-communities in which there is a good deal of individual variation in the evaluation of drug-use--say, the (perceived) class of "defensive middle-aged liberals" at whom the irony is directed in the first place. (Of course we may expect further changes in the use of dope to refer to opiates, so that its use to refer to marijuana must soon be re-evaluated.)

I don't expect that the reader will have agreed with all of the evaluations of uses that have been assumed here, or with the evaluations of social attitudes that were

invoked in explanation of them. The possibilities for variation are pretty much uncontrollable when we deal with examples like these; we can expect speakers to disagree over the degree to which a given use is motivated at all, and over the beliefs that other speakers will bring to bear in interpreting these metaphors. We can assume, for example, that use of grass have become in some measure conventionalized for most speakers, but we are likely to disagree about how much, and for whom. (And it doesn't matter what the facts really are, since it is our perception of them that determines how we will use the terms.) It is just this sort of disagreement, in fact, that makes the extension of labels like "slang" so difficult to determine, even in areas that are less volatile than this one.

But the existence of variation doesn't undermine our basic point, which is that the use-conditions attached to these metaphors are determined by speaker assumptions about how other speakers will perceive their rationales, and that these conditions are thus far more precise than could ever be indicated by the introduction of arbitrary markers of usage categories. In fact, we may suspect that we do not require nearly so many such categories in lexical description, and that the fact that many word-uses are appropriate only to a certain set of contexts can be independently explained. Not all slang words have metaphorical motivations, of course. In some cases--pot and reefer for

examples--the relation of the form to the referent is entirely conventional, and usage could be determined only by association of the use with a particular population. In other cases, usage may be determined in part by phonetic and phonaesthetic factors, as when someone says, "These houses are scrungy." A particularly interesting set of uses involve what we may call "formal metaphors." The use of maryjane for marijuana is an example; it was licensed, during the 1930s, by the assumption that the word marijuana was sufficiently unfamiliar as to be unrecoverable when translated. (But all translation is metaphor, it has been said.)

What we have said here about slang metaphors is particularly relevant to the examples of phrasal idioms that we talked about in the last section, since such expressions are invariably metaphorical in character. We saw there how assumptions about the beliefs that licensed the use of the idioms determined some aspects of their linguistic distribution; we could apply to them the kinds of arguments we have used in this chapter to show how their usage evaluations are determined as well. In this way we could hope to explain why kick the bucket, get one's hands on, and blow one's stack are evaluated as more "informal" than the roughly synonymous pass on, get hold of, and read the riot act. But this would be a complicated undertaking, since with phrasal idioms of this type we have to take several considerations

into account--not only the beliefs that license the use of each component, but also the beliefs that license the use of one relation to refer metaphorically to another, in the case of abnormally decomposable idioms like make tracks and give someone the gate. This is particularly difficult, since we do not have a way of showing how all of these factors conspire to determine evaluation. Moreover, it would be harder to link the usage of these expressions to differences in social evaluations. Expressions like get one's hands on have approximately the same usage conditions for most members of the community, and their status as "slang" is due to the evaluation that we consistently assign to a certain kind of behavior.

4.4 Notes from All Over

Obviously, we have not said all that there is to say about polysemy, syntactic identity, ostension, meaning, acceptability, idioms, metaphor, slang, and the do so construction. (Things might have gotten even more out of hand, if reason and Terry Langendoen hadn't conspired to hold the line.) And equally obviously, we could have departed from this particular itinerary at any of a number of points, and moved on to a wholly different series of topics, for surely there could be no theory just of the phenomena we have talked about here. Here are some unconnected notes and observations from the thesis I originally intended to write, before I got bogged down in the lexicon.

My interest in the problem of polysemy grew out of a paper I wrote with Chiahua Pan 1975 on the subject of generic sentences. In the course of looking at English indefinite NP's, I became interested in the analysis of the indefinite article, and hence in the mass/count distinction; this led naturally to a consideration of the question of how a word like chicken could be used to refer to a bird on one occasion, and a substance on another. I have not recovered from this distraction, and still do not have a satisfactory analysis of the English articles to offer; but the articles do provide a useful jumping-off point for discussion of some other semantic problems. For example, Pan and I were primarily concerned to give an account of the truth-conditions that attach to sentences (more fool we) like 1-4:

1. A symphony has four parts.
2. A beaver builds dams.
3. A baby-sitter gets \$2.00 an hour around here.
4. A Cadillac is expensive.

As has often been noted, the subject NP's of these sentences are not equivalent to universally quantified NP's; 1 is not falsified by the existence of Schubert's Eighth Symphony, nor 4 by the existence of one--or even a fleet--of cheap Cadillacs. But by the same token, the introduction of "generic" or "characteristic" quantifiers seems unrevealing, and doesn't help in discriminating the

indefinite singular from the indefinite plural (beavers build dams) or the definite generic (The beaver builds dams). What we finally came up with was this: a statement of the form "the $\emptyset$ is ψ " is true just in case every $\emptyset$ that is not ψ is not ψ for reasons that have to do with its own particular history. Thus, it could be the case that one beaver doesn't build dams because his teeth have fallen out, another because he is in a zoo, and so on, and the statement "a beaver builds dams" would still be licensed. This characterization works, and can be generalized to the indefinite plural generic, but at the time it seemed capricious, and we were not entirely successful in tying the interpretation to the indefinite morphology. In writing section 4.1, however, it occurred to me that another way of characterizing the interpretation of l was by saying that the referent of the subject was something like "any normal beaver," in our psense of "normal"--that is, a beaver about which we could assume nothing more than was given to us by culturally determined beliefs, an indefinite, unindexicalized, null-context beaver. "Unindexicalized" is the key here. It is the absence of the, not the presence of a, that makes the non-specific interpretation available. Note that mass terms can be used "generically" too, and that the truth-conditions are the same. Under the normal interpretation, the statement "Steel is expensive" is true just if any instance of steel that is cheap is so as a result of circumstances

particular to it.) Why do non-indexicalized NP's allow both interpretations, while definites don't? (The "generic" definite, as in "The beaver buidls dams," is a different bird altogether; the NP here refers to a type, not a non-specific token. Cf. Lawyer 1974, who cites examples like "The [*a] beaver is becoming extinct.")

In giving a proper analysis of the English definite, we would buy ourselves an understanding of the problems of syntactic identity that we talked about earlier. For the patterns with wh-words and definite pronouns, as in:

5. John tore up a newspaper, because Hearst had bought it the week before.
6. John tore up a newspaper, which Hearst had bought the week before.
7. John tore up a newspaper. Hearst had bought the newspaper a week before.

Earlier we suggested that the condition of linguistic identity, rather than the stronger identity of reference, should be sufficient to handle the problems raised by examples like 5 and 6. But why any "identity" conditions at all? The use of the definite in 7 must follow from its function; you use the when the value of the RF is computable at the argument picked out by the "indexing gesture" ([__Is], [__æt], [nyuspepr]), a determination that can't be made on the basis of linguistic information alone. By the same token, we could say that you use wh-words to indicate that the value of a

variable can be computed as the value of an RF at the referent of some other term within the sentence (or to ask for an indication of some argument from which the value of a variable can be computed). Which isn't the same as "binding" the variable at all: you can determine identity of a referent without being told what it is identical to. So we can ask the grammar to tell us how to use wh-words, but not when, and the "logical form" of a sentence must at every point contain indices for the functions that connect its terms. (Which is not to say that the grammar cannot also intrude to tell us under what configurations a particular variable can be "connected" to a particular term.)

A proper understanding of the articles is also a precondition for the univocal analysis of the copula that we mentioned above. The idea is this: all of what appear to be instances of the "be of predication" are really instances of the "identity be," where some function other than the identity is operating on the subject term. Thus the interpretation of "John is twenty-one" can be given as "Age of (John) = twenty one," and of "This is red," as "Color of (this) = red," and so on. So both subject and "predicate" are referring terms. But what about examples like "This is heavy"? If heavy can refer in predicate position, why not also in subject position; why don't we also get, "?Heavy is the weight of my true love's hare?"

A possible line of argument: the meaning of heavy is such that it can be used to refer to something only when the range of reference has been localized relative to the universe that surrounds the referent of some other term. Which is only to say, you can use heavy to refer to a certain range of a particular scale--"heavy for a person"--but not to a range on a general scale of weight.

Now let's make another assumption, about the "meaning" of word-order positions, following a line suggested by Keenan 1975. If be is truly the identity, then statements of the form "x is y" should always be equivalent to statements of the form "y is x." And this is wrong, we know, or else we would not be able to attribute two interpretations to a particular utterance of "A mighty fortress is Our God" (for one of which the last term must be written with lower case letters). Then either be is assymetric, or the assymetry is due to some difference in the "rhetorical" functions associated with subject and object positions. Rather than dragging in "focus" or "topic" (recall the problems we ran into in talking about passives in 4.2), let's say this: the range of reference of the object term in copular sentences must always be determined relative to the universe of discourse surrounding the referent of a subject term. Thus, in processing a sentence like "x is y," we must first compute the referent of "x" in context, and then compute the reference of "y" relative to it. Word-order signals the relative

scope of the domains from which the values of indexicals are to be assigned.

But it would require another thesis to work this out in detail; we would have to be able to use this principle to account for all the differences between the pairs in 8 and 9, as well:

- 8a. Editors are neurotics.
- b. Neurotics are editors.

- 9a. A friend of John's is a friend of mine.
- b. A friend of mine is a friend of John's.

And then, having disposed in the course of things with the difference between specific and non-specific indefinites, we could move on to talk about "Many arrows. . . ."

Suppose that heavy does refer in predicate position, then. What about its use in NP's like "a heavy chair"? There, we will say, it is used to indicate (or refer to, if you like) the RF, or one of the elements of a composite RF. "Modification" is required, not only when the value of RF is a member of a proper subset of the set of things to which the head can in this context be used to refer (as in "the red house"), but also when the RF takes us from sets of things to their simulacra (fake, counterfeit), or to sets of things that have been explicitly announced to be members of such-and-such a set (soi-disant, alleged, recognized). By distinguishing the uses of terms in adjective and predicate position, we may also be able to explain the properties of

"pseudo-adjectives," as in peace conference, executive secretary, and go-go girl. (Can we bring in scope here, too, to explain the difference between pre- and post-nominal position of adjectives, as in the chosen people/the people chosen?)

Would a univocal analysis of the copula lead us to still another indeterminacy? Consider a sentence like "She's not that kind of girl." We could analyze it in two ways, either as asserting the identity of a type with a type, or a token with a token; either interpretation will entail the other. Even in context, then, we may be unable to determine the referent of a term, beyond what is necessary to recover the interpretation of an utterance. This again, quite apart from Quine's qualms.

Another tack we might have taken: our discussion of the "use/mention" function has obvious relevance for the analysis of verbs designating propositional attitudes, and of complementation in general. Whatever sentences are used to refer to in assertions (we danced around this problem in 4.3), that is certainly not the only thing they can be used to refer to, and on occasion they can be used to refer to facts, utterances, propositions, or even people, as in "What did old 'I am not a crook' have to say for himself this evening?"

Not only "content words" are polysemous; we could have moved from newspaper to modals, aspectuals, cases, prepositions, conjunctions, tenses, or quantifiers, all of which routinely show multiple uses, connected by functions that are conventionalized to varying degrees. The semantics of such items has intrigued linguists for a long time, of course, and this framework doesn't really help when it comes to the central problem of explicating the meta-language in which we can talk about the meaning of the should or the Latin imperfect, since it must invariably make use of "basic" concepts that are accessible only through investigation of natural language, or of language acquisition. But we can make some questions a little easier to ask, such as, "Is the so of do so the same as pro-adverb so?", where linguists have heretofore been forced to cast their lot one way or the other, or to talk vaguely of "blends."

A note that I should have stuck somewhere back in Chapter Two, on the polysemy of "function." The difference between the "relations" with which linguists and psychologists have tried to express the connection between the psenses of polysemous words, and the functions that we found much more useful, is just the difference between

structuralist and functionalist accounts of meaning and use. Relations don't take you anywhere, and functions do. This is the best I can do by way of Conclusion.

FOOTNOTES--CHAPTER FOUR

1. An example much closer to home. It is a matter of fact that the only syntactic theory with which any linguist can presume another linguist to be acquainted, in the absence of further contextual knowledge, is the Aspects model. Accordingly, I have talked here about constructions like passive as if they were transformationally derived. And since I am obviously unconcerned here about syntactic theory, the intelligent reader will have interpreted such references non-tendentiously, and will not have attributed to me the Aspects position. Rather, he will have reasoned, "Well he needs some way of talking about passive, and that is the most accessible." Whereas if I had referred to passive in the terms made available by relational grammar, there would be a greater temptation to attribute to me that theoretical stance, since such use would not be justified by expectations about what theories the general linguistic reader is familiar with.
2. A brief polemic. Most linguists--even those who call themselves sociolinguists--have tended to consider prescriptivism only as a pernicious doctrine, based on fundamental misconceptions about the nature of language, which serves only to reinforce class distinctions (and confuse speaker judgments, in the process). (Cf. the abuse heaped on Bernstein for his earlier, somewhat naive formulations of the notions of "restricted" and "elaborated" codes.)

But this position is itself extraordinarily ingenuous; it is as if an anthropologist decided that primogeniture was a pernicious doctrine fostered by basic misconceptions about the nature of the family. In particular, as we come to lay fewer and fewer regularities in linguistic behavior to the structure of language, and more and more to the strategies whereby speakers exploit such structures, questions about differences in the linguistic practices of different sub-groups within the community assume a correspondingly greater importance and interest. There is every reason to believe that speakers who live in a community in which there is a rich diversity of belief-systems should become more skillful at certain kinds of language-use than those who live in a community which is relatively homogeneous, and even that these skills may enable the former to bend language to aesthetic and moral purposes in a way that is different in degree, if not in kind, from what the latter can do with it. Then if the culture places a high value on these specialized uses, and if they can be

mastered only through acquisition of the linguistic self-consciousness that must accompany the presence of certain kinds of belief-heterogeneities, linguistic prescriptivism must perform a vital service in maintaining notions of culture (with a "k") and civilization, just as the educators have been arguing all along. (Whether civilization itself has any intrinsic value, however, is a question that falls outside the scope of this inquiry.)

3. This discussion is based on Lewis's 1969 formulation. The 1976 formulation differs in important details, but not so as to concern us here.
4. One could say that handshaking is really just a linguistic convention. But what conventional behavior could not be said to have "meaning?" I confess that I don't like it when semioticians talk about the "meaning" of wearing a tie, or of a certain form of strip-tease ("By God," as Dwight Macdonald said on the use of contact as a verb, "let's hold the line on this one.") But they're probably right.
5. Lewis actually begins his discussion of convention by talking about "common knowledge," but ultimately rejects this as a basis for his definition, for reasons that he acknowledges have little practical importance.
6. The notion of "authority" has been problematic since Weber first introduced it into sociological theory. We might better talk about "perceived authority"; in that way, we can avoid a lot of difficulties, just as the notion of "normal belief" lets us finesse the problems raised by the "collective awareness."
7. I should note here that our remarks on the indeterminacy of meaning in Chapter Three appear to force us to revise the common-sense definition of "same lexical item" that we earlier argued was preserved by a pragmatic treatment of polysemy. In Chapter Two, we said that two uses of an item were "the same" for syntactic purposes just in case they had the same meaning. Let us now say that two uses R and R' of an item are the same for syntactic purposes just in case there exists some regularity R" in the use of that item such that R and R' could be predicted, given R", solely on the basis of the system of normal beliefs. Then R" could be either R or R', or some other regularity. If the meaning of the item is indeterminate, R" could be any of several regularities, which our formulation now allows for.

8. Katz's proposal will work only if there are no pairs of transformations T_i and T_j such that both might apply to a given constituent C of an idiom I , but only one of them yields a well-formed output. Katz maintains that he has been able to find no relevant examples, but relativization, which neither Katz nor Fraser considers, may create some problems:

- i. Headway has been made by all parties.
- ii. We were not surprised at the headway John made.
- iii. We won't be happy until a clean breast has been made of the entire affair.
- iv. *We were surprised at the clean breast made of things.
- v. *Tracks were made by all of the culprits.
- vi. *We were surprised at the tracks they made after the burglary.

(Examples after Brame.) But there may be disagreement about these judgments, and the way in which relativization is stated is sure to be crucial. In any case, we have bigger fish to fry.

9. This is a general objection to the introduction of negative exceptions in syntax; see below, where we argue that there cannot be arbitrary restrictions as to which verbs undergo passive.
10. Equivalently, we could express the interpretations of VP's as predicate expressions in Church's Lamda calculus, as suggested by Sag 1976.
11. As Katz has noted, judgments about sentences by these tend to be quite labile. Moreover, it has often been noted that a passive sentence which sounds odd in isolation may be improved if stressed in a certain way, or if additional material is added; see below. Since we are arguing that all of these sentences are generated by the grammar, we will not have to take a difference in judgments to reflect a difference in "rule-dialects"; speakers may be responding to ordinary differences in the assumptions they make about what sort of normal beliefs would be taken to license a use. Accordingly, it will suffice to note that some sentences are worse than others, whatever their absolute status. Clearly, 28 and 29 are less acceptable than examples like Tabs were kept on John, or The buck's been passed a little too often around here.
12. Such examples are less clearly acceptable where both uses are restricted to animates. The problems with

John married Mary right near where the preacher had ten years before under the crossed interpretation seems to be that it is very hard to construct a purely linguistic context in which a reading in which the RF is the identity is not plausible. I would scarcely want to claim that this was evidence for some sort of trans-derivational constraint; to the contrary, I think it suggests a line of argument whereby a number of phenomena that have been used in support of such constraints may be pragmatically explained. In fact, I think the identity condition as we formulated it in 3.1 may be invoked to explain a number of problematic aspects of anaphoric processes, but we will not have the opportunity to pursue this matter here.

13. Occasionally, it is true, we may be able to draw interesting generalizations about the way in which certain areas of the lexicon evolve. (Cf. Stern's 1931 discussion of the development of Middle English time adverbs, or Traugott's work on the way in which spatial terms develop temporal uses.) And there may be interesting parallels in lexical semantics to the other problem that WLH allude to under "constraints"--that of saying what possible form any innovation can take, which is equivalent to the problem of adducing linguistic universals. The notion "possible lexical item" has figured in linguistic discussion in several ways lately.
14. After 1940, we no longer encounter the expression "swing music." Swing thus patterns with bop, ragtime, R&B and bluegrass, as opposed to classical, country, rock, and disco, where we can use only the latter as adjectives with music. I used to have a theory about this, but I can't remember it now.
15. The present definition would also include some kinds of uses that were called metaphors by Aristotle, but which would probably not now be so regarded, such as when we say "Odysseus wrought ten thousand noble deeds," where "ten thousand, which is a particularly large number, is made to stand for the idea of large number in general." (trans. Wheelwright 1935)
16. As, for example, with Dylan Thomas's "A grief ago I saw him there," which is by way of becoming a locus classicus for recent discussions of metaphor; see, e.g., Ziff 1964, Putham 1961, Levin 1971, and others, all of whom have approached it with the following two assumptions:

- a. Ago "selects" terms that designate units of time, like hour, second, and month.
- b. Grief is not such a word, but it is interpreted here on the model of such words; i.e., moment, hour, etc.

(Putnam goes so far as to equate the statement that grief is here used as if it were a measure of time with the "normal, informal explanation of the line.")

Let's start with the first assumption. Ago is used to indicate the distance from the moment of speaking of a past event; it thus specifies both the point from which the measurement is made, and the direction in which it is to be counted off; it is the temporal analogue of an expression like "South of here." (Like yesterday, and unlike before, earlier, and soon, ago is a "pure" indexical; its value must be assigned relative to the moment of speaking):

- i. I got there January 10, 1957, and Peter had arrived two weeks earlier (two weeks before, two weeks ago, yesterday).

And as with all measure phrases, the NP used with ago must be indefinite:

- ii. He arrived a/*the week ago. (two weeks ago, *those two weeks ago).
- iii. The house is (*the) two miles South of here.
- iv. It is (*the) two feet long.

This much is given by the "meaning" of ago; since it is a "one-criterion" word, we may as well ascribe these properties to semantics, and say that ago carries a condition of reference which requires that its associated NP must refer to a measure of time. But words like hour and second are not the only words that can be so used. We can normally refer to intervals using words that designate periodic events, as in:

- v. He gave the same speech two meetings ago.
- vi. We had a blizzard two Christmases ago.
- vii. We last saw them three summers ago.

(Summer is here used, of course, not to refer to the duration of a summer, but to the interval between

summers.) We can also refer to intervals using words that designate things that appear periodically, as in:

viii. Three beers ago, I could have recited the whole poem.

ix. We ran a story on him two issues ago.

Nor is there anything particularly odd about referring to intervals with words that designate "inner" events, as in:

x. I tried that medication two headaches (hang-overs) ago.

xi. I thought the pain was getting better a couple of twinges (stitches) ago.

xii. That occurred to me a couple of thoughts ago.

xiii. Five New Years' resolutions (decisions) ago, I decided to give up smoking.

(If x-xiii are not all entirely normal outside of context, neither are they noticeably "literary" or metaphorical, as a grief ago is.)

The oddness of a grief ago, then, has nothing to do with the use of grief to refer to a measure of time; rather, it is odd to use grief to refer to an event at all. We can talk about "a pain" as an occurrence of the sensation of pain, but we don't use "a grief" in this way. ("A grief" can be used to refer to a type of grief, as "a (pure) gold" or "a (great) happiness" can be used to refer to types.) Thus the oddness of grief in:

xiv. I had a sudden pain (sensation of sickness, ?grief?ache).

xv. The third pain (sensation of sickness, ?grief ?ache) is the most intense.

This can be explained, I think, by reference to our normal theories of pain and grief, but it is obviously impossible for us to go any further than that here. It would appear that we do not regard grief as a sensation, but as an enduring state--an ache, not a pain--but we could make sense of these notions only in terms of the vast literature on the question of how we perceive our inner states.

Nor, for that matter, can we assign an interpretation to the line with any certainty. Byron's use of

frown was licensed by a view of nature that is common to most of the literature of the period, but Thomas's use of grief requires that we reconstruct a belief-system that is peculiar to his poetry. I should note, however, that it was a particular mishugas of Thomas to talk about his inner states as if they were generated by natural law:

The force that drives the green fuse through
the flower
Drives my green age . . .

So that the advent of grief is not materially different from the advent of summer or frost. But I won't try to defend this reading, nor does it matter. I am mainly interested in showing how foolish it is to suppose that a linguistic theory, or a theory of interpretation modeled after linguistic theory, should be able to tell us when a word-use is deviant, or under what circumstances a deviant use may have an interpretation in context, much less assign the interpretation.

17. I have seen pot linked to Mexican Spanish potaguaya, but this may be fanciful; I have also seen it linked to Southern U.S. pot, "moonshine." In any event, the use has no synchronic motivation, real or apparent.
18. In this connection, a personal anecdote: In my freshman year at college--this was 1962--I ran into a high-school friend who was leaving a party with a group of people. I asked her where they were going, and she said, "We're going sailing." When I asked in all innocence if I could come along, she turned to her friend, who said, "Well, it's sort of a small boat." I have never heard this use anywhere else, and it is even more coy than tea, since it presumes mutual familiarity with the sensations that marijuana induces. But the principle is roughly the same.

BIBLIOGRAPHY

- Algeo, J. (1973). On Defining the Proper Name. Gainesville: University of Florida Press.
- Alston, W. P. (1971). "How Does One Tell Whether a Word Has One, Several, or Many Senses?," in Semantics. Cambridge: Cambridge University Press.
- Aronoff, M. (1976). Word Formation in Generative Grammar. Cambridge: The MIT Press.
- Bierwisch, M. (1967). "Some Semantic Universals of German Adjectivals," in Foundations of Language, 3:1-36.
- Binnick, R. I. (1971). "Bring and Come," in Linguistic Inquiry, 2:260-65.
- Bloomfield, L. (1933). Language. New York: Holt.
- Bolinger, D. (1975). Aspects of Language, Second Edition. New York: Harcourt Brace Jovanovitch.
- Bolinger, D. (1974). "Meaning and Form," in Transactions of the New York Academy of Science, 36:218-33.
- Bolinger, D. (1970). "The Meaning of Do So," in Linguistic Inquiry, 1:140-44.
- Borkin, A. (1972). "Coreference and Beheaded NP's," in Papers in Linguistics, 5:28-45.
- Bréal, M. (1900). Semantics. New York: Henry Holt.
- Chafe, W. L. (1970). Meaning and the Structure of Language. Chicago: The University of Chicago Press.
- Chomsky, N. (1965). Aspects of the Theory of Syntax. Cambridge: The MIT Press.
- Chomsky, N. (1975). Reflections on Language. New York: Pantheon.
- Chomsky, N. (1972). "Remarks on Nominalization," in Studies on Semantics in Generative Grammar. The Hague: Mouton.

- Christophersen, P. (1939). "The Articles: A Study of Their Theory and Use in English." Copenhagen Dissertation.
- Cole, P. (1975). "Referential Opacity, Attributiveness, and the Performative Hypothesis," in Papers from the Eleventh Annual Meeting, Chicago Linguistics Society. Chicago: Chicago Linguistics Society.
- Fillmore, C. J. (1975). "Santa Cruz Lectures on Deixis." Bloomington: Indiana University Linguistics Club.
- Fraser, B. (1970). "Idioms within a Transformational Grammar," in Foundations of Language, 6:22-42.
- Dong, P. D. (1971). "The Applicability of Transformations to Idioms," in Papers from the Seventh Regional Meeting of the Chicago Linguistics Society. Chicago: Chicago Linguistics Society.
- Dougherty, R. (1970). "A Grammar of Coordinate Conjoined Structures: I," in Language, 46:298-339.
- Goodman, B. and Kolodin, I. (1939). The Kingdom of Swing. Harrisburg: Stackpole.
- Goodman, N. (1968). Languages of Art. New York: Bobbs-Merrill.
- Green, G. M. (1974). Semantics and Syntactic Regularity. Bloomington: Indiana University Press.
- Hankamer, J. and Sag, I. (1976). "Deep and Surface Anaphora," in Linguistic Inquiry, 7:391-428.
- Hintikka, K. J. J. (1973). "Grammar and Logic," in Approaches to Natural Language. Dordrecht: D. Reidel.
- Hockett, C. F. (1958). A Course in Modern Linguistics. New York: Macmillan.
- Hoeningwald, H. M. (1960). Language Change and Linguistic Reconstruction. Chicago: The University of Chicago Press.
- Jespersen, O. (1924). The Philosophy of Grammar. New York: H. Holt.
- Kaplan, D. "Dthat," unpublished ms.

- Katz, J. J. (1973). "Compositionality, Idiomaticity, and Lexical Substitution," in A Festschrift for Morris Halle. New York: Holt, Rinehart and Winston.
- Katz, J. J. (1977). Propositional Structure and Illocutionary Force. New York: Thomas Y. Crowell.
- Katz, J. J. (1972). Semantic Theory. New York: Harper and Row.
- Katz, J. J. (1964b). "Semantic Theory and the Meaning of 'Good'," in Journal of Philosophy, 61, no. 235.
- Katz, J. J. (1964). "Semi-sentences," in The Structure of Language. Englewood Cliffs, N.J.: Prentice-Hall.
- Katz, J. J. and Fodor, J. A. (1964). "The Structure of a Semantic Theory," in The Structure of Language. Englewood Cliffs, N.J.: Prentice-Hall.
- Keenan, E. L. (1975). "Some Universals of Passive in Relational Grammar," in Papers from the Eleventh Regional Meeting, Chicago Linguistic Society. Chicago: Chicago Linguistic Society.
- Kirsner, R. S. (1972). "On Deixis and Degree of Differentiation in Modern Standard Dutch." Unpublished Doctoral Dissertation, Department of Linguistics, Columbia University.
- Kripke, S. A. (1972). "Naming and Necessity," in Semantics of Natural Language. Dordrecht: D. Reidel.
- Lakoff, G. and Peters, S. (1969). "Phrasal Conjunction and Symmetric Predicates," in Modern Studies in English. Englewood Cliffs, N.J.: Prentice-Hall.
- Langendoen, D. T. (1966). "The Syntax of the English Expletive it," in Georgetown Monographs on Languages and Linguistics 19. Washington, D.C.: Georgetown University.
- Langendoen, D. T. and Bever, T. G. (1973). "Can a Not Unhappy Person Be Called a Not Sad One?," in A Festschrift for Morris Halle. New York: Holt, Rinehart and Winston.
- Lawler, J. (1974). "On Generics." Ann Arbor: Department of Linguistics, University of Michigan.
- Lehrer, A. (1969). "Semantic Cuisine," in Journal of Linguistics, 5:39-55.

- Lewis, D. (1969). Convention: A Philosophical Study. Cambridge: Harvard University Press.
- Lewis, D. (1972). "General Semantics," in Approaches to Natural Language. Dordrecht: D. Reidel.
- Lewis, D. (1976). "Language and Languages," in Minnesota Studies in the Philosophy of Science, VII. Minneapolis: University of Minnesota Press.
- Leech, G. N. (1969). Towards a Semantic Description of English. London: Longmans.
- Lees, R. B. (1960). The Grammar of English Nominalizations. The Hague: Mouton.
- McCawley, J. D. (Ms.) "The Role of Encyclopedic Information in Dictionary Definitions."
- McCawley, J. D. (1968). "Lexical Insertion in a Transformational Grammar without Deep Structure," in Papers from the Fourth Regional Meeting, Chicago Linguistics Society. Chicago: Chicago Linguistics Society.
- Miller, G. A. and Johnson-Laird, P. N. (1976). Language and Perception. Cambridge: Harvard University Press.
- Newmeyer, F. S. (1972). "The Insertion of Idioms," in Papers from the Eighth Regional Meeting, Chicago Linguistics Society. Chicago: Chicago Linguistics Society.
- Nunberg, G. and Pan, C. (1975). "Inferring Quantification in Generic Sentences," in Papers from the Eleventh Regional Meeting, Chicago Linguistics Society. Chicago: Chicago Linguistics Society.
- Parsons, T. (1960). "Durkheim's Contribution to the Theory of Integration of Social Systems," in Emile Durkheim, 1858-1917. Columbus: Ohio State University Press.
- Partee, B. H. (1972). "Opacity, Coreference, and Pronouns," in Semantics of Natural Language. Dordrecht: D. Reidel.
- Pleasants, H. (1969). Serious Music and All That Jazz. New York: Simon and Schuster.
- Postal, P. M. (1969). "Anaphoric Islands," in Papers from the Fifth Regional Meeting of the Chicago Linguistics Society. Chicago: Chicago Linguistics Society.

- Putnam, H. (1975a). "Explanation and Reference," in Mind, Language and Reality. Cambridge: Cambridge University Press.
- Putnam, H. (1975b). "The Meaning of 'meaning'," in Mind, Language and Reality. Cambridge: Cambridge University Press.
- Putnam, H. (1974). "Some Issues in the Theory of Grammar," in On Noam Chomsky. New York: Doubleday Anchor.
- Quine, W. V. (1971). "The Inscrutability of Reference," in Semantics. Cambridge: Cambridge University Press.
- Quine, W. V. (1960). Word and Object. Cambridge: The MIT Press.
- Ricoeur, P. (1969). Le Conflit des interprétations, Paris: Seuil.
- Ross, J. R. (1972). "Act," in Semantics of Natural Language. Dordrecht: D. Reidel.
- Sag, I. A. (1976). "Deletion and Logical Form." Unpublished Dissertation, Department of Linguistics, MIT.
- Schuller, G. (1968). The History of Jazz. New York: The Oxford University Press.
- Scott, W. (1819). Ivanhoe; or, the Jew's Daughter. London: Roach.
- Searle, J. R. (1969). Speech Acts. Cambridge: Cambridge University Press.
- Stearns, M. W. (1956). The Story of Jazz. New York: Oxford University Press.
- Stern, G. (1931). Meaning and Change of Meaning. Göteborg: Göteborgs Högskolas Årsskrift XXXVIII.
- Trench, R. C. (1859). English Past and Present. New York: Redfield.
- Vendler, Z. (1967). "Singular Terms," in Linguistics in Philosophy. Ithaca: Cornell University Press.
- Weinreich, U. (1972). Explorations in Semantic Theory. The Hague: Mouton.

- Weinreich, U. (1969). "Problems in the Analysis of Idioms," in Substance and Structure of Language. Berkeley: University of California Press.
- Weinreich, U., Labov, W. and Herzog, M. I. (1968). "Empirical Foundations for a Theory of Language Change," in Directions for Historical Linguistics. Austin: University of Texas Press.
- Wheelwright, P. (1935). Wheelwright's Aristotle. New York: The Odyssey Press.
- Wiggins, D. (1971b). "A Reply to Mr. Alston," in Semantics. Cambridge: Cambridge University Press.
- Wiggins, D. (1971a). "On Sentence-sense, Word-sense, and Difference of Word-sense," in Semantics. Cambridge: Cambridge University Press.
- Ziff, P. (1960). Semantic Analysis. Ithaca: Cornell University Press.