The interaction of structural and inferential elements in characterizing human linguistic communication

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Ellen Irene Lucast

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Jeanette Gundel, Advisor

January 2012

Acknowledgements

I'm grateful to the following people:

My committee, of course, including Jeanette Gundel, my advisor, who read carefully through many, many drafts in all stages of development; Apostolos Georgopoulos, for his encouragement; Nancy Stenson, for her practical advice; and Brian Reese, for always managing to recommend the exact source that it turns out I needed;

My friends and colleagues in the Program in Linguistics, past and present, including Eden, Sara, Paula, Mark, Suzanne, Kate, Chris, Mamadou, Peter, Kirsten, Sven, James, Mahmoud, and especially Kaitlin, my best comrade-in-arms throughout this project;

Scott Slattery, my academic counselor, who provided me with both encouragement and practical advice, and without whom I very literally would never have been able to write this dissertation;

Barbara London, longtime administrator in the Institute of Linguistics, who made everything happen;

Caroledith Olsen, the instructor of my first linguistics class at the U of M back when I didn't even know what linguistics was, and to whom I credit my habit of talking about angels singing when the solution to a linguistics exercise becomes clear;

Zack, Josh, Kevin, Riff, and Roy of Asymmetric Publications, LLC, for their consistent humor, media recommendations, and truly awful advice;

My community of online and Facebook friends, for their patience and encouragement;

My sisters Rebecca and Erica and brother-in-law Aaron;

And, of course, my parents Karen and Don, with my deepest gratitude for their boundless love and support. It is to them that this dissertation is dedicated.

Abstract

Is human linguistic communication different only in degree from other animal communication, or is it different in kind? If it is different in kind, can this difference best be attributed to one or a small number of core features? If so, what are these features? What role does the code itself play in characterizing human linguistic communication and what role is attributable to its communicative function?

To answer these questions, I argue the following: Human linguistic communication is in fact different in kind from other animal communication; its difference can be attributed to two main factors, one coded and one communicative, that lie at the core of the phenomenon of human language; and these two factors are a discrete combinatorial system and the ability to infer others' mental states. I demonstrate that these two factors limit the function of systems which do not display them in ways that are characteristically different from the function of human linguistic communication.

This work serves to update existing research on language features by integrating insight from the cognitivist research paradigm that currently prevails in linguistics. It also integrates two traditionally separate areas of inquiry, those of the functioning of the language code itself and of the inferential mechanisms that humans employ when using language for communication, to provide a more comprehensive theory on the nature of human linguistic communication.

Table of Contents

List of Figures

1 Introduction

1.1 Overview

In any discussion of what sets humans apart from the rest of the animal kingdom, "language" may be the most frequent suggestion. Humans certainly seem to have the ability to think of and communicate complex, abstract ideas in a way that no other species does. The recognition of this aspect of being human pervades art, mythology, and religion, and of course in more recent times has been a foundation for scientific research in areas such as cognitive psychology, linguistics, anthropology, evolutionary biology, and the cognitive sciences in general.

To assert that "language" is the characteristic property--or, to be fair, one of a handful of possibilities that might also include "culture" or "self-awareness"--that defines the human species has an undeniable commonsense appeal. Although by definition every species is different from every other, and asking what sets humans apart may not be any more interesting than asking what sets some particular species of mosquito or finch or gibbon apart from the rest of the animals, it is hard to deny a priori that human language[1], even if it turns out not to be qualitatively different from similar abilities found within the animal kingdom, is at least vastly different in a quantitative sense from anything else thus far observed. The perceptibility of this difference justifies attempting to identify more specifically what exactly the difference is.

The main sources of comparison against which human linguistic communication, henceforth HLC, is contrasted are the communication systems of non-human animals (and for some points, human communication that occurs through non-linguistic systems such as gestures and facial expressions). Such non-linguistic systems are also part of the normal development of individual animals and as such are shaped by evolutionary

[1] I use the term "language" here in a pretheoretical and nontechnical sense. Without a consensus on what exactly the word "language" means (aside from any working technical definition given in a specific context), it is dangerous to use it within an analysis concerned with pinning down the nature of the phenomenon that it names. Nevertheless to avoid this term altogether would be awkward and artificial.

constraints and whatever contingencies are placed on them by their instantiation in biological organisms. Comparing human linguistic communication with systems related to it along a shared biological dimension will serve to control for many features that HLC and non-human communication have in common, in order to best isolate those features in which the different systems contrast.

Nevertheless, the characteristic human ability that we refer to when we use the term "language" proves under scrutiny not to be a self-contained, neatly-bounded phenomenon that one might isolate by, say, subtracting a chimpanzee from a human and examining the pile of leftover parts. If we were to try to give a meaningful definition for "language" in its broadest sense--what humans do when they use a linguistic code to communicate that other animals do not--the most precise way to do so is to point to the total output of a set of interconnected systems that most likely include skills like reasoning, memory, sensory processing, motor control, information storage and retrieval, and so on. Some of these systems can undoubtedly be found in other species with varying degrees of resemblance to their counterparts in humans. Even put into grossly simplistic terms the difficulty of the question of human uniqueness becomes apparent: any other given species might display some of these abilities to a high degree of resemblance to their instantiation in humans and others to a lesser degree, and thus to show exactly which ingredients in what proportion make up the specialized cocktail of human language is a thornier problem than it may at first seem.

Specifically, as mentioned briefly above, the totality of what a human is doing when he or she is, say, chatting with a co-worker over lunch requires many simultaneous processes, from muscle control over breathing and mouth movements to knowledge of the shared social context, and many others. In a similar sense, we can describe the process of building a house as requiring skills and information in a variety of domains from joining wooden beams to knowing the zoning restrictions in the community. However, just as we can usefully sort the homebuilding processes along a continuum from the immediate, physical, and mechanical to the broad and abstract, so too is it possible to view the systems of human linguistic communication in this way.

This idea of language as a set of systems in the context of this paper deserves some elaboration. I suggest that the human language structure, which is to say the code itself, is the most central element in human linguistic communication: like the timber frame of a house, without it linguistic communication as a whole would be unrecognizable. For example, roughly speaking, the code allows us to say things, but in order to have anything to say we have to have some higher-level scheme that includes information about the world that we might want to talk about, and social or instrumental motivations that lead us to want to talk in the first place. Just as we could study the system of wood joinery that involves wood, metal hardware, and the skills to combine them, we could study a system that includes only the code. We could also look at how the wooden frame pieces combine to make the rooms of a house, or how they support the house's electrical and ventilation systems, or how the wooden structure fits into its context aesthetically or according to the terms of a homeowners' association rulebook, and so on. In the same way, I argue that it is possible to look at some of the subsystems of language in depth in order to establish at least a rough understanding of their relative contributions to human linguistic communication as it is broadly understood. Of course it is possible to take the reductionist viewpoint that if every element makes some unique contribution to human linguistic communication, then removing that element fundamentally changes the overall system. However, by listing features of human linguistic communication and comparing it to real or hypothetical systems that lack each feature in turn, it is possible to argue that some features make a contribution that seems to be more central to the essence of HLC than others.

1.2 Research questions, thesis, and approach

The broadest formulation of the overall questions that this dissertation addresses is as follows: Is human linguistic communication different only in degree from other animal communication, or is it different in kind? If it is different in kind, can this difference best be attributed to one or a small number of core features? If so, what are these features? What role does the code itself play in characterizing HLC and what role is attributable to its communicative function?

My basic answer to these broad questions is as follows: Human linguistic communication is in fact different in kind from other animal communication; its difference can be attributed to two main factors, one coded and one communicative, that lie at the core of the phenomenon of human language; and these two factors are *a discrete combinatorial system* and *the ability to infer others' mental states*. Working from a list generated by Hockett (1960), I first evaluate the listed features in order to create a revised list including only those items that are crucial to human linguistic communication. Then, I demonstrate that the interaction of the discrete combinatorial system and the human ability to make inferences is responsible for these crucial features. This reduction of human linguistic communication to two core elements allows a new perspective on the nature of human linguistic communication, parsing apart the contribution that the language structure itself makes and the contribution attributable to communicative abilities of humans that are not part of the language code.

This dissertation thus contributes to the literature on human linguistic communication in several ways. First, it takes a well-known and still widely cited paper on the topic, Hockett (1960), and reanalyzes its material from a perspective that takes into account the advances in linguistics and communication studies that have occurred in the years since its publication. It is hoped that this reanalysis will allow those of Hockett's ideas that maintain currency to be more easily integrated into contemporary discussions of the nature of human linguistic communication, while weeding out those ideas that subsequent research has deemed to be of little importance to the essence of linguistic communication.

Second, it picks up a thread in current research on human communication, specifically the work on linguistic pragmatics that Sperber & Wilson's (1986/1995) Relevance Theory represents, and elaborates on the mechanism by which human linguistic communication is achieved, namely the ostensive-inferential communication that they posit underlies human linguistic interaction. I demonstrate the way that the structural and coded elements of human language function as a flexible and precise tool by which speakers can demonstrate the content of their intended message, and through which hearers can accurately infer a speaker's intended message.

Lastly, by demonstrating that even the features that appear to be crucial to human linguistic communication thus far identified can themselves be derived from fundamental properties--the discrete combinatorial system that forms the basis of the language structure on the one hand, and on the other hand the ability to rely on the conventional meanings assigned to language structures to infer speakers' intended meanings--this dissertation represents an attempt to identify the most basic elements out of which HLC is built. One of the difficulties in addressing basic, broad questions such as 'what is language?' is that the various kinds of research done in this general area have traditionally remained fairly isolated, with the result that different disciplines end up talking past one another rather than sharing insights that allow their work to move forward toward a mutual goal. The development of a common framework would alleviate many of these problems, so that ideas from one field could transfer easily to another. Thus, identifying and justifying a set of basic assumptions about HLC will serve the aims of future research in linguistics, the psychology of human and nonhuman animals, philosophy of language, and any other discipline that concerns itself with the broad questions on the nature of language to which this dissertation contributes.

I will lay out the dissertation as follows. Chapter 2 presents a review of important papers in the areas of language evolution, the structure and function of human language, and perspectives on the descriptivist vs. cognitive paradigms that influence the contemporary state of research in this area. This serves to set the stage in a broad sense for the material that follows.

Chapters 3, 4, and 5 represent two parallel tracks, each of which develops one of the two parts of my thesis. Chapter 3 is structured around the design features of language posited by Hockett (1960), as I explain, evaluate, and develop further ideas about the contribution of each one in determining HLC's fundamental characteristics, and Chapter 4 develops the ideas of discreteness and combination in more detail. Hockett's work has been essentially taken for granted for a long time, but advances in both theoretical linguistics and animal research mean that it is time for a reinvestigation of Hockett's work. While many of the specific points that Hockett makes are made irrelevant through more recent research, many of his general ideas hold up when language and

communication are seen through a cognitive lens. The basic premise that human language is qualitatively different from any other known systems also holds up under scrutiny. His analysis of the language system into individual features, although flawed in its specifics, nevertheless remains a lucid approach to the task of identifying salient characteristics of language.

My treatment of the same topic, therefore, takes up this discussion where Hockett left off, which serves to ground my arguments in established literature while reconciling it with advances in the field that have been made in the years since his work was published. In my commentary and analysis of Hockett's work and my own further development of the ideas that he brings up, I argue first that some of his posited features do not in fact provide any useful insight into the characteristics of HLC, and second that those that do are not necessarily atomic but rather that many of them can be derived from others, in particular from the feature of discreteness which I argue is basic to many of the other properties that Hockett identifies as unique to human language.

In Chapter 5 I develop the other half of my argument, taking Sperber & Wilson's (1986/1995) Relevance Theory as a starting point for my position that humans' inferential ability is the other factor that characterizes HLC as compared to communicative systems of other animals. I lay out Sperber & Wilson's basic premise and then elaborate on it, showing that communication without inference would be impossible while demonstrating the specific mechanism by which coded information feeds the human inferential system.

Chapter 6 is a synthesis of the previous two chapters, showing how a discrete combinatorial code interacts with the inferential process to create the important features that characterize HLC.

Having introduced the idea of the code interacting with the inferential process, it is important to carefully delineate the distinction between the structural or formal elements of a coded system and the communicative use to which the code is put. In the contemporary view of linguistic and communication research this distinction is both foundational and functionally useful: separating these two aspects of HLC allows a finer-grained analysis of each one, and sets the scene for theories such as Chomsky's

generative grammar, which looks at the code as a self-contained system irrespective of its being used for anything at all, and for theories of communication that can account for the transmittal of both linguistic and non-linguistic messages.

Terminology plays some role in disguising the difference between these ideas in that ordinary usage of the word 'language' covers either or both of these two meanings. As will be shown below, the distinction between the language code and communication has played virtually no role in pre-generative work such as Hockett's but has gained in importance as researchers have explored the cognitive underpinnings of language form and function. Thus, when Hockett suggests that the use of the vocal-auditory channel and the ability of its users to lie are two features of 'language', these two separate aspects of human linguistic communication are confounded in such a way that some interpretation of Hockett's assertions is necessary before they can be evaluated with respect to more recent work in this area. Chapter 3 consists of a point-by-point discussion of each of Hockett's listed features, including my commentary on the positioning of the feature relative to the code/communication dichotomy.

Having asserted the importance of understanding the distinction between form and function, however, I proceed to demonstrate that the existence of clear theoretical differences does not necessarily imply that there is no interplay between these two aspects of HLC. In fact, both the role of the code in fostering communication and the influence of language's communicative function on the structure of the code can be demonstrated, and the combination of the two is the focus of this dissertation.

1.3 Defining terms and scope

Before embarking on any actual analysis it is necessary to define the terms that will be used and to specify the scope of the object to be studied. In this dissertation, the phenomenon under examination is specifically that which is a part of the repertoire of a typically-developing human. Except in cases of disability or neglect, every human develops language, and by adulthood is a fully competent user of at least one language. It is this biologically-based behavior which is meant by the term "human linguistic communication". I exclude such human-created systems as writing and Morse code,

which themselves are derivatives of biologically-based human language, and systems such as computer programming languages, which are human tools but do not emerge naturally in a biological organism.

It is necessary to provide definitions for several terms that I will use throughout this study in varying technical and non-technical capacities. 'Language', clearly a loaded term, will be used mainly because avoiding the word altogether would be artificially constraining. However, the term 'language' will only be used in a broad and non-technical sense for human spoken or signed codes and their written counterparts.

The term 'system' is used in a nontechnical sense to denote any collection of elements that interact in an organized way with one another.

I will use the term 'structure' to refer to the elements of a system that have to do with form, such as the inventory of signs, the rules for combining them, and their connection to real-world meanings. Discussion of the properties of a system's structure is generally meant to contrast with discussion of the system's usage, to emphasize the independence of these two aspects of the system.

'Human linguistic communication', or HLC, will be used to refer to human communication by means of a human language code such as Japanese or Hindi, and is a term meant to exclude facial expressions, body postures, and other such non-language indicators. (Signed languages are included under 'human linguistic communication' despite their gestural modality, since they arise naturally when the vocal-auditory modality is blocked and since they have all of the complexity and communicative capacity of spoken languages. This point is examined in some detail later in the paper but it is worth mentioning now as well.)

While the most precise terms for the participants in a communicative exchange are 'sender' and 'receiver', I sometimes vary these with 'speaker' and 'hearer' for stylistic reasons. In doing so I do not mean to exclude non-spoken language modalities such as sign or writing, or even necessarily to exclude non-linguistic communication: communication of any kind will have a sender and a receiver, and some of the points I make about HLC hold true of non-linguistic human communication as well.

Lastly, the term ‘meaning’ can be ambiguous, and so I specify ‘conventional meaning’ or ‘semantic meaning’ for the formal and systematic mapping of form to concept, and ‘speaker meaning’ when discussing the message that a sender intends to convey to a receiver in a given context of use.

I note one additional convention: following Sperber & Wilson’s (1986/1995) model, I use feminine pronouns to refer to senders and masculine pronouns for receivers, in an attempt to avoid confusion when both roles are under discussion.

2 Background, context, literature review

A dissertation such as the present study must be situated in relation to the major works in related fields. The two most important for my purposes are Hockett's (1960) feature-based description of language and Sperber & Wilson's (1986/1995) Relevance Theory, and I discuss these two works within a chapter devoted to presenting, interpreting, and elaborating upon the major ideas of each. However, other influential areas of research for a discussion of the characteristic features of human linguistic communication include language evolution, non-human animal communication research, comparative work in descriptive and theoretical linguistics, human language universals, and the philosophy of language. In this chapter I discuss a selection of works that together serve to provide context for the specific thesis that I put forth and defend in this dissertation.

2.1 Philosophy of language

Inquiries about the nature of language are not restricted to scientific approaches. Philosophers have also concerned themselves with trying to understand the relationship between language and thought and to identify the locus of meaning. Some philosophical work has hit on much the same dichotomy that I am concerned with reconciling in this dissertation: that of the coded versus communicative features of human linguistic communication. Taylor (1985: 215 and onward) summarizes the historical basis for, in particular, theories that he terms "designative" and "expressive" theories of meaning. Taylor's designative theory of meaning is roughly what linguists might call semantic meaning or sentence meaning, whereby the meaning of a sentence is composed of the conventional meanings of its parts. Taylor contrasts designative meaning with what he terms "expressive" meaning, the capacity of speakers to use language to express their thoughts or emotions, which is somewhat different from the pragmatic or speaker meaning that a linguist might identify as the counterpart to semantic meaning. By way of contrasting the two, Taylor claims that music has expressive but not designative meaning.

It is clear that a musical phrase is not assigned a propositional content via concrete mapping between meaning and form--it cannot be used to talk about the weather, for example--even though the idea that the composer of a piece and its performer are expressing something to the audience is central to musical theory.

The difference between "expressive meaning" as Taylor discusses it and the speaker meaning that is commonly discussed in the field of linguistic pragmatics deserves a brief note. While speakers are often--or, as I will argue later in this dissertation, always--communicating more to an intended hearer than just the coded meanings of the words and sentences that they utter, the difference between the coded meaning of what is said and the speaker's intended meaning goes beyond simply communicating the speaker's emotions or attitude. The mechanism for this is discussed more fully later.

In addition to recognizing two separate ways of defining 'meaning', Taylor makes another point that parallels more technical work in linguistic pragmatics, such as Sperber & Wilson's (1986/1995) Relevance Theory, in the sense that when one person speaks to another, information that may have been known to each of them privately becomes part of a "public space" (259), or as Sperber & Wilson have phrased it, becomes "mutually manifest" (1995:42).

The thesis that I present in this dissertation addresses a similar question to that which Taylor discusses in general terms: basically, the roles of form and function in understanding the essence of language. While Taylor takes a philosophical approach to the question, sketching out the different aspects of language in broad terms, my work capitalizes on the discussion he presents while pinning it down in more precise terms and grounding it in current linguistic theory.

2.2 Language universals

Because an analysis of the features of language is meant to describe human language in general and not be specific to any particular variety such as Zulu or German, linguistic research on language universals--specific features that all human languages share--is also relevant background information for the present work. Discussions about language universals often center around a list of proposed universals such as those found

in Greenberg (1963, 1966) or Pinker & Bloom (1990). These universals are largely concerned with specific properties of languages, naming elements that seem to be common to all known human languages, such as the presence of nouns and verbs, and implicational scales for phonological and syntactic features such as 'if a language has voiced stops then it also has voiceless stops'. Indeed, this is what is usually meant by the term "language universals".

This type of inquiry has its origins in a descriptive linguistic research tradition that dominated the field of linguistics through the middle of the 20th century. With the cognitivist shift toward mentally-based explanations for linguistic phenomena, the emphasis changed from simply describing observed patterns among languages to finding explanations for these patterns from human evolution and the cognitive architecture of the human mind.

Despite the differences between the earlier, descriptive tradition and the current cognitive paradigm, both research models provide evidence for linguistic universals: descriptive linguists found implicational and (apparently) absolute universals by comparing the features of languages that they described, while a central tenet of the cognitive research paradigm is that human minds share a cognitive architecture that is responsible for the abstract structural properties shared by all human languages. Thus the pairing of a bank of empirical evidence for absolute and implicational universals and near-universals with the theoretical framework of Universal Grammar that suggested a reason for the observed cross-linguistic similarities, has led to much fruitful research into the nature and form of human language.

However, the idea that an innate mechanism in the human mind that governs the shape of language is responsible for cross-linguistic similarities is not espoused by all scholars in the field. Evans & Levinson (2009) take a different perspective, suggesting that similarities among languages are "better explained as stable engineering solutions satisfying multiple design constraints" and suggesting that "the constraints of human cognition" are only one factor in the emergence of the attested patterns of languages (429). Evans & Levinson make the point that "we are the only known species whose communication system varies fundamentally in both form and content" (2009:431).

Bybee (2009) takes this same stance, saying that "[l]anguage universals (in the weak sense of statistical tendencies) are…emergent rather than given a priori" (2009:18). Basically, Evans & Levinson make the argument that the similarity in structure among various human languages can be explained by looking at the constraints on its formation and use, such as the types of things that need to be communicated and the physical parameters of human perception and production without resorting to an explanation based on any particular cognitive predisposition toward language in humans.

It should be noted that both agreement and disagreement with Evans & Levinson's premise is lively within the study of linguistic universals. Such an abundance of viewpoints is to be expected in an area of study which is so necessarily speculative and which touches so closely on the idea of what it means to be human.

This dissertation is concerned with HLC as it exists in modern humans, and thus in one sense the steps involved in its evolution do not matter; only the final product matters. However, unless we are to believe that a generation of fully linguistic and communicative offspring was born to a community of non-linguistic and non-communicative parents, there must have been intermediate steps in both the development of communication and the development of a language structure.

The details of this debate are largely irrelevant to this dissertation, because I am looking at universal properties of human linguistic communication at a much broader level than the syntactic and phonological details that are generally discussed under the heading of linguistic universals.

Yet even among a set of scholars interested in finding the key essence of human language, some aspects of the question have remained unexamined. For example, Greenberg (1963) discounts from his list of universals those elements which are "universal by definition--that is, we would not call the object in question a *language* if it lacked these properties" (E&L 437 citing Greenberg et al. 1963 p73). Here again is the myth that 'language' has some sort of definition by which any given "object in question" can be labeled as being or not being a language. Indeed, the whole premise of searching for language universals rests on the assumption that the field already has a consensus on the identity of 'language' in the first place, simply because to look for universal

properties among a set of entities, one first has to define what belongs in the set and what does not.

It is not hard to understand why such an assumption could and still does go unchallenged, since human language appears to be so qualitatively different from any other known communication system that the heuristic of 'knowing it when we see it' works well for most purposes. When the goal of a research program is to identify, classify, and describe the features common to the set of human languages, this prescientific view of 'languagehood' works just fine.

However, given the slightly different goal of systematically describing human language based on features that are in many cases shared by other, non-human communication systems, the 'know it when we see it' approach obviously falls short. Research in a cognitive paradigm is able to ask and search for the answers to the main question that the descriptive tradition ignores: not just what the properties of the system are, but crucially *why* they emerge. To understand this it is necessary to look deeper than surface characteristics of the language system to the cognitive architecture that might underlie them. This is the distinction that Chomsky makes by dividing language into "E-language", the external forms of language as produced by speakers, and "I-language", the mental systems that lead to the production of this external language (1986).

It is this deeper level which I intend to examine with this dissertation: taking an essentially cognitivist position, I look at the general design features of language such as those laid out by Hockett (1960) but with the intention of extending the concept of I-language to these features rather than simply describing their E-language forms as has been done in the past.

2.3 Language evolution

In studying the features that make human language useful for communication, it is important to pay some attention to the way those features emerged in the competence of modern humans. The evolution of human language is a murky topic at best, necessarily involving a good deal of speculation. Language leaves very little if any direct trace in the fossil record (some scant information about brain size and vocal apparatus being about

the only physical clues to a species' communication system), thus ruling out the major source of information in the study of evolution. The other main technique of evolutionists is to compare the characteristics of species closely and distantly related to the species in question for genetic, anatomical, and behavioral similarity, but here again language evolutionists are stymied: There are no surviving members of hominid species intermediate between humans and our nearest relatives, the chimpanzees. These "missing links" have left little in the way of fossil evidence much less extant populations whose communication and language use could be studied.

Such a lack of evidence, however, has not stopped linguists and anthropologists from formulating hypotheses that could direct future research in these fields. One well-known exchange on the subject is that between Hauser, Chomsky, and Fitch on one side and Pinker and Jackendoff on the other (Hauser, Chomsky, & Fitch 2002, Pinker & Jackendoff 2004, Fitch, Hauser, & Chomsky 2005, Jackendoff & Pinker 2005), which touches on several relevant aspects of the debate about the nature of human language and its evolution. Hauser, Chomsky, & Fitch open the debate as follows:

> …We submit that a distinction should be made between the faculty of language in the broad sense (FLB) and in the narrow sense (FLN). FLB includes a sensory-motor system, a conceptual-intentional system, and the computational mechanisms for recursion, providing the capacity to generate an infinite range of expressions from a finite set of elements. We hypothesize that FLN only includes recursion and is the only uniquely human component of the faculty of language. We further argue that FLN may have evolved for reasons other than language… (2002).

Hauser, Chomsky, & Fitch present this as a hypothesis to be empirically tested, the premise being that if either a) recursion is found in non-human animals' cognitive or communicative abilities, or b) something other than recursion is demonstrated to be unique to language, then the boundaries of the narrow language faculty will have to be redefined.

Much of the thrust of Hauser, Chomsky, & Fitch's discussion involves hypothesizing about evolutionary mechanisms that would have shaped the development of the narrow language faculty, and the extent to which elements of the broad language

faculty are shared with non-human animals, both in the great apes that are humans' nearest evolutionary relatives and in animals such as dolphins and birds which are only more distantly related to humans. Crucially, another claim that Hauser, Chomsky, & Fitch make is that human language did not necessarily evolve for communication, but rather could be useful in a private, internal domain for problem-solving and other cognitive functions.

Pinker & Jackendoff provide a counterpoint to Hauser, Chomsky, & Fitch's (2002) arguments by suggesting that the proposal that only recursion is "uniquely human and uniquely linguistic" is "problematic" because it "ignores the many aspects of grammar that are not recursive, such as phonology, morphology, case, agreement, and many properties of words. It is inconsistent with the anatomy and neural control of the human vocal tract. And it is weakened by experiments suggesting that speech perception cannot be reduced to primate audition, that word learning cannot be reduced to fact learning, and that at least one gene involved in speech and language was evolutionarily selected in the human lineage but is not specific to recursion" (2005).

Specifically for the present purposes, Pinker & Jackendoff claim that "[t]he system of sound distinctions found in human languages is both specific to language and uniquely human" (2005: 202). They also speculate that "the language faculty evolved in the human lineage for the communication of complex propositions (2005: 204).

Pinker & Jackendoff's argument takes a strong form when they claim that "major characteristics of phonology are specific to language (or to language and music), uniquely human, discretely infinite, and not recursive. Thus phonology represents a major counterexample to the recursion-only hypothesis" (2005: 212).

Word learning is another aspect of human language that Pinker & Jackendoff cite as potentially unique both to humans and to language (2005: 214-215), partly because "a good portion of people's knowledge of words (especially verbs and functional morphemes) consists of exactly the kind of information that is manipulated by recursive syntax, the component held [by Hauser, Chomsky, & Fitch] to make up the narrow language faculty" (2004: 215). In addition, Pinker & Jackendoff point out that while

human language is recursive, "the fact that actual human languages are a minuscule and well-defined subset of recursive languages is unexplained" (2005: 217).

The debate between Hauser, Chomsky, & Fitch and Pinker & Jackendoff represents an attempt to address the same broad questions as my work does, questions about the nature of language and the nature of being human. The details, however, are somewhat different: Hauser, Chomsky, & Fitch and Pinker & Jackendoff discuss human language in terms of evolution and the resources that the human brain uses to acquire and use language. My work looks at the use of language for communication, a point downplayed by Pinker & Jackendoff and set aside entirely by Hauser, Chomsky, & Fitch. I look at the features that allow humans to communicate with language, which include the structural property of discrete combination--a requirement for recursion, the significance of which is Hauser, Chomsky, & Fitch's main point--but also human communicative ability in the form of inference, which is an element of my discussion that Hauser, Chomsky, & Fitch and Pinker & Jackendoff do not address.

2.4 Non-linguistic human communication

Another perspective, and one that sheds light on the importance of separating the formal language system from the communicative capacity, is found in Burling (1993). The core of Burling's argument as relevant to the present inquiry is that humans communicate with (at least) two different fully-developed systems: 'language', and what he refers to as the "primate gesture-call system", which includes gestures, facial expressions, body language, and vocal sounds such as laughs and screams that correspond systematically to emotional states and can be reliably interpreted as such. The two systems are used simultaneously in face-to-face communication but convey different information: propositional content can only be transmitted through language, but intentions, emotion, states of mind, and other interactional messages can be conveyed through the gesture-call system as well.

The difference between language and communication is not a concept that is original to Burling (1993), but he enumerates the distinction with clarity and insight. The

human language system is rightly put into perspective as only one of the tools available for humans to use when conveying information.

Burling (1993) continues in his elaboration of the differences between the human gesture-call system and human language by evaluating them based on previously-posited design features drawn from Hockett's (1960) list. Among the differences he cites are the fact that the gesture-call system cannot be used to discuss entities removed from the time and place of speaking (that is, it does not show displacement); it cannot be used to convey propositional content such as, in his example, that "the bird we are both looking at is red" (1993:32). In addition, Burling explains that the gesture-call system is not productive in that new signals cannot be generated; it is not overall culturally transmitted, in that facial expressions, vocal cries, and gestures such as pointing are human universals; it is graded rather than discrete in form, in that gestures and expressions shade into one another; and it cannot easily be used to lie, in that many gesture-calls are involuntary and difficult to conceal or control.

Both the language system and the gesture-call system are true systems in the sense that there is an organizing principle that links signs to meanings: in language it is a convention shared among speakers of the particular language, and in the gesture-call system it is some more basic principle of human social perception that allows a scream to be interpreted as a distress call, for example, or wide eyes and an open mouth as an expression of surprise. This comparison between two systems usable for communication, then, provides an excellent case study for the task of separating the formal properties of language in particular from those of other systems.

2.5 Language abilities in nonhuman animals

My claim is that the characteristic features of human linguistic communication arise from two fundamental elements: a structural element, the discrete combinatorial system; and a usage element, the ability to draw inferences about others' communicative intentions. In chapters 3 and 4 I analyze existing claims about the key features of human linguistic communication and presented a revised list of features that are present in HLC and that name aspects of its flexibility and utility, such as displacement, productivity, and

so on. In order to highlight the role that each of these features plays in HLC and the way that each derives from the two features that I claim are fundamental to the nature of HLC, it is necessary to contrast HLC with the systems that other animals use in communication.

The goal of this section is to present an overview of some of the most 'advanced' nonhuman animal communication systems and language experiments in the literature, in order to provide context for later discussion and analysis of specific similarities and differences that these bear in relation to HLC. Note that this dissertation is not concerned with the most commonly asked framing of this type of investigation: the question of whether a given species 'has language'. As I have mentioned previously, this question relies on the definition of 'language' and is ultimately uninteresting: it is clear that by definition no other species has HLC, and moreover to phrase the question this way implies a value attached to 'having language' that is unnecessarily anthropocentric. In looking at HLC in comparison with other animals' communication systems, the interesting question to ask is what, if anything, makes HLC unique. Moreover, the nature and functioning of HLC is best highlighted by contrasting it with systems in which one or more of the HLC features in question is different or absent, and in this way nonhuman animal communication systems can provide useful case studies. Note too that discussions of human uniqueness, even in purportedly objective contexts, tend to become polarized along lines that are more political than scientific: one camp appears to have a stake in defending human uniqueness, the other in breaking down the traditional division between humans and the rest of the animals. This dissertation will appear to belong to the former camp, because it makes the claim that two basic elements of HLC originate the key features through with HLC operates and that other animals' systems are limited because they lack one or more of these features. However, in pointing out contrasts between HLC and other systems I do not intend to imply any sort of judgment against the nonhuman animals that I argue lack a given feature in their communicative system. And, in fact, to argue rigorously that a certain feature is present or absent in a given nonhuman animal communication system is not in itself the focus of this dissertation. Rather, I will review what the literature has to say about the communication and language abilities of a given species, and make comparisons between the functioning of that species' system and

HLC in those cases in which the understanding of a particular point would benefit from the comparison.

Some animals tend to be mentioned frequently in contexts in which communication systems are discussed. In particular, bees, vervet monkeys, dolphins, and great apes are often cited as nonhuman animals whose communication is particularly interesting in comparison with HLC. I discuss each of these in turn.

2.5.1 Bees

The bee dance system, as described by von Frisch (1967) and elaborated in Gould & Gould (1988), allows a bee to tell its hivemates the location of a food source that it has found, or less often for sources of water or a potential location to build a new hive. The dance consists of a figure-eight pattern in which the bee executes a series of waggling steps in between rounding each of the two lobes of the figure-eight. The length of the waggle portion corresponds to the distance from the hive to the food, and the orientation of the dance on the vertical wall of the hive indicates the direction of the food, with vertical corresponding to the location of the sun along the 360-degree circle of the horizon.

The structure of the bee dance uses partially iconic symbols--the length of the waggle run, and the orientation of the dance--which are discrete from one another but not combinatorial: the bee can select a value for length and a value for orientation, but only one value for each of those two parameters is available per dance and not, for example, a series of long and short waggles or an orientation that changes according to some pattern. Bee dance is also both productive and displaced, although since the only messages that bees transmit are about locations, new messages only consist of previously unmentioned locations rather than including any other content such as comments or warnings or descriptions. In addition, locations mentioned are only displaced in physical space, not time, and are only described by their position relative to the location of the dance rather than simply being uncoupled from a particular space or time.

Bee dance has no ambiguity: the direction and distance to the indicated location are precise and leave no room for misinterpretation, and no need for inference beyond the

overtly expressed direction and distance to the food source. Moreover, there is very little evidence that bees attribute states of mind to other bees or that they do so in making inferences about what is being communicated. For one example, scout bees were released from one side of a building and shown a food source on the other side, and thus forced to fly around the building rather than traveling straight between the source and the hive. When these scouts danced about the location of the food, however, they indicated the straight line path that would pass through the building, which was not a possible route for the bees receiving the dance information (Gould & Gould 1995:63-64).

Of course, a confounding problem may be that the bee dance code does not appear to include a way to describe a roundabout path to a particular location, but only the distance and direction to the end goal: if the bee were able to represent the idea of going around an obstacle and then returning to the indicated path but simply lacked the means to "say" that, the net result would be the same. As such, it is impossible to tease apart the difference between the dancers' simply not having the "vocabulary" to describe the roundabout path and their neglect in taking the perspective of their audience by indicating the actual means of accessing the food source.

2.5.2 Vervet monkeys

Cheney & Seyfarth (1990) provide an account of the social and communicative behavior of vervet monkeys in the wild, which includes discussion of the monkeys' system of calls to warn the troupe of predators. Specifically, Cheney & Seyfarth (after Struhsaker 1967) identify three different calls, which they dub the "leopard alarm", "eagle alarm", and "snake alarm" (102-103). Each of these calls is followed by a specific behavior in the other monkeys to look for or evade that particular predator.

This communication system consists of a set of symbols that are discrete but not combinatorial: the entire message in each case consists of one of these three different calls. Vervet monkey calls also show neither productivity nor displacement: their communication appears to be limited to these three calls which are used in situations in which one of the three predators is in the vicinity at the time of calling.

There has been a good deal of analysis about the referential role of these calls in the vervets' communication, which Cheney & Seyfarth summarize: The calls could be referential, referring to the predator in question or to the evasive behavior that the troupe should adopt; they could simply announce the caller's future action (that is, that the caller is about to flee to the treetops or to the ground and so on); or they could just be an emotional response, essentially screams of surprise or fear. Likewise, hearers of the calls may understand that the caller is providing them with a warning, or they may simply have learned the behavioral response that is the most likely to allow them to escape a given predator. Cheney & Seyfarth acknowledge these possibilities although they take a firm stand that the calls are indeed referential:

> "…we suggest that "semantic" and indexical information are combined in vervet alarm calls, much as they are combined in human speech. Each type of alarm refers to, or denotes, a particular type of predator. .. Supplementing and enriching this semantic information are indexical features that provide information about, for example, a caller's identity, her level of fear and anxiety, or the probability that she is likely to flee. Semantic information is of primary importance, but it is by no means the only sort of information conveyed." (110)

There is some evidence for intentionality in vervet communication. Cheney & Seyfarth note that vervets have been observed not to call in response to a predator if they are alone (145) and that females call more frequently when their own offspring are nearby than if unrelated juveniles are present (146). These behaviors suggest that the caller intends the call to be received, and that the caller 'understands' at some level that providing the call increases the chances that its hearers will survive. Although much more research on vervet psychology would have to be done to verify or refute claims of communicative intention, a generous interpretation of vervet calls might allow that the monkey that makes a particular call "knows" that its troupemates might not be aware of the predator, and that monkeys who hear the call "know" that the caller is calling for this reason.

Cheney & Seyfarth devote a good deal of their work to asking questions about vervet call "meaning" and other debated topics and designing experiments to test some of the suggested hypotheses, of which the summary I present here is only a small sample.

With vervets too, as with bees, it is necessary to ask whether there is any ambiguity in the communication system. Cheney & Seyfarth (1990) specifically claim that the leopard, eagle, and snake calls refer to their respective predators (110), but the evidence that the animals are actually referring to these dangers, rather than simply reacting to them, is shaky: the calls only have meaning in the presence or potential presence of the associated predator, and thus might equally be analyzed as warnings or calls to action with only a tenuous connection to the identity of the predator itself. When a monkey utters a "leopard" call, then, the call can be analyzed as ambiguous among all of these meanings. However, to the monkeys hearing the call, an ambiguity between interpretations such as "There is a leopard nearby" and "Climb into the trees!" is not likely to be behaviorally relevant, since simply being aware of a predator's proximity is mainly useful as a trigger for evasive behavior on the part of a prey animal such as a vervet. Because we cannot ask the monkeys what they are thinking, it is difficult to parse apart different meanings that the calls might have.

2.5.3 Apes

Research on great apes' language abilities tends to be different from research on animals such as bees and vervets in that the object of study is more often the ape's ability to learn and use some system--usually a human or human-invented one such as sign language or a symbol board on which the ape points to shapes that roughly represent words--rather than the native communication system of the animal in the wild. Being examples of, or at least modeled after, human language, the structural properties of the systems that these animals are taught to use are similar or identical to those of HLC, and as such need not be discussed individually. The communication that apes perform with the codes they are given, however, is a much more interesting object of study, as presented below.

2.5.3.1 Chimpanzees

The evidence for chimpanzee intentionality is some of the strongest to be found in non-human animals. Chimpanzees can be taught to lie and to interpret lies, as when captive chimpanzees were trained to show a "friendly" trainer where food was hidden but to misdirect an "unfriendly" trainer (Premack & Premack 1983:51-57). One specially language-trained chimp, Sarah, was also able to infer the intentions of human actors by selecting an appropriate solution to the problem she observed them struggling with, for example by selecting a picture of a key after observing an actor trying to escape from a cage. Moreover, Sarah selected good solutions for actors she liked, and mishaps for actors she didn't like (57-66).

Savage-Rumbaugh (1986) describes an experiment that may show that chimpanzees can understand that another actor can know something that the chimp itself doesn't know, and vice versa. This ability to understand that another's state of mind may be different from one's own is another prerequisite for intentional communication. The chimpanzees Sherman and Austin each learned to tell the other the contents of a closed container filled with food. First, the "informer" chimp was shown the food being put into the container. Then the "informer" chimp had to tell the "learner" chimp the contents of the container using the lexigram system that the two chimps had already learned. When the "learner" then told a teacher the contents of the container, both chimps got a taste of the food as a reward. The two evidently learned that one knew the contents and that the other did not, and also learned how to transmit that information in order to get a reward (136-148). The act of telling the contents to the other chimp thus seems, according to Savage-Rumbaugh, to be a case of intentional communication.

A thorough critique of this research is outside of the scope of this paper. However, the limitations and implications of the work described here have not been discussed from a linguistic point of view. If special training is required before a chimpanzee such as Sarah can infer actors' intentions, and she is the only chimp so far to have demonstrated this ability, it is difficult to generalize about the abilities of chimpanzees in general from this one data point. In the case of Sherman and Austin telling one another the contents of a container of food, the behavior of these two

chimpanzees may have arisen from their knowledge of the other's state of mind, but may also be an artifact of the task design or simply a pattern of behavior learned by rote or by reinforcement. Nevertheless, the presence of some evidence for inference in chimpanzees suggests connections between chimpanzee communication and that of humans that may prove fruitful for future research.

2.5.3.2 Gorillas

Language research with gorillas, such as that done with the famous Koko, has provided possible evidence of intentional communication in this species as well. The most basic prerequisite of intentional communication is self-awareness, because in order to understand that another entity has a mind it is first necessary to recognize the boundary between oneself and the external world. The classic test for self-awareness is one in which an animal is shown to recognize itself in a mirror by its pointing to a spot of paint applied surreptitiously to its face in a location that it cannot see except through its reflection. Koko has been shown to pass this test, as well as use more elaborate variations such as using the mirror to apply lipstick and to groom otherwise-invisible parts of her body (Hillix & Rumbaugh 103-104).

There is also evidence from research with Koko that gorillas can take advantage of signed language's discrete combination to create new expressions, which is to say that Koko's communication displays productivity. Koko has been trained to sign, and can also understand some spoken English. She has been reported to use sign innovatively, inventing terms for new objects such as "eye hat" for a mask (Hillix & Rumbaugh 109). In this way Koko genuinely exploits the discrete combinatorial structural properties of the system that she has been taught in order to create new expressions.

Gomez (1990) also describes some gorilla behavior as intentional communication, such as the gorilla's extending its arms toward a keeper as a request to be picked up. The assessment that this is intentional communication is based on two factors: One is eye contact between the gorilla and the human with which it is interacting; and the other is that the intended effect, the human's lifting the gorilla, is brought about not as a physical

consequence of the gorilla's gesture but as a consequence of the human's interpreting the gesture as a request to be lifted (346).

Eye contact can be an indication of intentional communication for several reasons. The direction of an animal's eye gaze can be a fairly reliable indication of the direction of its attention. A sender's gaze toward a receiver both indicates to the receiver that the sender is attending to him, and allows the sender to assess whether the receiver is in turn attending back to her. Eye contact, the gaze of the sender and receiver toward each other, allows each party to know that the other is attending to him or her. The gorilla's gesture of stretching its arms toward its keeper would be much less effective as an act of communication if it did not simultaneously establish that the keeper was attending to the gesture. The idea that the lifting comes about as a result of the keeper's recognition of the gorilla's desire to be lifted is also strong evidence of the gorilla's communicative intentions. The gorilla could presumably bring about its desired outcome by physically climbing up into the keeper's arms, and in fact this is the only way for the gorilla to achieve its aims without resorting to help from a cooperative partner. However, the fact that the gorilla chooses the communicative gesture to achieve its goal shows that the gorilla must believe that the keeper will behave a certain way--by picking the gorilla up--as a result of the gorilla's gesture.

There is of course much more written about ape communication than just the works briefly discussed here. There is also more analysis necessary and more questions to be answered before any conclusion about apes' abilities can be reached. One of these questions is the comparison of lab-trained apes' abilities to use an artificial language system to the spontaneous native communication between apes in the wild. Clearly an animal cannot be trained to do something that is impossible for it, and as such the lab experiments do show a real ability of the apes to take another's perspective. On the other hand, without comprehensive studies of ape communication in the wild, it does not seem accurate to say definitively that intentional communication forms part of the ape's repertoire of communicative abilities.

2.5.4 Dolphins

Dolphins have also been the subjects of language-training research, due to their high brain-to-body size ratio, or encephalization, which is also high in other communicative animals such as humans and other great apes and which may be correlated with linguistic or communicative ability.

In the dolphin research done thus far, the object of study is not the dolphins' native interspecies communication but the ability of individual dolphins to master artificial language systems taught to them by human trainers. Herman (1986) details training experiments done with two dolphins, one of which was trained in an audio-broadcast artificial language and another in a gestural system performed by human trainers (228-229). Herman claims that both dolphins were sensitive to the syntactic information encoded in the word order of commands such as "Take object A to object B" in both the acoustic language, where the word order was OBJECT-A FETCH OBJECT-B, and the gestural language, in which the order was OBJECT-B OBJECT-A FETCH, and further claims that what he terms the "inverse grammar" of the gestural language was more difficult for its dolphin to interpret than the "linear grammar" of the acoustic language which he explains as presenting ideas in the order in which they are acted upon (229), which is to say iconically. Herman further claims that the two trained dolphins responded accurately when new words, individually taught to each dolphin, were introduced in familiar sentence patterns; that the dolphins likewise correctly interpreted new sentence structures that consisted of familiar elements; and that the dolphins generally performed well when the commands were two conjoined sentences such as "PIPE TAIL-TOUCH PIPE OVER", which requires the dolphin to touch the pipe with her tail and then swim over the pipe (235).

Herman's aim in performing these studies is to investigate the capacity of dolphins to learn to interpret commands given by means of a human-invented signal with a certain structure, and his research does not extend to describing the structure of a signal native to the dolphins. Thus, critique of his claims about dolphins' linguistic capacities is largely irrelevant to the present discussion. The artificial languages themselves are both discrete combinatorial systems, and the dolphins using them appeared to be sensitive to the information accessible by both the meanings of the individual words and the

significance of their ordering into novel patterns, thus demonstrating that both the discrete and the combinatorial features of the system were salient to the dolphins.

There seems to be no evidence either way about the role of inference in the dolphins' communication. The dolphin research described is performed on one-sided exchanges and not real conversations, since the dolphins are only interpreting utterances from human trainers rather than producing their own utterances. In addition, the simple commands and yes/no questions that the dolphins are asked to interpret create situations in which there is little room for the kind of vagueness or ambiguity that requires inference to bridge.

As mentioned at the beginning of this chapter, the survey presented here is only a rough overview of a few different avenues of research on the communication systems of non-human animals and their ability to learn systems with at least some of the characteristic properties of HLC. While it is human linguistic communication that is the object of study of this dissertation and not the linguistic or communicative capacities of other species, nevertheless when the discussion of a particular feature is best illustrated with examples of systems that do or do not exhibit such a feature. The brief descriptions of nonhuman animal communication systems presented here will serve as useful examples for the discussions, found later in this dissertation, of individual features of HLC.

3 Survey of posited features of human language

3.1 Generating the list of features

In this section I present and analyze a list of proposed features of human 'language' in order to set the stage for a synthesis of the key structural and communicative features of HLC in chapter 5. An obvious place to start is with the well-known discussions of design features of language from Hockett and Hockett & Altmann (Hockett 1959, Hockett 1960, Hockett & Altmann 1968). Hockett's list of the design features of human language, sixteen in all, is still widely cited when the nature of language is discussed.

The sixteen features from Hockett (1960) are as follows, with my elaborations or explanations in the right-hand column:

Figure 1: Hockett's (1960) design features of language

Vocal-auditory channel	Signal is transmitted through vocal sound rather than e.g. gesture or chemical means
Broadcast transmission and directional reception	The signal emanates in all directions but is detected from one direction
Rapid fading	The physical signal does not linger
Interchangeability	Any individual can both send and receive messages
Complete feedback	The transmitter receives his own message in real time
Specialization	The message is not the byproduct of another action
Semanticity	The signal means something
Arbitrariness	Signals do not physically resemble their referents

Discreteness	There is no continuum of signs
Displacement	Users may refer to things removed in time and space
Openness	New messages may be created
Tradition	At least part of the system is learned
Duality of patterning	A set of meaningless units combine into a set of meaningful symbols, and these meaningful symbols combine into messages
Prevarication	The system can be used to lie
Reflexiveness	Users can talk about the communication system itself
Learnability	The system can be acquired by adults

Hockett's approach to analyzing language by breaking it down into individual features is so clearly useful that I begin with his sixteen features as I perform my own analysis. This serves two purposes. One is that, in the time since Hockett created his list, research in both non-human animal communication and in human language and cognition have added much to our understanding of these phenomena. Turning a critical eye on Hockett's work from a perspective that benefits from this newer research will allow recent insights to contribute to other fields to which Hockett's features have been relevant.

The other reason to use Hockett's list is that doing so situates my work in relation to existing literature, and allows readers who are familiar with Hockett as well as those who are not a way to relate this dissertation to other works that cite Hockett's features.

While undeniably insightful, Hockett's (1960) list should, as Hockett himself says, be considered a starting place for investigation rather than a definitive pronouncement on the nature of language. More importantly, the list deserves a critical eye when used as a basis for an investigation such as the present analysis, done from a

cognitivist point of view. Hockett makes no distinction between the types of features he lists: some are physical features of the medium of transmission, some are structural features of the code, and some reflect human language's use for communicating.

My discussion in this chapter, therefore, looks at each feature individually, explaining Hockett's definition of the feature, situating it with respect to its contribution to human linguistic communication, and providing a critical judgment on its continued relevance in the face of the advances that linguistics and animal studies have made since the list was first presented.

As influential as Hockett's (1960) list has been, a thorough look at the features of human language requires more than just those sixteen identified features. In addition to Hockett's features, I include two that have become prominent in the cognitive paradigm: reference and recursion. Because of the descriptive, anthropological tradition in which Hockett was working, he emphasized surface features of human language rather than anything that might be considered abstract or cognitive in nature. Some of Hockett's features are very superficial, namely those that are characteristic of the physical medium of sound (and, incidentally, light, which carries the signal for signed languages, although it is not clear and probably unlikely that Hockett meant to include signed languages in his analysis). Others, while more abstract, are still features that can be directly observed rather than having to be inferred through a deeper analysis of language. However, the work by Chomsky that emphasized the internal structure of sentences has brought recursion to the forefront of human language characteristics being discussed today. In addition, the feature of reference, which especially in its pragmatic conception is integral to a cognitive view of language (cf. Gundel, Hedberg, & Zacharski 1993, among others), must be discussed if the present work is to be relevant to contemporary work in linguistics and particularly pragmatics and linguistic communication.

3.2 Evaluating the list of features

In order to best explain the role that each of Hockett's features plays in the present discussion, I divide them into three categories, "physical and incidental features", "functional features", and "miscellaneous features". Additionally at the end I present two

features, recursion and reference, in a category that I simply call "beyond Hockett". In this chapter I retain the original order in which Hockett presents his features because it is the simplest way to organize a systematic discussion of the sixteen items and will be more intuitive to those readers who may be familiar with Hockett's original list. However, as will be seen, I consider some of the features to be important enough to deserve their own chapter, and their entry in the discussion list is little more than a placeholder. Furthermore, Hockett's ordering lends itself well to the three categories that I posit in order to better organize the list.

The first category of features that I present, "physical and incidental features", includes *vocal-auditory channel*, *broadcast transmission and directional reception*, *rapid fading*, *interchangeability*, *complete feedback*, and *specialization*. These are part of Hockett's list and therefore deserve discussion if for no other reason than to justify dismissing them as contributing to the core nature of HLC. The overall conclusion that I draw from discussing these features, however, is that they do not fundamentally differentiate HLC from other codes and communication systems. To continue an analogy made earlier in this paper, the physical and incidental features of HLC can be compared to the specific materials used to build a house: whether masonry is made of stone or bricks, or joists made of wood or steel, lends certain characteristics to the house (for example, these materials vary in their ability to withstand fire or wind damage) but is ultimately less important than the form of the structures that these materials are used to create. Moreover, since many non-human systems share these features, they do not serve to demonstrate the qualitative difference between human linguistic communication and other animals' communication and thus have little relevance for the present discussion.

The features that I identify as "functional", which are *semanticity*, *arbitrariness*, *discreteness*, *duality of patterning* (reanalyzed as a feature that I call *combination*), *displacement*, and *productivity*, are the most useful for the present analysis, as I demonstrate that the discrete combinatorial system on which the code is based allows the derivation of the other structural features frequently considered to be characteristic of human language.

An additional category, which I am calling "miscellaneous features", covers the four remaining features on Hockett's list, *tradition*, *prevarication*, *reflexiveness*, and *learnability*. These have varying impact on the function of linguistic communication, as will be discussed.

Last, the features of *recursion* and *reference* that I am adding to the list are contained in a category simply called "beyond Hockett". These two features come from perspectives in linguistics other than the descriptive tradition in which Hockett worked: recursion gained prominence in linguistic discussions with the work of Chomsky (1957) and is considered by him and others to be the single element that defines human language and distinguishes it from the communication systems of other species (Hauser, Chomsky, & Fitch 2002). Exactly what is meant by 'reference' varies somewhat according to the perspective of the field in which it is used: in semantics, it describes an abstract relationship between a word and the concepts included in that word's conventional meaning, whereas in pragmatics, reference is an act in which a sender invokes a particular concept in a receiver's mind through linguistic or non-linguistic means. In the present discussion I use the term in its pragmatic usage.

3.3 Physical and incidental features

3.3.1 Vocal-auditory channel

The first item on Hockett's list of design features is 'vocal-auditory channel', which is to say that the signal is created using the vocal apparatus and perceived by hearing. Information can be carried through a variety of media, and a receiver could be sent information by any modality through which it can receive sensory input: the visual, olfactory, gustatory, tactile, and auditory modalities are among those by which biological organisms transmit information, and non-biological systems send information through electric charges and electromagnetic radiation as well as many other modalities.

Human language reliably emerges in the vocal-auditory channel when that modality is available, and as such it is an accepted generalization that communication by means of the vocal-auditory channel is a biological characteristic of our species.

Nevertheless, fully-formed human language reliably emerges in the visual-gestural modality in the case of the signed languages that arise in communities of speakers where the vocal-auditory channel is blocked, usually in the case of deafness.[2]

How does the use of the vocal-auditory channel serve human language's communicative function? Given that there has to be some modality or another through which the signal can travel--that communication cannot happen telepathically--the only way to evaluate one modality's expedience is to compare it to other possible channels. The advantage of auditory over visual signal reception is that sight is blocked by darkness and physical obstructions between the sender and the receiver while sound is not. Thus, speakers can communicate with visually hidden hearers and hearers whose location they may not know.

Vocalizing is also a more efficient way to send signals than signing, for the simple reason that a signer's hands and body must be free to sign while a speaker's hands can be occupied with another task. Technically, of course, a signal sent through the visual modality need not be created by posture or movement. Some animals, such as octopi, squid, and baboons, send signals through changes in color. (Indeed, human blushing is a signal of this sort.) Many animals react to aggression by puffing up feathers or fur (or the entire body, as in pufferfish) to appear larger. These means of transmitting signals are available to humans and indeed are used in body language or what Burling (1993) calls the 'gesture-call system'. However, such channels are not well suited for transmitting complex codes because they do not offer the selection of different states that is necessary in practical terms for differentiating signals of the code (see section 3.2.9 below for a discussion of the feature of discreteness). It would be difficult for a receiver to quickly and reliably differentiate between, for example, full blush, 80% blush, 60% blush, and so on. An animal can blush or not, or puff its fur or not, and thus it is possible to imagine a system built of sequences of puffs and deflations or blushes and pallors. Still, each of these modalities is only a binary system through which transmitting

[2] That signed languages are fully-formed codes with equivalent components and the same capacity for expressing speaker-meaning as spoken languages is taken for granted among linguists; see Stokoe (1972) .

complex messages is likely to be slow and error-prone. Pushing this line of thought to its limits, even an extremely detailed skin color pattern--imagine writing or a series of diagrams appearing on a sender's skin for a receiver to read --still has the disadvantage of being unavailable in darkness and among obstructions.

Signing is not the only means of communicating through the visual modality: it is also possible for an organism to manipulate the environment in a way intended to be interpreted visually. Written human languages are an example of this, as are other human visual codes such as trail signs made of stacked rocks and bent branches, and the elaborate nests that male bower birds create out of plant material, stones, and found objects such as bottle caps in order to attract mates. Writing systems are derived from spoken languages, but the others are not. Visual codes like these are more permanent than signed or auditory codes (see section 3.3.3 on rapid fading), which means that they can convey a message even without the physical presence of their creator. However, they have the same disadvantages that signed codes have in that they are imperceptible in darkness and when blocked by obstacles; in addition, the fact that they persist in the environment detaches them from time and therefore from the other clues that are present in the immediate vicinity in time and space of the communication of a message. For example, a vervet monkey that hears another make an eagle warning call behaves as if the call is relevant to an eagle that is present now. A written message or pile of stones made as an eagle warning, on the other hand, lacks the effectiveness of the vocal call because it will linger in the environment even after the eagle danger has passed.

To return to the basic question: How does the use of the vocal-auditory channel serve HLC? Vocalization is an efficient way to rapidly transmit a detailed signal: humans are capable of making and discriminating fine distinctions in sound made by the different positions of the tongue, lips, nasal passages, and vocal folds, so a large inventory of different atomic units is available to feed the systems that govern the patterning of these units into meaningful messages, as discussed below. The fact that human vocalizations are under voluntary control also makes the vocal channel extremely useful for communication. However, it is difficult to know the direction of causality between these two factors: it is plausible to speculate that, for example, the advantages

that came with using the vocal tract for communication put evolutionary pressure on the development of finer degrees of control and variation in the vocal channel. As such, we could also speculate that the same development might have occurred in any other physiological human feature, such as the blushing mechanism discussed above, so that the mechanism that causes blushing could hypothetically have come under voluntary control and become fine-tuned enough that detailed patterns of flushed and pale skin could have served as the modality by which the abstractions of language are manifested. And, of course, rapid manipulation of the hands, arms, and face, piggybacking on the manual dexterity that humans must have already used for toolmaking, gathering, and so forth, did in fact evolve as a system of expression. Because these varied modalities are all capable of conveying all of the nuance of communication of any of the other modalities, the role of the vocal-auditory channel in HLC must be analyzed as a fluke of evolution based on non-essential advantages rather than as a defining feature of linguistic communication.

3.3.2 Broadcast transmission and directional reception

Hockett's next enumerated feature of human language, broadcast transmission and directional reception, can be simply restated as the observation that the physical signal carried by sound or light waves is broadcast in all directions from a sender, and a receiver can detect by means of the same physical signal itself where the sender is located in space.

Broadcast transmission and directional reception is an epiphenomenon emergent from the use of either the vocal-auditory channel or the visual-gestural channel and results from the physical properties of sound and light waves. Most communication modalities operate by broadcast transmission; information that is transmitted by taste is an example of an exception.

Broadcast messages have the communicative advantage that the message, or at least the presence of a message, is perceptible even to potential receivers who are not actively seeking the message. A creature that advertises its toxicity through bright colors or a foul smell might thereby avert an attack in the first place, whereas a creature whose

taste shows its toxicity must, of course, be tasted by a predator, and likely wounded or killed, for the message to be received.

In HLC, wherein a sender's making manifest to a receiver her communicative intention is part of the act of communication, a modality that can be broadcast allows the transmission of the message's content to simultaneously serve as notice of the message's existence. It is a property of the auditory system that a sound can be heard by a receiver from any direction around it, including above and below. Thus a sender anywhere in the receiver's vicinity can alert the (as yet potential) receiver to the existence of a message simply by creating the message. The visual modality is also broadcast in the sense that a visual signal can be received by anyone within a line of sight, but sight in most animals is limited in direction and cannot monitor all of the three-dimensional space around a receiver that hearing can: a sender outside of a receiver's current visual field must first attract the receiver's attention to the fact that she wishes to send a message, as is done when a speaker of signed language touches a deaf person's shoulder or waves a hand within his line of vision in order to indicate that she would like to speak. Likewise, directional reception--the fact that a hearer can detect the location of a speaker based on the sound itself--helps to focus the hearer's attention on the speaker in order for the message to be successfully transmitted.

Modality, as discussed above, is only incidental to language, and the broadcast nature of hearing, and to a lesser extent seeing, is a feature of the modality and thus just as incidental. Broadcast transmission and directional reception contribute to HLC by helping to make manifest the intent to communicate. However, they are not crucial to HLC and cannot be considered central features.

3.3.3 Rapid fading

Hockett explains rapid fading by stating that "[t]he physical nature of sound yields another design-feature in any communicative system, that uses a sound channel: a signal has to be received just at the right time or it is irrecoverably gone" (1960:133). Although Hockett does not specifically say so, information carried in a visual channel

such as the signals of signed languages is subject to the same rapid fading as that of symbols carried in sound channels.

The rapid decay of the physical signal is due to the physical properties of sound or light waves as well as the perceptual, cognitive, and memory demands on the receiver to take in and parse the signal in real time. This requires efficiency in communication on the sender's part, so that the message may be interpreted as quickly and easily as possible. Rapid fading limits the theoretically infinite inventory of sentences to those which can be processed in a few seconds. It also requires that the message be encoded in such a way that the receiver reaches what Sperber and Wilson (1986/1995) term "maximum cognitive effects" as efficiently as possible.

Sound and gesture signals, unlike written or chemical signals, do not persist in the physical environment. One way in which this fact contributes to HLC is that, no matter the propositional content of the message, the fact of its existence indicates immediacy: when someone speaks, the message is being communicated here and now in a way not necessarily true for media such as odors or written signs, which may persist for minutes, years, or millennia. However, while spoken (or signed) language is the most basic form of HLC, human language can also exist in written form, which records the same information contained in the speech or sign stream in enough detail to allow messages to be recoverable. As such, rapid fading cannot be considered central to HLC.

3.3.4 Interchangeability

An interchangeable system is one in which any individual can function as both a sender and a receiver. Most biological intraspecies communication is interchangeable; the most common examples of non-interchangeable systems are probably those in which one sex signals to the other, as in the light signals of female fireflies or many pheromone systems advertising sexual availability.

The fact that humans can both send and receive the messages of HLC does not appear to play a significant role in shaping the way HLC is performed. Within phonology, patterns of sound change are traditionally explained as resulting from a balance between ease of articulation and ease of comprehension. The easiest sound

system to articulate is one with minimal movement of the articulators, which in an extreme case would likely have a single vowel and one or possibly no consonants. Such a system would have little capacity to transmit information, however, since all signals would essentially be the same. A maximally articulated system, on the other hand, could convey large quantities of information to the receiver but would require a relatively large investment of energy from the sender as she creates the sounds of the language. However, this balance between the effort of sending and receiving does not hinge on the same individuals' performing both roles. It is possible that experience as a receiver may help an individual hone her skill as a sender and vice versa. Senders who also have experience receiving others' messages may learn how to communicate more effectively, by whatever measure effective communication is defined. However, interchangeability cannot in itself be considered fundamental to HLC.

3.3.5 Complete feedback

The fact that a sender can perceive his own message at the moment of sending it contributes to HLC by allowing real-time error correction. However, 'complete feedback' is really an epiphenomenon of the use of the visual or auditory channel: because a human can hear her own vocalizations or see his own hand motions, once it is produced a signal feeds back into the perceptual system. Thus, the same arguments that demonstrate that the use of the visual or auditory channel is not crucial to HLC apply in eliminating complete feedback from the list of significant features as well.

3.3.6 Specialization

The fact that the communicated message is not the byproduct of another action aids in the receiver's perception of the sender's communicative intention. Signals such as the sound of an individual's footsteps, for example, might be used intentionally, and indeed stepping heavily (to cause vibrations felt through the floor rather than heard) is used in some Deaf communities as a conventional device to announce one's presence when a person enters a room and the deaf people present cannot see them. On the other hand, the sound made by footsteps might be involuntary, just an incidental phenomenon

of walking. Thus when an individual perceives footsteps, he does not automatically know that the walker intends to communicate her presence. On the other hand, since speech, signing, and writing are all actions specialized for communication, perception of one of these physical signals is a fairly reliable indication that the speaker, signer, or writer intends to communicate, and as such fosters the act of communication itself. The fact that humans speak or sign independently of other actions is certainly useful for HLC for this reason, but the basic nature of the communication system would not be altered in the absence of specialization, and as such this feature cannot be considered crucial to HLC.

The features so far discussed, which I have termed *physical/behavioral* properties of human language--vocal-auditory channel, broadcast transmission and directional reception, rapid fading, interchangeability, complete feedback, and specialization--clearly belong to a pre-cognitive perspective on language, when an inventory of the directly observable features of individual languages, and of human language in general, was the focus of study. These physical and behavioral properties of human language are described without invoking mental or cognitive processes to account for them, and as such these features have not retained much importance for defining human language in a cognitivist paradigm. However, the features in the next category, which I term *functional* features, are more difficult to describe without resorting to an explanation based on the mind of a user of the system.

3.4 Functional features

3.4.1 Semanticity

'Semanticity' is simply the quality of having meaning. The fact that the signals sent in communication of any kind are meaningful is obvious enough to be a basic assumption in a discussion of human linguistic communication. However, in human linguistic communication, semanticity plays a more complex role than such a simple description might imply.

Here I specify a distinction that Hockett does not make: the difference between what might be termed 'word meaning' or 'sentence meaning' on the one hand, and 'speaker meaning' on the other. The signals of human language--morphemes, words, and sentences--are conventionally assigned a semantic meaning, which links each one to a concept in the world. The meanings conventionally associated with words are independent of context. Speaker meaning, by contrast, is a product of a sender's intention to communicate some idea to a receiver. As discussed in chapter 4, speaker meaning is entirely separate from semantic meaning and can be communicated without the use of a code. As in Sperber & Wilson's (1986/1995:25-26) example, a person may reply to the question "How are you feeling?" by taking a bottle of aspirin from her bag and showing it to her interlocutor. The showing of a bottle of aspirin is not conventionally assigned the semantic meaning 'I am not feeling well', but this message can nevertheless be clearly conveyed by the gesture.

The semantic meanings of the signals of human language are extremely useful but neither necessary nor sufficient to convey speaker meaning. A premise of this paper, based on Sperber & Wilson (1986/1995), is that the coded message functions as evidence of the sender's intention to communicate as well as evidence of the intended content of the communication. Thus, any set of signs can be used to convey any meaning at all (a premise discussed in depth in chapter 5. At first glance this account of communication appears to leave no room for semantics in conveying a meaning from speaker to hearer. However, a receiver has to base his inference about a sender's intentions on some kind of shared information, or else it would be impossible to guess which of the infinite possibilities a sender intends to convey at any moment. The conventional linking of linguistic signs to meanings provides this shared information, in that anyone who speaks the particular language of the sender shares these conventions. Thus, if a sender wishes to invoke a certain idea in a receiver's mind, she need only provide the words and sentences that conventionally indicate that meaning, and need not know anything else about the receiver's experiences or state of mind in order to do so. This idea is discussed in much more depth in chapter 4.

Semanticity is not crucial to communication in general or even human communication more specifically, but it is responsible for the difference in utility of human *linguistic* communication as compared to other types of human communication. The mechanism by which the semantic content of a given utterance contributes to the communicative efficacy of that utterance is discussed in chapter X.

3.4.2 Arbitrariness

Hockett's use of the term 'arbitrariness' mirrors its common usage in linguistics, naming the fact that the signs used in the system do not physically resemble the entities that they represent. The semantic linking between elements of the code and entities in the world can be described on a continuum from iconic--physically resembling the represented object or idea--to completely arbitrary. Hockett presents the example of a map (1960:143), which I develop here: The most extreme example of iconicity would be a code in which entities somehow represent themselves: a completely iconic map, for example, is the terrain itself. Such a map would be unusable because consulting it would be identical to consulting the terrain itself, a task for which no map is necessary. Any attempt to represent an area of terrain with symbols necessarily introduces an element of arbitrariness, because the physical resemblance between symbol and idea is less than perfect. A map is of course not wholly arbitrary, since the shapes and relative positions of landforms on the map will imitate the shapes and positions of those landforms in nature. A typical subway map may be less iconic than a road map because the details about distance and relative position of landmarks is less important in the context of riding a subway than it is in driving a car, which may require taking distance and location into account. At the arbitrary extreme is a map whose images and diagrams bear no resemblance to the mapped terrain: roads could be as well represented by blotches of color or soap bubbles or musical tones as with lines oriented with respect to other landmarks on the page. This kind of map would require the user to learn associations between symbols and the things they denote without the processing help that some degree of iconicity would provide.

The pairing of sound to meaning in human language codes is generally arbitrary. This is probably largely because there is little possibility of physical resemblance between most objects or ideas in the world and the sounds made by the human vocal apparatus. The gestures of signed languages are more iconic than the sounds of spoken languages, probably because there are more gestures that look like objects in the world than there are sounds that resemble them. However, even when the iconicity of a signed word is apparent in retrospect, it is not always derivable up front. The ASL word for 'tree' is a vertical forearm with the palm and fingers spread, which looks somewhat like a tree with spread branches to receivers who know that it means *tree*, but that gesture also might resemble any number of other things equally well--*flag*, *satellite dish*, p*erpendicular,* or *tall*. In addition, the iconicity of signs is not especially important for their interpretation: signs tend to change over time to become more arbitrary (Frishberg 1975), and babies who learn the sign for *milk*, which is iconic to the gesture of milking a cow, certainly do so long before they are aware of the origin of the gesture (Amber Martin, p.c.).

Spoken language does in fact use more iconicity than is often attributed to it: onomatopoeic words such as *boom*, *snap*, and *meow* are obvious examples, but there are others including the use of high vowels in words meaning 'small'--*teeny*, *wee*--and low vowels in words meaning 'big'--*vast*, *large*--which is partially iconic to the shape of the vocal tract when these sounds are uttered and to the fact that small objects tend to make higher-pitched sounds than large objects. Even the presence of plural morphemes, such as 's' in English, which is absent in a singular noun such as 'table', forms a larger and more complex word that is to some degree iconic to the larger quantity of 'tables' that that word names (Haiman 1985).

It is not surprising that iconicity tends to appear where it is possible, since the resemblance serves as a cue that makes it easier to remember what meaning is associated with a given sound. However, because iconic resemblance is impossible in most cases, most words in a language end up being linked to their meanings by a convention shared by speakers of that language. With no way of calculating words from meanings or meanings from words, knowing the code involves storing lexical items in memory. This

has the functional result of requiring a relatively intense investment in learning word meanings before the code is usable. So, while once the code is learned it is as easy to use as it would be if meanings were iconically linked to words, the code is ultimately less user-friendly for communication than it would be if users did not have to first learn at least a critical mass of words.

On the other hand, if the human language code relied on iconicity to link symbols to meanings, a good deal of HLC would be impossible. Arbitrariness is crucial for crossing the boundary between vocal pantomime and an entirely symbolic system. As players of charades or Pictionary can attest, many ideas are difficult to represent even in a visual medium, and audio representations are likely to be as difficult or more so given humans' tendency to gather more information about the world through sight than through sound. By decoupling representation from resemblance, an arbitrary system allows us to talk about abstract concepts that cannot be mimicked with the sounds or gestures of human language.

Therefore, although arbitrary linkage between sound and meaning requires more effort as the code is learned and stored, without it we would not be able to speak about the vast range of topics demonstrated in HLC and unattested in other communication systems. Arbitrariness therefore should be considered an essential feature of HLC.

3.4.3 Discreteness and duality of patterning (combination)

Because the features of discreteness and duality of patterning (which I reanalyze as a feature I call 'combination') are crucial to the thesis of this dissertation, I discuss them within their own chapter, only mentioning them here as a placeholder so as to adhere to Hockett's list as closely as possible.

3.4.4 Displacement

A communication system with the feature of displacement allows users to communicate about entities that are not present at the time and place of speaking. Displacement is a relatively rare feature of communication systems, likely because things in the immediate environment are of the most interest to most animals. Honeybee

dances, which tell a bee's hivemates where to find a food source, are the best example of a non-human communication system in which the message is about something not present at the time of communicating. However, the location of food is always relevant to the immediate context, since gathering food is a neverending task, and moreover the location being danced about is given relative to the immediate context, which severely limits the scope of displacement in their dance system.

Unlike bees, humans talk about a large variety of displaced entities, not only physical things in our experience that happen to be removed from the place where we are speaking, but also events in the past and hypothetical future and imaginary, invisible, and abstract things and ideas. Displacement is crucial to HLC in that it describes the enormous scope of content that can be communicated.

3.4.5 Openness (productivity)

Openness, also called productivity or creativity, is the capacity of the system to support the creation of messages never before expressed. Productivity in human language takes two significant forms: *new phrases*, which are new combinations of elements built according to existing rules, and *new mappings*[3] between units of meaning and the concepts that they name.

New phrases are, of course, extremely common. An outlandish example such as 'Confucius took the last plane out of Las Vegas before snark-hunting season began' has almost certainly never been created before, and even many much more ordinary utterances that a given speaker makes in a given day are unlikely ever to have been said word-for-word by that speaker. Even those that do happen to have been said in the past are nevertheless unlikely, except in the case of stock phrases such as 'Have a nice day' or 'Pleased to meet you', to be sentences drawn fully formed from memory.

[3] I use the term 'new mappings' rather than 'new words' to emphasize that it is the new link between form and meaning that is under discussion, rather than a newly invented phonological form. New forms can o course be coined out of thin air, but as this is relatively rare (and anyway it is essentially governed by the same rules as govern the creation of phrases, i.e., a combination of existing units according to existing rules.

New mappings occur in several ways. One is that a word's meaning changes over historical time, such as the meaning of English 'silly' which has changed from *pure, innocent* to *inane*. Also, previously metaphorical usages may start to become more central to a word's definition, as in English 'mouse', which now describes a computer pointing tool as easily as it does a small rodent. Another example is in the borrowing of a foreign word, such as 'taboo' from Tongan, which means *forbidden* in English. Lastly, although the process is relatively rare, new phonological words can be coined wholesale and enter the general vocabulary, such as the name of the subatomic particle *quark*, which, when newly discovered, was named from a nonsense word in a poem (OED online).

No other species appears to have a communication system that displays productivity[4], and certainly none shows it to the extent that human language does, in which most messages are newly created in the communicative situation rather than drawn from a stock of premade signs. The creation of new sentences is a hallmark of HLC. This compositional means of building meaningful signals makes HLC adaptable to an extremely wide variety of communicative situations, and indeed the consensus among linguists is that humans can use language to talk about anything they wish to mention.

Productivity has the additional advantage of reducing burdens on an individual's memory by allowing messages to be created online in the communicative situation rather than simply stored in memory to be drawn upon when needed. The mix-and-match style of message formation also allows related meanings to share related forms, such as 'Henry ate three sandwiches at lunch' and 'Juliet ate three sandwiches at lunch', where the eating of three sandwiches at lunch is the same in both scenarios and is also represented by identical phrases in each sentence. This can also help in comprehension, especially of messages that have never before been heard.

A new word can be created whenever a certain series of sounds is used to express a certain semantic meaning by a critical mass of speakers. For example, many words for

[4] A bee's dance encodes a message that has never before been expressed when it finds food in a new location. However, I claim that since the bee is still restricted only to communicating about the locations and quality of food sources, the bee dance system is only creative in a limited sense.

new technologies work this way: 'mouse' is probably most commonly assigned the semantic meaning of 'computer pointing device' rather than 'small rodent' in many people's experience, even those who had to learn to apply this term to computer pointing devices well after they had learned to apply it to small rodents[5]. It is unlikely that everyone who learned to call the pointing device a 'mouse' when they first encountered it was explicitly told that 'mouse' was the name for the thing. Semantic extensions such as 'mouse' often rely on some sort of analogy or resemblance between an established meaning and a new meaning, such as the physical resemblance of the pointing device to the rodent, with the cord standing in for a tail.

However, when the new word involves coining a phonological word out of thin air, the word's first users must supply an explicit definition for it because there is not yet a consensus among the community of speakers just what semantic meaning should be assigned to this string of phonemes. Explicit coinages, in which words are made up wholesale with no recognizable link to existing words and morphemes, are relatively rare because of this very problem. However, they are clearly possible: 'blurb' and 'meme' are two examples (OED online).

Borrowings of words from foreign languages is much more common, especially when the concept that these words name is also borrowed and thus speakers have a new need to refer to that concept. The association between phonological word and semantic meaning that has already been established in the foreign language community in question and is known to speakers who speak both languages is enough to create a critical mass of consensus that can then spread to monolingual speakers of the borrowing language. To take a simple example, a flaky crescent-shaped pastry from France was introduced to the English-speaking world along with its French name, 'croissant'. A more subtle point is illustrated by the English-language borrowing of the German word 'schadenfreude', which is used to mean 'happiness at another's misfortune'. In the case of

[5] The word 'word' itself might be seen as ambiguous here between the ideas of "phonologically well-formed and cohesive string" and "sound-to-semantic-meaning mapping", so that while 'mouse' is the same phonological word in both cases, it represents a different sound-to-meaning mapping in the "device" sense than it does in the "rodent" sense.

'schadenfreude', the concept itself already existed in the minds of English speakers in the sense that monolingual English speakers could certainly experience happiness at another's misfortune before the borrowing of the German word--indeed, an emotion such as this is likely to be common if not universal among humans who live in social groups. English speakers could also, of course, describe this feeling with only English words and morphemes, just as I have done throughout this paragraph. However, what English lacked was one single word that named the emotion.

Productivity is a central characteristic of HLC. Having a code that can accommodate new ideas appears to be unique to humans, and is a tool that is extremely useful in accomplishing communication because it allows a sender to create a coded message that gives a receiver a very precise idea of what she intends to communicate. The mechanism by which coded elements function to convey speakers' messages is discussed more in chapter 5.

3.5 Miscellaneous features

3.5.1 Tradition

Hockett's feature 'tradition' describes the fact that some elements of the communicative system are learned through experience, a definition that excludes both explicit teaching and innate ability. Specifically, Hockett's definition of 'tradition' talks about initial acquisition of the system by the young of the species. In humans, the most obvious example of tradition in language acquisition is that babies learn the language to which they are exposed: babies in Zulu-speaking communities acquire Zulu, and so on. A system like the pheromones that many species use to indicate sexual availability, on the other hand, is entirely innate to the species and requires no learning whatsoever.

The respective roles of learning and innateness in human language have been debated, but the prevailing point of view is that some elements of the system are innate while others must be learned. The main learned element in the system is word meaning, which is to say the conventional symbol-to-denotation semantic mapping that allows a

speaker and hearer that share this mapping to use the code as useful evidence of the speaker's intentions and thus to communicate a message[6].

The fact that this mapping is learned rather than innate allows it to change within the course of a user's experience. Rapid language change would be impossible if all of the aspects of language were innate, because then change could only occur through genetic mutation and thus would happen on a timescale orders of magnitude greater than the lifetime of an individual. Sentence composition would still be possible if words and syntactic rules were all innate, but users would not be able to learn new forms or add new meanings to existing forms. With only strict and unchangeable mappings between words and their values, the utility of words would be severely limited: their meanings could not be extended or used metaphorically, and new words could not be borrowed. Any new concepts that speakers learned would be inexpressible. As such the feature of openness would be impossible without the feature of tradition. Thus, in that it is a prerequisite for openness, tradition is crucial to HLC. However, I will not consider tradition per se as a fundamental feature of HLC but consider it an element of openness.

3.5.2 Prevarication

A system that displays prevarication is one that allows its users to lie by means of that system. Prevarication is clearly a feature that describes a pragmatic usage of the system rather than an aspect of the code itself, because lying crucially depends on the assumption on the part of the receiver that the sender is inclined to be truthful. In contrast, most codes themselves have no particular qualities that help or hinder the act of lying, in that the code is simply a tool that a speaker uses for whatever pragmatic purpose she sees fit, lying being just one of these possible purposes. The bee dance code could be used to lie about the location or even existence of a food source, a bird call could be used to lie about an individual's sexual availability or the extent of its territory, and so on. The

[6] The set of mappings between sound and meaning are specific to the particular language that the baby acquires--French-acquiring babies must learn that [ʃɑ] maps to the same set of felines that English-acquiring babies learn as [kæt]--and thus the set of mappings can be roughly equated with the specific language that the baby is learning. Syntax also varies from language to language, although it is argued that at least some elements of human-language syntax are innate in humans.

only codes that cannot be used to lie are those that are a direct byproduct of phenomena they communicate, such as chemical odors linked to sexual availability.

While it is clear that many codes could, logically speaking, be used to lie, there is an important difference between that and the actual ability of an organism to use a code to lie. A sufficiently realistic motorized honeybee model can be manipulated by a human researcher to dance about a certain location (Michelsen et al. 1989), and receiver bees can be sent to the location that the model indicates, whether or not there is actually food at that location. Here the model can be said to have lied using the bee dance code. However, real bees have not been observed to lie about food sources. Of course, one possible if farfetched reason for this is that they have no motivation to do so, since each individual's well-being depends on that of the entire hive. However, if we assume that a bee dances about a particular location as an innate response to its finding food in that location, then the dance is tied to true information about its environment in a way that precludes the possibility of a lying bee.[7]

Some non-human animals do lie, such as birds that pretend to be injured in order to lure predators away from the young in their nests, nonpoisonous viceroy butterflies that mimic the warning coloration of poisonous monarch butterflies, and even vervet monkeys who have been observed making a predator call and then claiming food that other monkeys abandon when they react to the call (Cheney & Seyfarth 1990).

Burling (1993) argues that humans use two codes simultaneously in most communication, language and what he terms the "primate gesture-call system", which includes facial expressions, grunts, screams, laughter, and body postures. In comparing these two systems, he notes that most of the gesture-call system signs are innate and outside of voluntary control, which makes it very difficult for humans to use the gesture-call system to lie. While basic imitations of facial expressions, laughter, and so on are within most people's ability, only people such as trained actors who have practiced

[7] This line of reasoning comes dangerously close to an assertion that bees do not have intentions, which is not only virtually impossible to determine but also constitutes circular reasoning. I feel fairly safe in guessing that bees are unlikely to have anything resembling intentions, but I am aware of the sticky philosophical consequences of doing so.

suppressing their unconscious reactions such as facial expressions and substituting other, consciously produced signals can successfully lie with the gesture-call system. For example, when we say that someone is a "bad liar", it is usually not because he had trouble uttering an untruthful statement through language, but because his facial expressions and other gestures contradict the spoken message. Lying with the language code itself is trivially simple: anyone who knows a language can utter a sentence such as "I am happy" when she is actually sad, and so on. But the hearer is only convinced that the speaker is happy if he can determine that the speaker's gestures and facial expressions are congruent with her statement.

The fact that HLC can include lying is significant not because it indicates anything about the code itself but because of what it shows about a human's ability to judge the state of mind of another person. The observation that the human language code can be used to lie conflates two facts: that humans use the code, and that humans can lie. That human language is under more conscious control than, for example, bee dance or the human gesture-call system makes it a more useful tool than these systems for humans to use in lying, but for the present discussion about how features of the code facilitate HLC, there is little that is added by observing that the human language code can be used to lie.

3.5.3 Reflexiveness

A system that displays reflexiveness can be used to communicate about the system itself. This does not, however, appear to be a specialized phenomenon but simply one instance of displacement: to say that a communication system displays reflexiveness is simply to say that 'human language' is available as a topic of discussion just as any other entity is. Displacement, in turn, is possible because the code is built out of discrete signs, as discussed in chapter 4. That a bee cannot perform a dance about its dance code is part of the broader phenomenon that it cannot perform dances about much of anything other than locations in space, which in turn derives both from the limitations on 'vocabulary' created by its non-discrete code and its apparent inability to represent ideas in the first place. Reflexiveness is therefore not in itself a crucial feature of HLC.

3.5.4 Learnability

A system which can be acquired by adults is described as learnable in Hockett's (1960) formulation. This is different from the feature of tradition, because while tradition describes the word learning that goes on during first-language acquisition, learnability describes second-language acquisition. Human languages can be acquired by individuals who have aged past the critical period in which a normally-developing individual acquires a first language, although both the learning processes involved and the results of language learning in each case appear to be somewhat different. Most people who have passed the critical age for language acquisition, approximately those past puberty, can only acquire a new language imperfectly, and in cases in which people have reached adulthood without acquiring a first language, such as pathology or severe neglect, the evidence suggests that they will not be able to do so.

While learnability and tradition are different features, describing second- and first-language acquisition respectively, a learnable system must also be traditional: new languages are only learnable by adults because certain of their elements must be learned in the first place. It would, of course, be impossible to learn a system that is fully innate, just as it is impossible to learn how to become taller or to change one's blood type. Thus, tradition is necessary for learnability.

It also appears that adult second language acquisition relies on a different set of skills than children's acquisition of a native language during the critical period: while children seem genetically predisposed to acquire learned elements of the language but gradually lose this ability during adolescence, adults can compensate for the loss of this innate priming with more conscious and practiced domain-general skills such as the ability to think critically about language structures to perceive analogies in similar structures and expand them to new areas of the language.

How does the code's learnability facilitate HLC? An adult can learn a second human language because of domain-general learning abilities but also because of the ability to think about language as an abstract concept as discussed in the section about reflexiveness above. Just as with reflexiveness, the human language code is learnable

because humans use the code and humans can learn. Learnability is not fundamental to HLC.

3.6 Beyond Hockett

Hockett (1960) is an example of work within a descriptive paradigm that prevailed in linguistics until the so-called "cognitive revolution" began with the publication of Chomsky's (1957) *Syntactic Structures*. Broadly speaking, Hockett's list has two important shortcomings as seen through the cognitive paradigm favored in much of linguistics today. First, it lists features from a descriptive or observational point of view, presenting the results of whatever unaddressed internal structure may be present in the language system, without discussing the mental or cognitive mechanisms that create the phenomena. Second, it confounds the code with the act of communicating when features of the code--those that would be true of the system even if it were never used to transmit a message--with features of communication itself, which in general are not specific to the medium of the language code but could be displayed by non-language systems such as gesture and facial expression. Thus, as I attempt to sort, reevaluate, and make a more comprehensive version of a list of language features such as Hockett's, it is necessary to take some more recent developments in linguistics into account.

3.6.1 Recursion

A system that displays recursion is one that uses iterative rules to generate its outputs. Recursion is a feature of the human language code and has been argued to be the only feature unique to humans and unique to language (HCF 2002).

The feature of recursion is necessary for the infinitude of human language. For example, certain types of phrases can be embedded within phrases of the same type. A rule such as [S → X said that S] allows infinite embedding of sentences within other sentences, as in the following examples generated by this rule:

(1) Lawrence said that roses are red.

(2) Sara said that Lawrence said that roses are red.

(3) Robert said that Sara said that Lawrence said that roses are red. …

(4) Mary said that John said that Julie said that Martha said that Robert said that Sara said that Lawrence said that roses are red.

This example can clearly be extended indefinitely. The advantage that this brings to HLC is that it is impossible to run out of possible language structures, and thus there are at least no structural restrictions on what is expressible in HLC.

3.6.2 Reference

The second cognitive feature under discussion is reference. Reference is the act of pointing out a particular entity in order to bring it to the attention of a receiver. Reference is a communicative feature of HLC, which is to say that it is a feature of using language to communicate with other humans, rather than a feature of the language system itself; in fact, the act of reference can be done as easily by pointing a finger at something as by uttering 'that building' or 'the President'. Linguistic reference, of course, can take advantage of displacement to allow speakers to refer to absent and abstract entities to which they cannot physically point. The human language code contains structures that facilitate linguistic reference, and in fact the semantic meanings of words such as 'the' and 'it' are nothing more than their referential values (Gundel, Hedberg, & Zacharski 1993). However, the act itself is not dependent on the use of a code, but rather the ability of a sender to anticipate what linguistic or non-linguistic action will successfully draw the attention of a receiver to the entity to which she wishes to refer. The mechanism by which linguistic reference is accomplished is discussed further in chapter 5.

4 Discreteness and duality (combination)

The thesis of this dissertation is that HLC is distinguished from other communication systems by a structural property, namely a discrete combinatorial system, and a communicative property, the ability of users of the system to infer the mental states of other users. As mentioned in chapter 3, the features "discreteness" and "duality of patterning" that Hockett lists are a rough characterization of a discrete combinatorial system. However, given this feature's importance to HLC, a separate and more detailed discussion of discrete combination is warranted.

A discrete combinatorial system, of course, has two theoretically separate elements: the existence of discrete units within the system, and the combination of these units. I discuss each of these in turn.

4.1 Discreteness

Discreteness is the quality of being composed of units that exist as definably separate entities rather than as points along a continuum. The red, green, and amber of traffic lights constitute a discrete system, with no overlap among the encoded states 'stop', 'clear the intersection', and 'go'. Contrast this with, for example, the use of the same colors that can be interpreted as indicating the ripeness of a pepper: the pepper starts green, ripens to yellow, and then finally to red, with the degree along the continuum reflecting the degree of ripeness. Yellowish-green shows that the pepper is riper than green but not as ripe as yellow, and so on. This is characteristic of a continuous system. Traffic lights, however, change from green to yellow with no yellowish-green signal that corresponds to an instruction halfway between 'go' and 'clear the intersection'.

The classic example of discreteness in human language is the human ability to perceive discrete phonemes within the continuous medium of the acoustic signal. This categorical perception is a trick of interpretation, in which human listeners perceive boundaries within the signal that do not exist within the physical medium. Hockett's own description of this ability captures the basic idea of discreteness and clearly explains this complicated phenomenon, and is elegant enough to be worth quoting at length:

> Any utterance in any language consists of an arrangement of certain basic signaling units called phonemes, of which a given language has a definite and finite stock. Phonemes are not sounds, but ranges of sound quarried by quantization out of the whole multidimensional continuum of physiologically possible vocal sound. In different languages this quarrying yields different sets of phonemes--a difference of sound that is functional in one language may or may not be in another. We can illustrate both the quantizing and the ways in which languages differ by considering just one physiologically given dimension: voicing, the vibration of the vocal cords. In English the scale of degrees of strength of voicing is quantized into two subregions. Thus the initial consonants of the words *pat* and *bat* differ only as to voicing: the *p* is most typically voiceless, the *b* most typically voiced. Yet some occurrences of *p* are slightly voiced, and in some occurrences of *b* the voicing is very weak. But if a hearer hears something that is not clearly marked as a *p* or as a *b*, it is only this two-way ambiguity that he has to try to resolve. There is no further alternative, except to leave the ambiguity unresolved--and that is not an alternative within the system, but a breakdown of the system. In some other languages, such as Menomini, the scale of degrees of strength of voicing is not quantized into smaller contrasting regions at all: a speaker of Menomini does not at first hear any difference between English *pat* and *bat*. In still other languages, such as Hindi, degree of strength of voicing is quantized into two regions, and exact timing of onset of voicing is likewise quantized into two, to yield four contrasting units where English has two and Menomini only one. (Hockett 1960: 146-147)

Human listeners thus parse the continuous sound space into perceptually discrete units of sound, conditioned as described above by the boundaries by which a given speaker's native language divides the continuous signal. From this array of perceived sound units--called phonemes--human languages build morphemes and words.

From the point of view of semantics, morphemes are units of meaning, out of one or more of which words are built. From the point of view of phonology, morphemes are ordered clusters of phonemes that conform to the limits on sound patterning that each language requires. Combining these two definitions, a morpheme can be described as a chunk of sound to which a unitary meaning is assigned. Swapping one phoneme for another changes the morpheme entirely: the single-morpheme word '_at' made with an initial sound that is acoustically ambiguous between 'p' and 'b' can only be perceived as a strange instance of either the word 'pat' or the word 'bat'. In contrast to the example

above about a yellowish-green pepper being halfway between the ripeness states indicated by green and yellow respectively, a word that sounds halfway between 'pat' and 'bat' will not somehow be given a meaning halfway between the meanings of those two words (if such an idea could even be usefully considered).

4.2 Combination

The second element of a discrete combinatorial system, along with discrete units, is of course a combinatorial system. For convenience I will use the term 'combination' to name the feature of having a combinatorial system.

The human language code is in fact built of two separate combinatorial systems: the discrete set of phonemes combine into morphemes, and the discrete set of morphemes are the building blocks of phrases and sentences. The sentence is the largest structural unit of natural language, a sign whose value is an entire proposition. When phonemes are combined into morphemes and morphemes are combined into words and sentences, the identity of these individual elements remains. The meaning of a sentence depends of course on its syntactic structure, illustrated trivially by the fact that "dog bites man" and "man bites dog" are sentences with very different meanings. However, the meanings of "dog", "bites", and "man" are identifiable elements within the meaning of each sentence, and likewise [d], [ɑ], [g], [b], and the rest of the phonemes in each string remain individually perceptible.

Some notes about terminology. A minor point to make here regards the conflation of morphemes and words in the present discussion. A morpheme is the smallest unit of meaning in a given language. Words, from the point of view of phonology and morphology, are composed of one or more morphemes, each of which in turn is a series of phonemes. A word is a morpheme unit that can stand alone within a syntactic structure[8]. Thus, a sentence can be a structured sequence of single morphemes, or can consist of polymorphemic words that themselves combine into a sentence. The difference, for the present purposes, is negligible: the idea that one level of combination

[8] Another view is that the term 'word' names a form-to-meaning mapping rather than simply a structural form; this usage is not relevant in the immediate context but is discussed later in the dissertation.

is at work when morphemes are combined into sentences does not depend on whether some morphemes are first grouped together into words that then feed the generation of sentences.

A more important point is the difference between, and sometimes conflation of, 'phrase' and 'sentence' in the present discussion. The idea that sentences are not simply ordered strings of words but that the words are grouped more basically into phrases that represent hierarchically ordered elements within a sentence is a key assumption on which modern syntactic theory is based. Given the pre-generative tradition from which Hockett (1960) arose, it is unlikely that much consideration was given to the internal structure of sentences other than the basic fact that they are composed of multiple morphemes--whether they have internal phrase structure or are simply an ordered string of morphemes does not appear to have been given any particular treatment in that work.

This appears to be an important oversight in Hockett's (1960) treatment of the levels of combination in human language. However, with only the conflation, justified momentarily, of 'phrase' with 'sentence', I agree with Hockett's (1960) premise that there are only two combinatorial steps that are relevant to a description of the language structure on a fundamental level. The justification for this is that only two levels of combination result in a unit of language that is qualitatively different from the units of which it is composed. Phonemes are meaningless units which, when combined into morphemes, are assigned a semantic meaning. The addition of meaning makes morphemes are qualitatively different from the phonemes out of which they are built. The combination of morphemes according to the rules of syntax likewise results in a qualitatively different unit, the phrase. Both morphemes and phrases are carriers of meaning, but the significant difference between the two is that the mappings between word forms and their meanings are stored in memory, whereas the meanings of sentences are derived online from the stored meanings of the element morphemes and the relationships between them as defined by the syntactic rules operating within the phrase. Thus, the step from morphemes to phrases is another step whose output is qualitatively different from the units of its input.

Returning to the discussion of 'phrase' and 'sentence': I conflate these and thus uphold Hockett's (1960) assertion that combination in human language occurs on two significant levels by noting that a sentence is merely a specialized type of phrase, one that can stand alone to express a proposition. The relationship between phrases and sentences is analogous to that of morphemes and words, wherein the latter is a terminal unit resulting from the combination of one or more of the former units, and whose criterion for termination is the achieving of a form which can stand alone as a finished unit. A word is thus a type of morpheme unit, and a sentence is a type of phrase: neither is qualitatively different from its component parts.

It is the fact that the human language code has these two combinatorial systems that Hockett noted when he identified the feature of "duality of patterning"; however, I claim that the fact that human language has two levels rather than zero or one or three is less important than the fact that it uses combination at all.

The system of phonemes, morphemes, and sentences is not logically necessary for the expression of propositions. A code could have any number of tiers of patterning: along with the possibility of having indefinitely many, a given code could function to express propositions with one or no intermediate tiers between its atomic units and its complete propositions. For example, in a system with no intermediate layers, one sound, handshape, written symbol, or other unit of a code could represent the content of one human-language sentence, something like 'The weather is fine today' or 'The director's office is located at the end of the sixth-floor hallway of the main building'. Adding one intermediate layer between basic unit and proposition would yield a code in which a string of units represents a proposition, such as [aaaaa] to encode 'Please call me at your earliest convenience' and [aaaab] for 'First-class postage stamps will increase in price by three cents six weeks from now', or any other expressible message.

What does it mean for communication that the human language code builds signals out of two levels of patterning? Communication would theoretically be possible in a system in which one unit represents one proposition and cannot be further analyzed into smaller meaningful parts. Indeed, many nonhuman animal communication systems

work this way: vervet monkeys make one call when a leopard is present, another for eagles and yet another for snakes (Cheney & Seyfarth 1990).

In a system like that of vervet monkeys, any new signs added to the code must be different enough from existing signs that they are reliably produced and perceived as distinct from any other signal. This poses an obvious problem in that the continuous range of physical signals that are possible--the sounds that can be created by the vocal apparatus, or the gestures that can be made with the hands and body--can in practice only be divided so finely before the units become impossible to produce and perceive accurately. As mentioned above, a system for HLC that operated this way would assign one entire message to one basic unit, which in the case of HLC is a phoneme, such as a [p] or a [b]. The more signs that are added to the repertoire, the smaller the differences become between each sign and the sign that it most closely physically resembles. A system of this kind becomes impractical when signs become so close together in the physical space that it is functionally impossible to produce and perceive individual signs that can be reliably distinguished from their nearest physical neighbors: a [p]-like sound could be produced with slightly earlier voicing, and one with earlier voicing still, and so on, but at some point the system reaches a saturation point past which it is difficult to tell exactly what thin slice of the continuum a speaker is trying to produce. In addition, every message that a sender might wish to express would have to be pre-encoded and previously memorized in order to be available for use, as is true in logographic writing systems. It would be difficult for new signs to be created, because nothing in the signal itself gives any indication of that signal's value.

Adding one intermediate level of combinations of discrete units allows a much smaller, and thus more easily mutually distinguished, set of physical signals to form the utterances that express the messages that humans communicate to each other. Phonemes would be individually meaningless units in a system like this, and a semantic value would be assigned to a string of phonemes as a whole, either without regard to order or with

order as a meaningful proto-syntactic component[9]. The system could support the expression of an indeterminate number of messages by allowing strings of indeterminate length.

A system with one level of patterning such as this is more feasible for perception and production than one with a direct map from each unitary sign to an entire message. Still, this system suffers from the same difficulty as one with no intermediary levels, in that there is no link at all between the form of an utterance and its semantic content. If no particular string of units encodes 'queen' or 'England', then the signals encoding two messages such as 'Elizabeth is the queen of England' and 'Elizabeth is not the queen of England' would bear no more likeness to one another than those of any two randomly selected signals, and likewise the resemblance between strings such as [aaaaa] and [aaaab] would not be relevant to the messages that they encode.

Adding a second level of patterning alleviates this difficulty. Meaningless phonemes combine to make morphemes, the minimal semantic unit, and these individually meaningful units recombine according to the rules of morphosyntax to form phrases. These phrases take their meaning from the composite meanings of their elements, creating a resemblance between forms when meanings are related: the sentences 'Elizabeth is the queen of England' and 'Elizabeth is not the queen of England' share identical forms other than the one morpheme 'not', and because forms map to meanings in an identical way in both sentences, the meanings are also closely related. Some tens of thousands of words must still be memorized individually, but a speaker can form phrases based on these words and the rules of morphosyntax.

Thus, discreteness plays two roles in fostering HLC. Discrete units ease the burdens on production and perception that a continuous array of signals would create. In addition, discrete units and rules for combining them ease memory burdens as much as possible by allowing messages to be created online rather than having to be created and

[9] Spoken or signed natural languages are constrained to some linear order of the presentation of their individual signs due to the vocal-auditory or visual-gestural modalities in which they are presented. However, the order of units can be significant or not: ordered sets would be mathematical permuations of the constituent elements where unordered sets are mathematical combinations.

stored in memory before they are available for use in HLC. This is discussed in more detail in chapter 6.

5 Linguistic pragmatics

The ability to make inferences about another speaker's communicative intentions is one of two elements that I argue are central to the essence of HLC. In this chapter I describe, and then elaborate upon, a prevailing theory in the field of linguistic pragmatics on which I base my argument. Specifically, I provide a technical explanation of HLC by showing the specific role that the language structure plays in providing a basis for inference. I show that the structure alone is not enough to provide reliable transmission of a speaker's meaning, but that the structure's role is to provide evidence of a speaker's intended meaning by narrowing down the range of possible meanings until the remaining indeterminacy can be bridged by an inferential process.

I take as a premise of this paper that the fundamental means of communication between humans is ostensive-inferential communication as described by Sperber & Wilson (1986/1995) and discussed below. I extend this premise by claiming that HLC is accomplished by combining structural elements with this inference-making ability. I begin by reviewing and discussing the basic mechanism through which HLC is accomplished, and then develop the ideas that Sperber & Wilson (1986/1995) put forth in order to demonstrate the role that the language structure plays in HLC.

5.1 Pragmatics background

Linguistic pragmatics is the field that deals with the gap between "word meaning" and "sentence meaning" on the one hand and "speaker meaning" on the other: more formally, it is the study of the means by which messages are transmitted through language beyond simply the information contained within the language code itself.

A crucial work in establishing the field of linguistic pragmatics is Grice (1975), which serves as the starting point for Sperber & Wilson (1986/1995), on which my own research is based. Grice originates the concept of *implicature* to account for the communication of messages whose content is not explicitly contained within the words and sentences that compose a given utterance. In Grice's conception, implicatures arise

from the Cooperative Principle and its constituent maxims. The Cooperative Principle states that in order to communicate successfully, speakers obey a tacit requirement to cooperate with their conversational partners in accomplishing a communicative goal. The Cooperative Principle breaks down into four more specific rules that Grice (1975) terms "maxims": the maxim of Quality, which requires that communicators say only what they have evidence to believe is true; the maxim of Quantity, which requires them to provide the appropriate amount of information for the situation; the maxim of Manner, which requires brevity and orderliness; and the maxim of Relation, which requires speakers to make their conversational contribution relevant to the situation.

Grice's (1975) explanation of implicature is thus based on a tacit social agreement. Because hearers assume that speakers are adhering to the Cooperative Principle, they tend to make certain assumptions about the speaker's intentions. For example, if a speaker says that she has five dollars, a hearer will tend to assume that the speaker has exactly five and no more based on his assumption that she is giving as much information as is required. This assumption happens even though having fie dollars is not inconsistent with having seven, since having seven entails having five. Similarly, when a speaker says, for example, that X might be the case, her hearer will tend to assume that the speaker does not know for certain whether X is the case--not because she has said so outright, but because he assumes that if she was certain that X is the case she would not have weakened her assertion by using 'might'. Once again, the hearer tends to assume this even though it is true that if X is definitely the case, it is also true that it might be the case and therefore a speaker who says 'might' but knows for certain is not lying.

Speakers are not, of course, constrained by the Cooperative Principle, and can choose to disobey it either because they do not wish to cooperate with their interlocutor or in order to imply something that is not directly said. This is what Grice terms flouting a maxim. For example, a speaker who abruptly changes the subject of conversation is violating the expectation that her contribution be relevant to the situation at hand. However, aside from the content of the new subject she introduces, she may also be implicating that she does not wish to continue discussing the previous topic.

Whether a speaker explicitly flouts a maxim for effect or simply takes advantage of a hearer's likely implicature-based assumptions to communicate a message, the process of implicature for Grice derives from assumption that the Cooperative Principle and its maxims are being obeyed. Crucially for my work, Grice clearly delineates the difference between what is "said" and what is implicated by explaining, "In the sense in which I am using the word *say*, I intend what someone has said to be closely related to the conventional meaning of the words (the sentence) he has uttered" (1975:44). Implicit in this is the assumption that what is "said" and what is implicated are relatively distinct categories, and that communication can and often does happen without the implicature.

The work of Sperber & Wilson (1986/1995) is rooted in Grice's (1975) ideas, but their Relevance Theory reconstrues the problem of implicature as a cognitive, rather than a social, phenomenon, as they explain:

> "Ostensive behaviour provides evidence of one's thoughts. It succeeds in doing so because it implies a guarantee of relevance. It implies such a guarantee because humans automatically turn their attention to what seems most relevant to them. The main thesis of this book is that an act of ostension caries a guarantee of relevance, and that this fact--which we will call the *principle of relevance*--makes manifest the intention behind the ostension. We believe that it is this principle of relevance that is needed to make the inferential model of communication explanatory." (1986/1995:50)

Thus, Sperber and Wilson's (1986/1995) theory is based on the idea that every act of communication involves some inferential enrichment that goes beyond what is linguistically encoded. Sperber and Wilson, in discussing the possibility of devising a general theory of communication, comment that "most sentences can be used to convey an infinite number of different thoughts" and that "different utterances of the same sentence may differ in their interpretation; and indeed they usually do" (1986/1995:9). In addition, they state that "…the grammar can only help determine the possibilities of interpretation" (1986/1995:10).

It is this idea of Sperber and Wilson's (1986/1995) that I develop in this chapter: I first demonstrate their idea that sentences can be used to convey an unlimited set of

intended messages, and then expand on their idea that the grammar helps narrow down the possibilities of interpretation by proposing that the language structure provides the shared set of conventions necessary to allow messages to be successfully sent and received.

Sperber & Wilson (1986/1995) describe two models of communication: the code model, in which information is encoded by the sender, transmitted, and decoded by the receiver; and the inferential model, in which the information encoded and transmitted, along with information about the sender's intention to communicate, serves as evidence for the sender's intended message. In the inferential model, the receiver processes the combination of the coded information and the inferred communicative intention and chooses an interpretation that is maximally relevant given the context of the utterance. It is this model that is ultimately preferred, and which I will develop later in this chapter. First, though, it is important to examine the code model of communication, both to review the reasons why it is dispreferred in explaining communication, and to provide some basic exposition about what the language code is *not* doing in the service of communication, so that I can later discuss what role it *does* play in fostering HLC.

5.2The code model vs. the inferential model

As Sperber & Wilson (1986/1995) describe, a conventional model of human linguistic communication which has long had currency is that in which information is simply encoded by the sender, transmitted as a code, and decoded by the receiver. This model is appropriate for much of the world's varied communication systems, from the body postures of dogs to the red, amber, and green of traffic signals. Crucially, a system which functions by coding and decoding information relies on the system's having minimal or no indeterminacy, because there is no provision in the code for nuance or interpretation.

Take the example of traffic lights. On an initial pass, the traffic light example illustrates the code model of communication, in which the entire message is encoded, sent, and decoded by means of some conventional signs. In fact the situation of a driver approaching a traffic light is much more complex than just this, but for the sake of

illustration we can imagine a driver who wishes to proceed straight across an intersection but is unsure whether she has the right of way to pass through the intersection despite oncoming traffic, or whether she must stop to let oncoming traffic pass before she proceeds. For the sake of simplification we exclude a number of other possible behaviors on the driver's part, including turning at the intersection or indeed such things as backing up, getting out of the car and walking, or any of the infinite possibilities actually open to her at any moment. The situation of the driver thus approaching an intersection is reduced to a binary indeterminate state: proceed, or stop.

In a situation thus simplified, the code model applies perfectly to the traffic lights: Green encodes all of and only the message *You have the right and the obligation to go forward*, and red encodes *You must stop at the intersection for the duration of the red light before proceeding*. This code works as well as it does in this situation because there are only two possible messages to be sent in the context of this traffic intersection: there are no intermediate meanings between the two states nor any potential new meanings that would need a means of expression. Thus the conventional or 'semantic' meaning assigned to the red and green lights alone is enough to resolve the indeterminacy inherent to the situation.

The code model of HLC holds that the speaker's language, such as English or Cherokee, is essentially a more complex equivalent to the colored lights used to convey right-of-way information at a crossroads. Although the traffic lights have only three symbols to transmit while English and Cherokee can supply a much larger, even infinite number, the code model of HLC holds that the same principles govern both types of system.

The language code is certainly useful. Human linguistic communication relies on our ability to obtain information from various aspects of the language transmission. Much of the information conveyed in a given utterance is recovered from aspects of the language structure itself: the sounds of the language and the way they combine (phonetics and phonology); the way words are built from units of meaning (morphology); the way words combine into phrases and sentences (syntax), and the way words and sentences are conventionally associated with real-world meanings (semantics).

However, the structural information itself is not enough to determine only one unambiguous interpretation of any given message. Examples of indeterminacy abound. One type has to do with the structure itself. Lexical and structural ambiguity are sources of indeterminacy: in the string "I saw her duck", the word "duck" is ambiguous between "waterfowl" and "lower one's body to avoid an obstacle", and the whole sentence is ambiguous between the meanings of "I witnessed her lowering her body" and "I viewed her pet waterfowl" due to the different syntactic structures underlying those two sentences.

Even words that do not represent obvious homonyms such as "duck" can be ambiguous. The word "Russia", for example, usually names the nation whose capital is Moscow. If I requested that you "show me Russia", I could be asking for a tour of that country, in which case the word "Russia" does name the nation. However, "show me Russia" could be used if I wanted you to show me a map of Russia; if I wanted you to show me a puzzle piece that belonged to a world map puzzle; if you and I were playing a conquer-the-world board game; if you were a travel agent and I wanted to see prices for airline tickets to that country; and many more situations. Thus, although we would not typically say that the word "Russia" is ambiguous, in fact it can be used to name an unlimited number of items as these examples demonstrate.

Lexical and structural ambiguity are certainly a problem for the code model, but they are a relatively small one. Much more important is a pragmatic type of ambiguity, in which whole sentences can also be used in situations in which a speaker's meaning in uttering them goes further than the conventional meaning of the sentence. A common example is in indirect constructions often used for the sake of politeness. I might ask, "Would you mind giving me a ride to the airport?", which is a yes/no question about your preferences, when what I really intend is to request that you do so.

Sperber & Wilson (1986/1995) provide and discuss similar examples:

(5) Do you know what time it is?

(6) Coffee would keep me awake. (Sperber & Wilson 1986/1995: 11, (6) and (7))

As Sperber & Wilson explain, (5) could be a question about the hearer's knowledge or it could be a hint, perhaps that it is time for the hearer to leave for work or that a telephone call has caught the speaker while she was in bed. (6) could be either the acceptance or refusal of an offer of coffee depending on whether the speaker wants to stay awake or not, but it could also be uttered when the speaker is simply listing things that would prevent her from sleeping, naming the effects of coffee, and so on.

Now, identical sentences[10] with different interpretations present the same problem for the code model of HLC that identical traffic signals with different interpretations would for the traffic light system. If a green light could indicate either "stop" or "go", it is clear that the traffic light communication system would cease to convey any useful information at all. With traffic lights, this problem is avoided by giving each signal a single, fully determined meaning. In HLC, on the other hand, communication thrives despite this multiple signal-meaning mapping. The code model therefore fails to provide an acceptable explanation for HLC.

HLC is clearly very useful for transmitting information from a sender to a receiver despite the fact that it any number of possible meanings to attach to any particular signal. As such, it is clear that a model other than the code model is necessary to explain its success. Sperber & Wilson (1986/1995) develop an idea set forth by Grice (1975) that explains human communication, linguistic and otherwise, by positing that the crucial factor in the transmission of a message is that the sender intends the receiver to recognize her intention to communicate, and that the recognition of this intention is itself the means by which the message is transmitted. Sperber & Wilson continue by defining "ostensive-inferential communication" as follows:

[10] The term 'sentence' is most often used to name a unit of the code, a syntactic object, which has a conventional meaning but not a speaker-meaning. Thus if two different conventional meanings can be assigned, the objects in question are two different sentences (to be contrasted with the common usage that they are "the same sentence with two different interpretations").

> "…the communicator produces a stimulus which makes it mutually manifest to communicator and audience that the communicator intends, by means of this stimulus, to make manifest or more manifest to the audience a set of assumptions I." (1995: 63)

So, in ostensive-inferential communication, which includes human linguistic communication, the receiver's recognition of the sender's communicative intention is the fundamental element of the act of communication. The role of the structure is to make salient to the receiver a number of possible meanings, and the role of inference is to allow the receiver to choose the intended meaning from among these possible meanings based on some shared context or convention. Thus, although inference may be carrying much of the load of communication, an entirely structureless communication would amount to telepathy. It is clear that some sort of structure is a crucial part of HLC. Where, then, does this leave an examination of the structural elements of HLC? I base my work on a passage from Sperber & Wilson (1986/1995), quoted here:

> "Verbal communication is a complex form of communication. Linguistic coding and decoding is involved, but the linguistic meaning of an uttered sentence falls short of encoding what the speaker means: it merely helps the audience infer what she means. The output of decoding is correctly treated by the audience as a piece of evidence about the communicator's intentions. In other words, a coding-decoding process is subservient to a Gricean inferential process." (1986/1995:27)

Crucially, this ostensive-inferential communication model applies as well to linguistic as it does to non-linguistic human communication. Whether the information being conveyed is transmitted through the language system or some other coded or non-coded channel is irrelevant to the basic idea of ostensive-inferential communication. Following Grice, Sperber & Wilson (1986/1995) assume that this stimulus need not take the form of any conventionally coded information at all, with the example that it is possible for one person to show another that she is sick simply by showing her interlocutor a bottle of aspirin (1995: 25-26). There is certainly no coded link between the showing of a bottle of aspirin and the meaning 'I am sick'--showing a bottle of aspirin

might equally convey ‘I have been to the drugstore today’, or ‘This is the item responsible for the rattling sound you hear when I walk’, to name just two possibilities--yet in a given context the message may be appropriately sent and received.

5.3 Expanding the inferential model

So far, Sperber & Wilson’s (1986/1995) discussion has demonstrated the shortcomings of the code model of communication and proposed an alternative, the inferential model, to better explain the facts of human communication. Their demonstration that the language structure alone does not convey enough information for real communication, and that communication can also be accomplished by non-coded means such as showing a bottle of aspirin, provides a crucial contribution to the overall problem of explaining HLC. However, a problem that Sperber & Wilson (1986/1995) leave unaddressed is the role that the language structure does play in HLC. While saying ‘I am sick’ and showing a bottle of aspirin can both convey the meaning *I am sick*, it seems clear even on a naïve level that conveying this message with the utterance ‘I am sick’ is less likely to convey something other than the speaker’s intended message and more likely to be understood in a wide variety of situations as conveying *I am sick* than the showing of a bottle of pills. Thus, the next step in developing the study of HLC is to determine what it is about the language structure that makes it more useful than other means of conveying the same message.

The basic problem in communication is one of reducing indeterminacy: narrowing down the possible messages that a sender could communicate until just the most relevant message, and therefore that most likely to be intended, is left. We have seen above that the language structure cannot necessarily accomplish this task, due to ambiguous lexical items and syntactic structures. It is also clear that non-codified gestures alone, such as the showing of the pill bottle, are not sufficient to reliably communicate any specific message. Indeed, the party game of charades derives its entertainment value from the difficulty and imprecision of this sort of communication.

All of this indeterminacy might lead an observer to believe that human linguistic communication must be a hit-or-miss affair with limited utility. Indeed, if one were to

imagine an artificial intelligence whose processing ability were limited to only the language structure, it seems clear that such a machine would fail upon its first encounter with 'bank' or 'the wind blew over the small plane' or 'here'[11]. Yet human linguistic communication clearly does not systematically fail at these items and indeed can be used successfully to convey subtle shades of meaning.

Since, as Sperber & Wilson (1986/1995) have convincingly demonstrated, human communication relies on a shared set of assumptions between sender and receiver; and since this communication is clearly much easier and more precise when the language structure is used than by other means; it stands to reason that the language structure is providing the conversational participants with a shared set of assumptions. This, as I will demonstrate, is exactly the role that the language structure serves in HLC.

A metaphor will serve to illustrate this point. The difference between the sender's intended meaning and the meaning that the receiver understands can be represented as a gap between the two. Although a hearer often has some idea of what a speaker might say just because some utterances are more likely than others in certain situations, he almost certainly does not know exactly what the speaker is going to say. The receiver ultimately bridges the gap by making an inference about the sender's intended meaning based on some level of shared context or convention.

If two people share appropriate background knowledge, a message such as, for example, *Meet me at the bank* may be successfully transmitted even if no utterance with that conventional meaning is spoken at all. Imagine the following dialogue:

(7) A: Another letter came today.

(8) B: Three o'clock, then?

If we know that part of B's communicated message is *Meet me at the bank*, we can assume from this exchange that A and B have a good deal of shared knowledge and

[11] 'Bank' because this phonetic word is mapped to both *financial institution* and *river's edge*; 'blew over' could mean *caused gusts above* or *overturned*; and 'here' depends, roughly, on the speaker's location.

assumptions: that the letter came from the bank, that it requests that both parties come to the bank together, and so on, and thus that inference based on this shared information is enough to convey *Meet me at the bank*[12]. However, pure inference is not enough to bridge this gap in most cases in which the meaning *Meet me at the bank* is the intended content in a communicative exchange. If A and B in the example above were strangers chatting in line at a coffee shop, for example, and A wanted to arrange a meeting with B at some particular bank, A's utterance above, 'another letter came today', would certainly not suffice to communicate that meaning. The shared convention that bridges the gap has to come from somewhere else. The most obvious source for it is in the conventional meanings of the words of a given language. For the present purposes an individual language such as Polish or Farsi, or in the case of the examples given here, English, can be thought of as a set of conventions that relate forms to meanings.

It is crucial to note that even when a speaker utters 'Meet me at the bank' when she intends the message *Meet me at the bank*, the mechanism by which this utterance conveys this meaning is still technically inferential: if 'Meet me at the bank' invoked an inside joke or some other association for a speaker and a hearer, the intended meaning of the utterance could be quite different, and thus 'Meet me at the bank' does not automatically transmit the message *Meet me at the bank*. However, the most likely intended meaning for a given utterance is the conventional meaning of the words and phrases of which it is composed (Sperber & Wilson 1986/1995:22). When this is the case, the gap that must be bridged by inference is extremely small: the precision with which the structure of a given human language allows its users to encode meanings allows a sender to create an utterance that represents something very similar to her intended meaning.

At the moment before our hypothetical speaker utters the first sound of 'Meet me at the bank', the gap between intended meaning and received meaning is large enough to be unbridgeable--it would take an act of telepathy for a receiver to understand a sender's intended meaning before the sender even began speaking. The conventional meaning

[12] The word 'bank' is ambiguous between 'financial institution' and 'edge of a river'. However,

associated with the utterance 'Meet me at the bank', however, bridges a large part of this gap by allowing sender and receiver to draw on the shared conventions of whatever language they are speaking: rather than guessing that the message could be anything at all, the receiver can safely assume that the sender intends to convey a message about meeting the speaker at some particular bank. At this point the gap to be bridged is much smaller than it was at first, and the receiver has a much better chance of selecting the intended message from the possible meanings that a speaker might have for the phrase 'Meet me at the bank'.

At the very beginning of the sender's utterance, at the instant that it becomes apparent to the receiver that the sender intends to send him a message, the message's indeterminacy is at its peak: there is no evidence yet for what the message is, only that one is being sent. As the auditory signal is sent from sender to receiver the receiver begins to narrow down the sender's possible intentions based on what he hears. Once the word 'meet' has been uttered, the sender can now infer based on the semantic meaning of the word 'meet' that the message has to do with *meeting*.[13] He might also begin to represent the syntax of the utterance by recognizing that a sentence-initial tensed verb often begins a command or request, so that he can now at least tentatively assume that the sender is requesting or commanding him to perform the act of *meeting*. As 'me', 'at', and so on are added to the string, along with whatever potential syntactic arrangement can be guessed at from the words heard so far, the indeterminacy of the message is reduced from its initial totality to a much narrower range of likely meanings. The completed utterance 'Meet me at the bank' does not fully resolve the indeterminacy between the two possible semantic meanings that the sender might intend, but the result of her having sent this message is that, from an infinite possibility of meanings, the

[13] There are of course many things that one can mean with the word 'meet', including the two different common meanings 'make the acquaintance of' and 'get together with'. The phonetic word [mit] could also be 'meat' as in 'animal flesh'; 'mete' as in 'to dole out'; 'meet' as in 'proper'; 'meet' as in 'athletic match'; an example of onomatopoeia; a word in a foreign language; and so on. In the example above I set these aside for the sake of simplification.

sender has been able on the basis of the language structure alone to weed out nearly all of them.

It may seem that other than deciding whether a financial bank or a river bank is meant and possibly filling in some missing details about exactly what time and in what part of the bank the meeting is to occur, an utterance such as 'Meet me at the bank' is straightforward enough that it can only be assigned one basic meaning. However, the conventional meaning of the utterance still only provides evidence about the sender's intentions that feeds the receiver's inference-making ability. In fact, any utterance can be successfully used to convey any meaning at all, if the sender and receiver share the necessary convention: the utterance may form part of a prearranged code, or may invoke an inside joke that each of the two knows that the other will remember.

The ordinariness of these examples suggests that many more such items could be found, and in fact it is difficult if not impossible to find sentences which have only one possible interpretation. Even sentences expressing mathematical truths such as "one plus one equals two" can be used in the right context to mean something else. In fact, I claim that any utterance can be used to convey any proposition at all given the appropriate context. This claim is crucial to the claim that inference is one of the key elements of HLC. I demonstrate this as follows.

Inside jokes between close friends are the most common example of otherwise outlandish-seeming links between form and intended meaning, wherein the participants in the conversation share a set of background information and know that the others share this as well, and will use it to understand a speaker's meaning.

Taking the utterance "one plus one equals two", for example, it is possible to imagine a scenario in which this utterance is interpreted to mean any arbitrary proposition, for example *John is wearing a blue hat*. Imagine that you and I have a mutual friend, John, who enters the room wearing a blue hat. This reminds me of a certain professor whose class you and I both had as undergraduates, who also frequently wore a blue hat. This professor had a catch phrase, "one plus one equals two", that he often uttered, perhaps as a sarcastic reply in cases in which a student made a point that he considered to be trivially obvious. Now when our friend John enters with his blue hat, I

am reminded of our former professor and his blue hat, and in order to bring this professor to your mind as well I mention his catch phrase, "one plus one equals two". You hear this, think of the professor, think of his blue hat, notice John's blue hat, realize the resemblance, and smile knowingly. What I have done is point out to you that John is wearing a blue hat, not by uttering any sentence that conveys that proposition in its conventional form, but with a logically unrelated proposition that I knew would lead you, based on the immediate physical context (John and his hat in the room) and our shared background knowledge (the memory of our former professor), to understand my intended meaning.

If this example seems too farfetched to be plausible, consider another: You and I see our friend John enter the room with his hair sticking up all over, and I nudge you and whisper "E equals mc squared". The communicated meaning is something like *John looks like Albert Einstein today*, although the conventional meaning of the utterance's coded content is quite different. In this case the necessary context that bridges the gap is the common cultural knowledge of Einstein's appearance and of his famous equation rather than an inside joke unique to only the two of us, but the principle is the same. That these and many similarly derived instances of communication can occur shows that the communication of a speaker's intended meaning is clearly not just a function of the conventional meanings of the words and phrases that compose her utterance.

5.4 Chapter summary

This chapter provides a detailed explanation of the mechanism by which a speaker's intended meaning is successfully transmitted to and understood by a hearer. The structure itself does not carry enough information for a message to be reliably interpreted, because of ambiguities in the structure and the possibility that a given string could have more than one interpretation. However, between two individuals who share the same language code, the shared knowledge that a sender can reliably invoke is the conventional association between linguistic forms and meanings that she holds and that she assumes the receiver holds in a form similar to her own. The information contained

in structural elements narrows down the range of possible interpretations to the point at which the inferential process can bridge the remaining gap.

6 Synthesis

Having discussed the research that leads up to my main argument, I am now in a position to argue directly for a working conceptual definition of human linguistic communication as *a discrete combinatorial system used by individuals capable of inferring others' mental states*, and the related claim that *all of the important features of human linguistic communication thus far posited emerge from one of the two parts of this definition or from the interaction of the two*.

The first part of this claim is about a particular structural feature, a discrete combinatorial system, that I argue provides a practical basis for other characteristic features of human language, namely displacement, productivity, reference, and recursion, as discussed above in chapter 3. A discrete combinatorial system provides an efficient solution to the dilemma of encoding an infinite inventory of possible sentence-meanings with a finite and manageable set of starting units. Having a discrete combinatorial system to work with allows speakers to compose messages whose encoded sentence-meanings closely resemble the speakers' intended meanings, and therefore to create shared contexts with their hearers that allow the hearers to infer speakers' intended meanings with a high degree of efficiency and accuracy.

The second part of my claim is that humans' ability to make inferences interacts with the discrete combinatorial system to derive the other identified features of HLC. To continue a metaphor used in an earlier chapter, the discrete combinatorial system is equivalent to the lumber and hardware used in constructing a building: Each is a highly useful set of modular materials ready to be built into a house, or a message, but which requires the input of a skilled user--a carpenter, or an inference-capable communicator--to become a finished product.

In the various sections of this chapter, I demonstrate the way in which each characteristic feature of HLC--productivity, displacement, reference, and recursion--emerges from the interaction of the discrete combinatorial system and the inferential abilities of its users. An important disclaimer to present up front is that I am arguing these emergences from a practical, not necessarily a logical, perspective. That is, I am

not presenting the case that a discrete combinatorial system and inference are, together, necessary and sufficient for each of the other features. Rather, I show that these other features are available as epiphenomena of the discrete combinatorial system and inference.

6.1 Discreteness as a prerequisite for combination

First, a discussion of the features that are primarily structural: discreteness as a prerequisite for combination, and combination as a prerequisite for recursion.

An arbitrarily comprehensive communication system--one that can express any given message--requires an arbitrarily large inventory of signs in order to supply one sign for each communicated message. In a system in which one meaning is paired with one atomic sign, the problem arises that the inventory of physical signals available cannot support the increase in the number of signs necessary to express all of the desired messages, let alone an infinite number of messages.

In a fixed sign space, assigning a slice of the space to each possible message requires finer and finer subdivisions to the continuous medium. This requires users of the system to be able to produce and perceive finer and finer distinctions in the medium. On a practical level, the burdens placed on the individual perceiving or producing the signs become untenable. An example of a fixed sign space is the visible color spectrum. If a system requires two signs, it is easy to choose two colors that are easily perceived as different, such as red and green. Adding more signs requires finer divisions, such as adding yellow, violet, and so on. But as finding a color in a box of sixty-four crayons shows, there comes a point when it is difficult to reliably distinguish between related shades: imagine the difficulty in interpreting a message in which violet-blue conveyed an entirely different meaning from blue-violet, and so on.

Of course, not all sign spaces are fixed. If messages were conveyed by, for example, the continuous medium of distance between two points, each new message could be encoded in a sign that consisted of moving the two points even further apart. Such a system is theoretically unbounded. However, it is clear that this kind of system is limited by physical constraints on perception and production as well, in such theoretically

uninteresting but practically limiting ways as the time it takes to move the two points to the appropriate distance, the difficulty in perceiving a faraway point and in judging arbitrarily large distances, and so on.

Shifting to a discrete medium can help with some of the problems found in continuous media. Discrete signs imply an unbounded rather than a bounded sign space, because the creation of new signs requires neither finer divisions of a fixed space nor indefinite outward expansion of an expandable space. For example, the set of integers can be thought of as a discrete sign space, in which a new message could be encoded by the next integer in line. Arbitrarily large integers are of course not physically out of the reach of the users of the system, but they do require individuals to remember a huge and ever-increasing inventory of signs. (It must be remembered that the convention in the western numbering system with which we are familiar is to encode integers with a discrete combinatorial system, e.g. combinations of the discrete digits 0 through 9. The problem of remembering an arbitrarily large number of symbols is not that of decoding combinations of these ten digits, such as interpreting '11' as 'eleven', '12' as 'twelve', and so on, but akin to remembering another arbitrary squiggle as 'eleven', and yet another as 'twelve', and so on. It is clear that at some point the average human's memory would be severely taxed, and numbers in the thousands or millions would be functionally inaccessible.)

In fact, the example given above of the decimal numbering system shows an elegant solution to the problem of the ever-expanding sign repertoire: the trading of an arbitrarily large sign set for a limited set supplemented by rules that govern the arrangement of these signs in systematically meaningful sequences: which is to say, a combinatorial system. In order to represent any given number, we need not have memorized a unique symbol for that number, but rather we need only to know the ten digits of the decimal system plus the syntactic conventions of placing these symbols in set positions relative to one another to indicate a particular number (based on the additional convention of the base-ten system that breaks numbers down into how many thousands, how many hundreds, how many tens, and so on).

It is possible to apply combination to the units of continuous, not just discrete, systems. For example, two integers could be added to create a third integer, or yellow-green and green could be blended to make greenish-yellow-green. Let us say that the integer 5 is assigned a unit of meaning such as *eat*, and that 7 is linked to *tuna sandwich*. These two integers could be added (or multiplied, or averaged etc.) to create a message such as *eat tuna sandwich*, represented by the integer 12 (or 35, or 6...). We could do the same with yellow-green mapped to *eat* and green to *tuna sandwich* to arrive at greenish-yellow-green, *eat tuna sandwich*.

Before discussing further, it is important to note that this tuna sandwich example already has the problem that discrete units of meaning are being mapped to the continuous sign space: if the model were extended it could be demonstrated that somewhere in the spectrum, in the simplest case between yellow-green and green, there must be a boundary between *'eat'* and *'tuna sandwich'*, which means that the sign structure is in fact discrete. A truly continuous signal medium can only encode semantic meanings that themselves exist on a continuous scale. For example, the pitch on which I sing the syllable 'la' might indicate how happy I am: a low pitch indicates that I am not especially happy and a high pitch shows that I am very happy. Some human non-linguistic communication does in fact work this way, so that a bigger smile, tighter eyes, louder or higher-pitched cries, and so on indicate more extreme versions of the emotions to which they correspond. This system in which relative pitch indicates relative happiness is viable because both media are continuous, but a similar system in which the pitch code corresponds to discrete semantic meanings would not be. If a certain pitch or pitch range meant 'Deal me another card' and an adjacent pitch range meant 'The apple blossoms are abundant this year', and so on, then in a practical sense we could no longer say that the signal was continuous: we would be positing the ability to distinguish some sort of perceptual boundary between the pitch ranges that encode each of the different meanings, and thus the medium would be in effect discrete even if in a strictly physical sense it is continuous. This is in fact the case with the *p* and *b* example above, in which somewhere along the spectrum from a fully voiceless bilabial stop to a fully voiced one the perceptions of English speakers will flip from hearing p to hearing b, despite the lack

of any physically determined boundary in the signal. It is also true with the tones in tonal languages such as Hmong, in which the relative pitch at which a unit of sound is spoken can distinguish two different words. Pitch itself is a continuous medium, but Hmong--and tonal languages in general--have a discrete set of possible tones that can attach to a word. Speakers of these languages can reliably distinguish words with e.g. a high tone from those with a high-mid tone and so on, and in that sense the continuous medium of pitch is functionally used as a discrete medium to carry tonal information.

Let us nevertheless continue extending the example in which the integer 12 has the meaning *eat tuna sandwich* by virtue of the adding of 5, *eat*, and 7, *tuna sandwich*. One problem with this system is that certain slices of the sign space would have to be reserved for these combined meanings, and could not be assigned their own atomic meanings. In this case, we might be able to assign individual meanings to the numbers 1 through 11, but in order for 12 to be available as a combination of 5 and 7, it could not already have its own unitary meaning. Of course, it becomes clear that this problem extends over the entire sign space: 5 or 7 or any integer could be a combination of others.

In addition, 12 could be reached as the sum of 5 and 7, but also as 4 and 8, 3 and 9, and so on. If 4 and 8 each have their own unitary meanings, the system breaks down: 12 cannot be assigned to the combination of the meanings of 4 and 8, having already been assigned to the sum of 5 and 7.

The advantage of combining discrete rather than continuous units is that discrete units do not have the restriction of being bound along the dimension of continuity. In a series of integers, or color spectrum, or any other continuous sign space, combined signs must coexist along the same dimension as unitary signs. Combinations of discrete signs, on the other hand, because they are transparent in their derivation (a phrase such as ‘eat tuna sandwich’ clearly contains the elements ‘eat’ and ‘tuna sandwich’, for example), are not crowded into the same array as their units the way that 12 exists in the same string of integers as 5 and 7, or that greenish-yellow-green exists along the same line from red to violet that its factors, green and yellow-green, do. New chunks formed from discrete combinations can occupy another tier in the model that can be thought of as laterally displaced from the original array of atomic units: if we assign basic unanalyzable units to,

say, the first tier, we can put two-unit combinations in a second tier, three-unit combinations in a third tier, and so on. In this way, having discrete units as the input for the combinatorial process allows a much more efficient use of the sign space than having continuous units.

Two ideas sum up this section. The first is that combination relieves the production and perception difficulties that arise from a system in which each of many messages is encoded in a single, unanalyzable unit. The second is that combination is much more usable when the units to be combined are discrete from one another rather than continuous.

6.2 Combination as a prerequisite for phrase structure and recursion

The aim of this dissertation is not to describe the complexities of the language structure itself, which of course is the realm of the linguistic subfield of syntax. Nevertheless, with so much emphasis on recursion specifically, and phrase structure in general, as the element that is unique both to humans and to language (see HCF 2002), it is necessary to address these structural features at least briefly in order to situate this dissertation with respect to other work in related areas. As will become clear, I do not include either phrase structure or recursion among the properties without which human linguistic communication would be unrecognizeable. However, I discuss these properties here in order to show how my investigation meshes with other work in these areas.

First, some basic terminology. The property of being a discrete combinatorial system, as discussed throughout this dissertation, implies no more than that the elements are individual units that can be at minimum concatenated together. The simplest form of concatenation is a simple jumble with no respect to order, as mentioned in chapter 4. A dinner salad, for example, might consist of lettuce, tomatoes, and carrots, but rearranging the elements by tossing the salad does not change the salad in any meaningful way. Likewise, in the game of Bingo, the winning squares must all be crossed off but the order in which this happens is irrelevant.

A slightly more complex form of concatenation is one in which the linear order of constituent elements is significant. DNA is an example of a discrete combinatorial

system with simple linear concatenation as its only rule (Pinker 1994); if the elements of a strand of DNA are reordered, the DNA would now encode the plan for a different organism than it had with its original order.

Human language is structured in a yet more complex way, into phrase structures, on several levels (see Kinsella 2009). The combination of phonemes into syllables shows a phrase structure as the nucleus and coda of a syllable combine to make a unit called a rhyme, and this rhyme combines with an onset to form the syllable, displaying a phrase structure than can be represented as [onset [nucleus coda]]. Likewise, in a sentence such as 'The funny monkey scratched its belly and ate an apple', some elements are more closely related into phrases than others: even without going into detail about the structure of this sentence, it is clear to speakers of English that [the funny monkey] and [scratched its belly] hold together in a way that other linear sequences, such as [funny monkey scratched] and [belly and ate an] do not.[14] Likewise, it is clear that [scratched its belly and ate an apple] consists of the units [scratched its belly] and [ate an apple] but is also a unit that functions together with [the funny monkey] to form the whole sentence. Individual words are the discrete units that are combined into sentences, but between the word level and the sentence level are indefinitely many phrase levels. It is clear that a system displaying phrase structure must be based on a discrete combinatorial system, since the creation of phrase structures requires that individual elements be combined. Note that the implication is unidirectional: as mentioned above, a system can be discrete and combinatorial but limited to simple linear concatenation in its combination of elements.

One step father in complexity is the recursive system, in which phrase structures can be created with elements identical to the output element. Just as it is possible to have a discrete combinatorial system without phrase structure, it is possible to have a phrase structural system without recursion: the song of the Bengalese finch has been described as such a system (Okanoya 2002). In a recursive system, a phrase type can include

[14] Intuition alone need not be the only indicator that shows the cohesion of some units and not others. Simple tests such as whether another unit such as 'it' for 'the funny monkey' and 'did so' for 'scratched its belly' among others, can demonstrate this on a more rigorous level.

another phrase of its same type. Recursion sets up the possibility of infinite embedding, and thus accounts for the possibility of infinite output forms from a finite set of input forms. Hornstein (2009) accounts for recursion as a process of concatenation and labeling: A and B can be concatenated into A^B, and then this unit is relabeled as a type of A, as in [$_A$ A^B]. This new A-type unit can be concatenated with C, and so on (Hornstein 2009:58, example (8)).

The various levels of organization can be represented abstractly as in figure 1, from the simple existence of discrete units in (a) through unordered combination in (b), ordered combination in (c), hierarchical structure in (d) and a specialized type of structure generation, recursion, in (e).

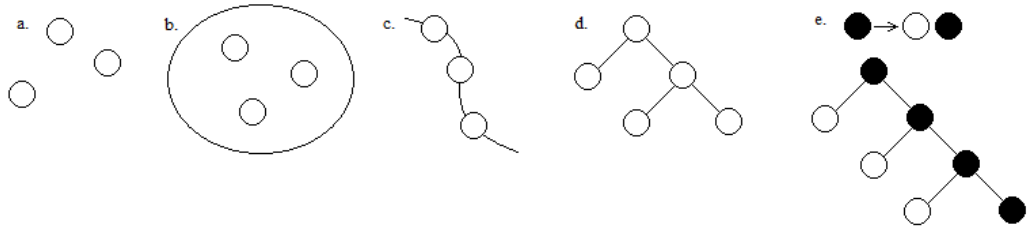

Figure 2. a. discreteness, b. non-ordered combination, c. ordered combination, d. hierarchical structure, e. recursion.

Taking these descriptions into account, both phrase structure and recursion can be seen as features that require a discrete combinatorial system as their base. In fact, to sum up the discussion so far, discreteness is a prerequisite for combination; discrete combination is a prerequisite for phrase structure, and discrete combination organized into a phrase structure is a prerequisite for recursion.

The presence of these features in the human language structure, and their almost-categorical[15] absence in other species, is not in dispute, nor is their contribution to the structural uniqueness of human language. However, this dissertation is not concerned with describing the language structure past the level that is required for the emergence of

[15] Recall of course the description from Okanoya (2002) of the phrase structure of Bengalese finch song.

other features of the system, such as the discreteness, productivity, and so on discussed in this chapter--features that bear on the communicative aspects of HLC.

The question relevant to the present discussion is this: What role does the presence of phrase structure and recursion play in engendering the features of HLC? It is my perhaps controversial claim in this dissertation that the communicative functions of HLC hinge on the presence of the first two elements in the structural hierarchy, discreteness and combination, the presence of which in human language are sufficient, when combined with the ability of a user to infer another's mental state, to allow the emergence of the functional features of HLC. I remain agnostic about the role of phrase structure and recursion in communication, saying neither that they are crucial for the level of function that I describe here, nor that they can be entirely dismissed as far as contributing functionally to HLC. I do, however, follow the lead of Kinsella (2009) in suggesting that while "recursion…permits essentially two communicative functions to be realised. It is used (i) to make reference to a finer level of precision, and (ii) to provide additional modifying information" (151). Sentences generated by a recursive process can be expressed with non-recursive sentences as well, as in the following:

(9a) The hermit lived in a cabin under a willow in the bog near the hedge along the highway from the mountain next to the river…

(10a) The cat [that] the dog [that] the snake [that] the boy [that] the girl liked saw bit chased died.

(9b) From the mountain next to the river there ran a highway. In the bog near the hedge along this highway, a hermit lived in a cabin under a willow tree.

(10b)The cat died. The dog had chased the cat. The snake had bitten the dog. The boy had seen the snake. The girl liked the boy.

Here, the non-recursively generated structures may be clearer for a receiver to understand than the recursive structures. This is not to say that recursion plays no role in communication: as Kinsella (2009) points out, "a non-recursive expression of a recursive concept is more structurally complex than its recursive counterpart. Repetition and/or numerous anaphora are required, which from both the speaker's and hearer's points of view are relatively costly. Thus, recursion furnishes the most efficient or optimal solution to the problem of expressing a recursive concept" (152). The basic idea is that communication may be more cumbersome but is certainly possible without recursion.

The more general property of having a phrase structure is arguably more basic to the human language structure than the specific property of recursion, in that a phrase structural system without recursion would be recognizable as human language in a way that a system without phrase structure would not. Items (9b) and (10b) above make use of phrase structure, and it would be difficult to paraphrase them in any way that did not display this feature. Examples of word strings that could easily be used to communicate, though rare, can be found, as in Pinker's (1994) "Skid crash hospital" and "Drum vapor worker cigarette flick boom."

Pinker's (1994) examples display ordered discrete combination, as they are essentially word-beads strung in a line. Removing one more layer of complexity from this hypothetical system, namely order, would leave a sort of 'word salad' from which some comprehensibility is lost: compare the original "Drum vapor worker cigarette flick boom" with the rearrangement "Flick worker boom vapor drum cigarette", in which the implied cause and effect from the original version is now missing.

Taking yet another step backwards into simplicity, we are left with nothing but nothing but discreteness as a property of the system. With only discreteness left, any proposition intended to be expressed would have to be contained within one unit of the discrete-but-not-combinatorial language. Even this restriction does not entirely rule out the possibility of communication, since the exclamations "Fire!" or "Run!" or "Boom!" can easily convey an entire message. However, relatively few messages can be reliably communicated by one-word utterances.

It is at the point at which ordered combination is eliminated that I judge communication based only on the shared knowledge of the language code to be impaired past usefulness. Thus, in reducing HLC to functionally atomic properties, I have included discreteness and combination (in its ordered form) as basic while excluding phrase structure and recursion. It is left to future research to demonstrate whether the division between essential and non-essential elements of HLC should be placed somewhere else along the hierarchy of structural complexity.

6.3 Discreteness and inference are sufficient for productivity

As mentioned above, productivity in a communicative system takes two forms: new sentences, and new mappings between words and concepts. For each of these, productivity relies on two features: the ability of the structure to support the creation of new forms, and the ability of users to understand both the semantic meanings and the speaker-meanings when new forms are used.

Let us first discuss the creation of new sentences. Productivity in creating sentences, on the level of the language structure, derives directly from discrete combination: each sign encodes a separate component of a proposition, along with its possible morphosyntactic roles. These components then combine to fill both the morphosyntactic requirements of a grammatical sentence and the semantic requirements of a complete semantic proposition. A well-formed sentence can be built out of just about any compositional elements, mix-and-match style, as long as the syntactic and semantic requirements of the sentence elements are mutually congruent: verbs may be required to agree with nouns, for example, or a transitive verb may require two arguments, among many possibilities. Given a vocabulary such as [Fido, Felix, eats, sniffs, ice cream, bacon], each element labeled with its possible syntactic category such as noun or verb, and the rules of English syntax that allow us to combine these into sentences that express an action performed by a subject upon an object, one can create and interpret sentences such as ‘Fido sniffs ice cream’ and ‘bacon eats Felix’ with little trouble. If the inventory of signs is augmented by adding [Clio, crackers], it is now

possible to form ‘Clio eats Fido’, ‘Felix eats crackers’, and many other new sentences. This is all possible even if the sentence ‘Clio eats Fido’ has never before been formed.

The limited symbol inventory and relatively simple set of syntactic rules for combining them that compose a discrete combinatorial system confer another powerful advantage to biological organisms, who must produce and interpret messages online in real time and often quite rapidly in response to environmental stimuli, by allowing the sign for any given message to be calculated systematically in the communicative situation rather than having to be retrieved from memory. The inventory of tens of thousands of morphemes in any given human language has to be stored in memory, but the inventory of possible messages is unlimited.

Contrast this with the difficulty that would be found if each new proposition that a speaker wanted to form required an entirely new, unanalyzable sign for its expression. In theory a user tinkering with such a system could create a new sign (a new arbitrary squiggle, the next larger integer, or a finer division in the color spectrum, to extend earlier examples) and map a meaning onto it.

However, since we assume that these new signs are created in order to help a sender convey a proposition to a receiver, some communication difficulties are likely to arise. An arbitrary, unanalyzable sign that no one else has ever seen before cannot, by definition, have a conventional semantic meaning attached to it. Because a message made from the parts of a discrete combinatorial system derives its conventional meaning from both the meanings of its parts and the meaning imparted by the syntactic structure of the message, a new message formed in this way can have a semantic meaning that a hearer can then use to infer a speaker’s intended meaning. The discrete combinatorial system allows new semantic meanings to piggyback on the meanings of existing units, thus making conventional signs available even for messages never before expressed.

So far the discussion has related only to the coding and decoding of messages on the conventional semantic level. The inferential process mainly occurs outside of the level of sentence productivity, governing instead the formation of the phrases and sentences uttered as evidence of a speaker’s meaning. Productivity in HLC is not limited to the creation of new sentences, but also includes the assignment of new sound-to-

meaning mappings for existing morphemes, such as the English word 'text' which has recently gained the additional semantic meaning 'to send a text message'. The means by which new words are created is different from that by which new sentences are made in that the creation of new words is an expansion of the structure where the creation of new sentences is simply a use of the existing structure: in the extended homebuilding metaphor used elsewhere in this dissertation, phrase generation from morphemes is akin to using hardware and lumber in various configurations to build all shapes of buildings, whereas the creation of new words is equivalent to the creation of new materials such as new sizes and shapes of boards and nails and so on. It seems intuitively clear that the process of creating new words is different in kind than that of creating new sentences.

I take care in this section to differentiate terms for forms from terms for form-to-meaning mappings. Since the term 'word' can be used for either of these in various contexts, I use the term 'word form' or 'phonological word' when referring to the form itself, and 'semantic word' or 'mapping' to talk about a linked word form and its conventional meaning. For example, the phonological word 'bear', represented [ber], can be described as a consonant-vowel-consonant sequence or as consisting of one syllable, etc. On a semantic level, the form 'bear' takes part in several different words, each of which uses the same phonological form but which map to *ursine mammal*, *carry*, *endure*, and so on, and thus these are three different semantic words.

The creation of new words is the establishment of new mappings, and as such relies much more on the inferential process than the creation of new phrases does. The mapping between a word form and a meaning is conventional, and must be learned. New words can be learned explicitly, as when a biology textbook presents and defines a term such as 'mitosis'. However, many words that a hearer learns are acquired simply by the speaker's hearing them in a particular context and assimilating them into his or her vocabulary. For example, the term 'cool' exists in the casual vocabularies of many American English speakers with the approximate meaning of *interesting*, *popular*, *desirable*. Yet most people did not learn this word by being told that it means *interesting*, *popular*, *desirable* but by hearing it used in situations when they could infer that the speaker held that opinion about the object to which 'cool' was applied. The productivity

that allowed the new word ‘cool’ therefore must be attributed directly to the human inferential ability.

Likewise, even words that do not spread to the community of speakers at large but are used only in a one-time situation can be understood by hearers who have never encountered them before. I might be sitting with a shedding cat on my lap, and when the cat hops down I might look at my fur-covered pants and say, ‘The cat haired me!’ (Caroledith Olsen, p.c.) The meaning *to leave hair on something* is not included in a general semantic definition of the word ‘hair’, and yet my intended meaning would be perfectly comprehensible to a speaker of English. This example illustrates the productivity of the code in that I created a past-tense verb using the suffix ‘-ed’, but also the role of inference in the sense that this utterance is incomprehensible if a conventional semantic meaning were the only interpretation available to a hearer of the word ‘hair’.

This kind of accommodation is common enough that it may be easy to take for granted, but some analysis shows that a sophisticated process of ostensive-inferential communication is at work. If understanding a speaker’s meaning were purely a process of decoding the semantic values of the elements of her utterance, a hearer could never understand any sentence that contained a word that he did not already know. There would be no way to repair the conversation without asking for the speaker to define the word explicitly, and even so it is impossible to be sure that the concept that a hearer develops from hearing the word’s definition is really the same as the concept that the speaker believes that the word represents.

In fact, except for explicitly defined words such as in the ‘mitosis’ example above, a hearer can never be completely sure that his mental definition of any word is the same as that of any given speaker. The knowledge of a word’s semantic meaning is accumulated through knowledge of the types of situations in which using that word gets a speaker’s intended meaning across in the most efficient way possible.

In this sense, even uses of words that a hearer knows in the ordinary sense will invoke productivity in that the hearer must accommodate the application of a word to a concept that he may not have linked to that word in his previous experience.

For example, a hearer might encounter a certain piece of fabric, guess that it is meant as a cover for a dining table, and therefore think of it as a 'tablecloth', but if another speaker sees that same piece of fabric and guesses that it is meant to be spread on a bed, she might call it a 'quilt'. If these two people are conversing about this item in a location that does not suggest one or the other usage for the fabric, for example at a garage sale, when the speaker says, "What a lovely quilt!" the hearer will probably not have much trouble understanding that she is talking about the same piece of fabric that he thought of as a tablecloth. (It may take the hearer an extra moment to grasp the speaker's meaning, depending on how strong his belief originally was that the item was a tablecloth and not a quilt, but what is crucial for my argument is that this mental readjustment can be done at all.) The mechanism by which the hearer successfully interprets the speaker's meaning is not based on any features of the code but rather by other cognitive features of the users of the system.

Likewise, existing words can take on new or altered semantic values over time, which is only possible if the speaker's intention when she uses a particular word is the crucial element for conveying a meaning to the hearer rather than the strict semantic value of the word. For example, the most basic semantic meaning of 'princess' is *daughter of a king or queen*. But if this were the only concept accessible to a hearer interpreting what I mean when I use the word 'princess', then the sentence 'She's a real princess,' uttered about a spoiled (non-royal) child, would make no sense. This usage of the word might gradually become common enough that *spoiled child* becomes another semantic meaning mapped to the word form 'princess'. However, this taking on of variations in semantic meaning would not be possible at all if hearers were not able to accommodate speakers' intended meanings during the process.

6.4 Discrete combination and inference are sufficient for displacement

A communication system that can support the transmission of messages that refer to entities removed in time or space from the communicative situation has the feature of displacement. As discussed earlier, displacement is a crucial feature in granting HLC its flexibility and usefulness.

Strictly speaking, displacement is a relatively trivial feature for a system to display. The semantic meanings of the squares of a wall calendar that represent days of the month are displaced, and if a sender were to point to a square on the calendar as a way of invoking some future date in the mind of a receiver, the system being used shows displacement. To give an even more trivial example, if 'X' encodes *the location 50 yards north of where we are* and 'Y' encodes *yesterday*, then any message sent in this system is necessarily displaced even though the system is otherwise extremely unsophisticated.

It is important to note that discreteness and combination are not logical requirements for displacement. The signals of a continuous system can correspond to entities that are not present in space or time. All that is necessary is for the signs to map to values removed from the communicative situation. In the bee dance code, the angle of the dance corresponds to an angle from the sun and can be set to any point along the 360 degrees of the circle. Likewise, the quality of the food source is indicated by the enthusiasm of the waggle portion of the dance, with a greater enthusiasm indicating a richer source. Thus, the bee dance has the feature of displacement: The food source is removed in space from the dance, and in time to the extent that the bees are not experiencing it now but will do so in the time that it takes to fly to the source.

So, displacement itself is simple. What is key to HLC, however, is not displacement per se, but productive displacement: the ability of the system to support new displaced messages. A system that shows productive displacement--where new signals can be used to talk about things removed in space and time--requires discreteness.

Because the different elements of a message are encoded in discrete parts, individual ideas can be mentioned in the absence of any particular predication. That is, humans can mention things without being compelled to say any particular thing about them: the entity being discussed and what a speaker is saying about it are coded in separate parts of the message. If a speaker were to say to a hearer, "Eagle," the hearer would probably have various reactions depending on the context. On a birdwatching hike, he might look around on the assumption that the speaker meant to tell him that she sees an eagle and is calling it to his attention. If the hearer had asked the speaker to name some national symbols or to list as many raptors as she could, he would assume that she

meant to indicate to him that she knows that an eagle is a national symbol or is a raptor. However, if she suddenly blurted out "Eagle," in the absence of any (linguistic or non-linguistic) context in which that utterance was expected, he is most likely to shrug his shoulders and ask, "So what? What do you mean by that?" That is, while inference can bridge the gap between the utterance of the word 'eagle' and an intended message given an appropriate shared context, the word itself does not have any particular message other than referring to a category of entity, and thus speakers who use the word are not thereby committing themselves to saying anything in particular, or anything at all, about an eagle.

This is not the case with other types of communication systems. Contrast the example given above of HLC with the communication of vervet monkeys. In a system like that of vervet monkeys, one unitary utterance expresses an entire message without the possibility of separating the topic and what is predicated about it. This is quite different from HLC, in which the topic and what is predicated about it are discrete and each is separately retrievable from the signal. A message such as the eagle call causes monkeys who hear the call to look around for eagles or to become less vulnerable to attacking eagles by moving out of the treetops. As such, the unitary, indivisible eagle call transmits a message that is best glossed in English as 'There may be an eagle nearby' or 'eagles are a particular threat right now'. The message cannot be divided, as the English sentence can, into a part that invokes the idea of an eagle and another part that predicates of the eagle that it may be nearby. This is a feature of the discrete combinatorial system: In order for a structure made of discrete units to be useful, each unit must be independent of the others. Any predicate must be equally available when 'eagle' is the subject; if this is not the case, then 'eagle' and its predicate are not truly separate and independent units. 'Eagle' thus contain no information about whether any eagle is present here and now, and a speaker is free to use it without any necessary implication about what she may mean in using it.

6.5 Displacement and inference are sufficient for reference

Reference, in linguistic terms, is the act of picking out an entity or concept in the world through use of a linguistic sign. Reference is a speech act and is performed by

speakers: while it is common to say that a *word* refers to some thing or idea in the same way that we might say that a hammer pounds in a nail, the view from the field of linguistic pragmatics is that the word is a tool that a *speaker* uses to refer to a thing or idea the way a carpenter uses a hammer to pound a nail. To discuss whether a communication system has the feature of reference, then, is to discuss whether speakers can use that system to point to specific things or ideas.

It is debated whether a vervet monkey that makes an eagle or leopard or snake call is actually referring to that predator. But vervet calls, whether or not they have a referential element, are clearly not exclusively referential, and whatever reference is involved in a vervet call cannot be separated from the call's other function as a warning or call to action. A vervet's eagle call does not simply mean 'there's an eagle' or invoke the idea of an eagle in order to comment on it; the call is always used when the caller perceives an eagle nearby.

Humans can also warn others to be careful of an eagle, but the difference between HLC and the communication of vervets is that humans can invoke the idea of an eagle without being constrained in what we say about it. We can point out an eagle in the context of birdwatching without implying anything further than just observing its presence, or say in the summer in Minneapolis that we saw an eagle last winter in Alaska, or even explain that eagles are fish-eating raptors, invoking the concept of an eagle without referring to any particular individual eagle. This much is simply displacement, as discussed above.

But how does a hearer understand what particular eagle, if any, a human speaker (of English, in this example) means when she says the word 'eagle'? In order to successfully refer to something, whether an eagle present at the time of speaking, an eagle seen in the past or in a different place, eagles in general, or any of the many other possible referents for 'eagle', a speaker must know that her hearer will be able to pick out the eagle that she is talking about.

Certain clues carried by the language structure make this process easier. Information carried by determiners in phrases such as 'an eagle', 'the eagle', 'this eagle', and others can signal to a hearer exactly how close to his immediate awareness the

speaker believes the eagle in question to be (see Gundel, Hedberg, and Zacharski 1993). However, the key element in a successful act of linguistic reference is not the encoding of information within the structure of the utterance but the inferential process whereby a speaker can infer what the hearer is likely to understand from her utterance and the hearer can guess what the speaker most likely means by her utterance.

Among vervets, to provide a contrasting example, hearers almost certainly lack the ability to infer what referent a speaker means when there is more than one possibility, and speakers almost certainly do not go through the (unconscious) process of calculating what referent a hearer will identify when he hears an eagle call. Because of the lack of inferential ability, the references of vervet calls (if 'references' is the right term) are limited to eagles present, or implied to be present, in the immediate environment.

7 Conclusion

As an object of study, human linguistic communication can seem to defy analysis, being so fundamental and automatic to our experience and yet so complex as to evade attempts to pin down its nature. It is not a neatly-bounded phenomenon that can be isolated as a cognitive system among other, completely separate systems, nor does it consist only of the formulaic deployment of a preexisting code.

In a broad sense, the work of this dissertation applies to very broad questions about human language and communication, such as: What is the nature of language? How does human communication happen? What makes human language different from other formal and cognitive systems, and other communication systems? What are the basic elements of human linguistic communication? Why do humans communicate with language in ways in which other animals do not?

As a way of addressing these broad questions, I have formulated narrower questions, namely these: Is human linguistic communication different only in degree from other animal communication, or is it different in kind? If it is different in kind, can this difference best be attributed to one or a small number of core features? If so, what are these features? What role does the code itself play in characterizing HLC and what role is attributable to its communicative function?

I have shown that HLC has properties that the communication systems of other animals do not seem to display, namely productivity, displacement, reference, and the structural features of a discrete combinatorial system and higher-level organization into potentially recursive phrases. I posit that the combination of the structural features of discreteness and combination, which together make a discrete combinatorial system, combine with the human ability to make inferences about others' intentions in a way that allows the system to display the features of productivity, displacement, and reference. Structural units of human language can indicate a particular conventional meaning very precisely, based on the meanings of their constituent parts and the rules of syntax. This precision allows speakers to construct a message specific enough that the shared

knowledge of its conventional meaning is enough to allow hearers to infer the speaker's intention in communicating this message.

The key contribution of this dissertation is the demonstration that properties often considered basic to HLC arise neither from the language structure alone nor from communication alone, but specifically from the use of human language for communication. The language code itself is merely a tool that would remain inert and lifeless if it were not used by organisms who can use it to convey their communicative intentions and infer the communicative intentions of others. The beauty of HLC is that it pairs the system of the language code, a tool that can be wielded very precisely by skilled users, with the fundamental inferential system of human communication to yield a new, larger and marvelously useful system for transmitting messages.

As is always the case with a study that discusses a wide range of subtopics in relatively broad terms, there are certain avenues of thought that have necessarily been left for future research. A very important one among these is a thorough investigation of the inferential abilities of nonhuman animals. While there is some research that suggests that some other animals, primates such as vervet monkeys and chimpanzees chief among them, may have communicative intentions and the ability to infer others' communicative intentions, the literature is far from conclusive. Some of the claims that these animals are using inference in their communication may well be due to a tendency of primate researchers to read more into their subjects' behavior than the evidence strictly demonstrates, especially since the desired result in much nonhuman animal cognition research seems to be to find more complex abilities than have been attributed to these animals in the past. However, both the animal cognition literature and studies of human linguistic communication such as the present work would benefit from a thorough investigation of the abilities--and the use of any such abilities in communication--that other animals may have.

Another direction for future studies is the role that the feature of recursion plays in HLC. It is clear that recursion plays a large role in the theoretical account of the construction of sentences within the human language code. What is less clear is the

contribution it makes in day-to-day use of the code for communication. It is left to future scholars to evaluate the importance of recursion in the functioning of HLC.

The work in this dissertation is intended to provide a new perspective on the nature of HLC, one that updates the traditionally-cited features of language to be relevant within the prevailing cognitive paradigm, and that corrects some of the imprecisions in past work by carefully delineating a boundary between coded and communicative features even as it demonstrates the interrelatedness of the two.

8 References

Burling, R. 1993. "Primate calls, human language, and nonverbal communication." *Current Anthropology* 34(1):25-53.

Bybee, Joan. 2009. "Language universals and usage-based theory." Pp. 17-39 in *Language Universals*, edited by M.H. Christiansen, C. Collins and S. Edelman. Oxford: Oxford University Press.

Cheney, D. and R. Seyfarth. 1990. *How monkeys see the world.* Chicago: University of Chicago Press.

Chomsky, N. 1957. *Syntactic Structures.* The Hague: Mouton.

------. 1986. *Knowledge of language: its nature, origin, and use.* New York: Praeger.

Evans, Nicholas and Stephen C. Levinson. 2009. "The myth of language universals: Language diversity and its importance for cognitive science." *Behavioral and Brain Sciences* 32(5):429.

Fitch, W. T., M. D. Hauser and N. Chomsky. 2005. "The evolution of the language faculty: clarifications and implications." *Cognition* 97(2): 179-210.

Frisch, K. 1967. *The dance language and orientation of bees*. Cambridge, MA: Harvard University Press.

Frishberg, N. 1975. "Arbitrariness and iconicity: historical change in American Sign Language." *Language*: 696.

Gómez, Juan C. 1991. "Visual behaviour as a window for reading the mind of others in primates." Pp. 195-207 in *Natural theories of mind: Evolution, development and simulation of everyday mindreading.*, edited by A. Whiten. Cambridge, MA, US: Basil Blackwell.

Gould, J. L., and C. G. Gould. 1988. *The honey bee.* Scientific American Library.

Greenberg, Joseph H. 1966. *Universals of language.* 2nd ed. Oxford: MIT Press.

------. 1963. "Some universals of grammar with particular references to the order of meaningful elements." Pp. 73-113 in *Universals of human language*, edited by J.H. Greenberg. Cambridge, MA: MIT Press.

Grice, H. P., ed. 1975. *Logic and conversation.* New York: Academic Press.

Gundel, J., N. Hedberg and R. Zacharski. 1993. "Cognitive status and the form of referring expressions in discourse." *Language* 69:274-307.

Haiman, J. 1985. *Iconicity in syntax.* Amsterdam: John Benjamins.

Hauser, M. D., N. Chomsky and W. T. Fitch. 2002. "The faculty of language: What is it, who has it, and how does it evolve?" *Science* 298:1569-1579.

Herman, L. M. 1986. "Cognition and language competencies of bottlenose dolphins." Pp. 221-252 in *Dolphin cognition and behavior: A comparative approach*, edited by R. J Schusterman, J. A. Thomas, & F. G. Wood. Erlbaum.

Hillix, W. A., and D. M. Rumbaugh. 2004. *Animal bodies, human minds: Ape, dolphin, and parrot language skills.* New York: Klüwer/Academic Press.

Hockett, Charles. 1959. "Animal languages and human language." *Human Biology* 31(1):32-39.

------. 1960. "Logical considerations in the study of animal communication." Pp. 392-430 in *Animal sounds and communication*, edited by W.E. Lanyon and W.N. Tavolga. Washington, D.C.: American Institute of Biological Sciences.

Hockett, Charles and Stuart A. Altmann. 1968. "A note on design features." Pp. 61-72 in *Animal communication: techniques of study and results of research*, edited by T.A. Sebeok. Bloomington: Indiana University Press.

Hornstein, N. 2009. A Theory of Syntax: Minimal Operations and Universal Grammar. Cambridge: Cambridge University Press.

Jackendoff, R. and S. Pinker. 2005. "The nature of the language faculty and its implications for evolution of language (Reply to Fitch, Hauser, & Chomsky)." *Cognition* 97:211-225.

Kinsella, A. R. 2009. *Language Evolution and Syntactic Theory*. Cambridge: Cambridge University Press.

Michelsen, A. 1989. “Honeybees can be recruited by a mechanical model of a dancing bee.” *Naturwissenschaften* 76(6):277.

Okanoya, K. 2002. “Sexual display as a syntactic vehicle: The evolution of syntax in birdsong and human language through sexual selection.” In Alison Wray, ed., The Transition to Language. Oxford: Oxford University Press.

Pinker, S. and P. Bloom. 1990. "Natural language and natural selection." *Behavioral and Brain Sciences* 13(4):707-784.

Pinker, S. and R. Jackendoff. 2004. "The faculty of language: What's special about it?" *Cognition* 95:201-236.

Premack, D. 1983. "The codes of man and beasts." *Behavioral and Brain Sciences* 6(125):67.

Savage Rumbaugh, E. S. 1986. *Ape language: From conditioned response to symbol.* New York: Columbia University Press.

Shannon, C. E. 1948. "A mathematical theory of communication." *Bell System Technical Journal* 27 (July & October): 379.

Sperber, Dan and Deirdre Wilson. 1995. *Relevance: Communication and cognition.* Oxford: Blackwell.

Stokoe, W. C. 1972. *Semiotics and Human Sign Languages.* The Hague: Mouton.

Struhsaker, T. T. 1967. "Social structure among vervet monkeys (Cercopithecus aethiops)." *Behaviour*: 83.

Taylor, Charles. 1985. *Human agency and language.* New York: Cambridge University Press.