The Puzzle of Ambiguity[†]

Thomas Wasow, Amy Perfors, and David Beaver
Stanford University

0. Introduction

Montague's celebrated claim that no "important theoretical difference exists between formal and natural languages" (Montague 1974; 188) implies that ambiguity is not theoretically important, for ambiguity abounds in natural languages, whereas formal languages are unambiguous by design. More generally, the pervasiveness of ambiguity in natural languages seems to be widely regarded as unremarkable. Our objective in this paper is to argue, to the contrary, that the highly ambiguous character of natural languages is surprising, and that the very existence of ambiguity calls for an explanation.

Section 1 clarifies what we mean by ambiguity, discussing the distinction between vagueness and ambiguity. We go on to identify several distinct types of ambiguity. Section 2 presents evidence that English is massively ambiguous[1]. Section 3 elaborates our central argument: if (as is widely claimed) ambiguity impedes efficient communication, then one would expect languages to evolve so as to reduce ambiguity; but this does not appear to have happened. Section 4 responds to some possible objections to the argument in Section 3. Section 5 explores some possible strategies for explaining ambiguity, concluding with pointers to our ongoing research on the subject.

1. Characterizing Ambiguity

Ambiguity is a semantic property. Semanticists argue over exactly what meaning is, but it surely involves associating expressions in a language with something else (things or events in the world, mental representations, sets of possible worlds, or what have you). For our purposes, it is not necessary to decide what that something else is. We will assume that meanings are regions in a space, remaining agnostic as to its dimensionality, the metaphysical status of the points in it, and exactly how linguistic expressions get associated with regions. We will use the term 'denotation' for the association between expressions and regions.

An expression is ambiguous if it has two or more distinct denotations – that is, if it is associated with more than one region of the meaning space. The most obvious instances of ambiguity involve expressions with (at least) two

[†] Perfors's work on this paper was supported by the Center for Integrative Research in Science and the Humanities at Stanford. We are grateful to Peter Sells for helpful comments on an earlier draft, but absolve him of blame for any errors in this version.

[1] We assume that this is not an idiosyncracy of English. Although we have not investigated the matter systematically, we are confident that ambiguity is widespread in other natural languages.

denotations that are disjoint from one another. A standard example is *bank*, which can denote the rim of a river or a financial institution; another example is *long*, which as an adjective denotes size along one dimension but as a verb is roughly synonymous with *yearn*. But there are plausible cases of expressions that have multiple non−disjoint denotations. For example, the verb *eat* can denote (*inter alia*) either the ingestion of food (*We ate the cake*) or gradual destructive consumption (*Salt ate the paint on the right fender*); both denotations are appropriate to (1), although *eat* does not seem ambiguous in this sentence.

1. Termites ate part of the kitchen floorboards.

The gradual destructive consumption meaning of *eat* probably originated as a metaphorical extension of the ingestion meaning; in contemporary usage, however, they have distinct, but overlapping denotations.

We even leave open the possibility that an expression may have two meanings one of which is properly included in the other. A plausible example is *cow*, which can denote either a mature female bovine or cattle of any age or sex. Despite the fact that one meaning is a special case of the other, they are listed as separate definitions in the *American Heritage Dictionary*; we share the intuition that these are two distinct interpretations of the word.

Ambiguity is not the same thing as vagueness. Expressions are vague if the regions they denote do not have perfectly well−defined boundaries. Most expressions in natural languages are vague – that is, the denotations of most expressions are fuzzy around the edges. As lexicographers know (but linguists and philosophers of language seem reluctant to accept), it is rarely possible to provide necessary and sufficient conditions for membership in the denotation of a natural language expression. Exceptions to this tend to be technical terms of science or mathematics, which have been assigned precise definitions for some restricted use. Whether sucking a lozenge counts as eating has no definitive answer: it depends on the context and the reason the question was raised. Whether an investment club or a check−cashing store should be included in the denotation of *bank* is likewise context−dependent. These questions are orthogonal to the ambiguity of the words. The denotations of linguistic expressions typically have fuzzy boundaries, making them vague; but ambiguity requires more than one denotation.

Our characterization of the ambiguity/vagueness distinction is only a sketch; much additional detail would be necessary if we wanted to provide a diagnostic for distinguishing them (see Zwicky and Sadock, 1975, for a different characterization of the distinction and a discussion of diagnostics). There are cases of clear ambiguity, most obviously homophones like *bank* or *long* (whose two meanings have no intuitive connection). There are also many unambiguous, but vague expressions, e.g. *tiny*, which always means very small, although exactly what counts as tiny is almost never precisely defined. But not all cases are so clear. For example, the word *house* may conjure up an image of a structure designed for human habitation, but we also speak of bird houses,

beaver houses, etc. The word is surely vague, but is it also ambiguous? We leave this question open. In what follows, we focus on cases that are clearly ambiguous, and not just vague.

There are several types of ambiguity. The examples given so far have all been lexical – that is, words with more than one denotation. But both smaller and larger expressions may exhibit ambiguities. Inflectional morphemes are ambiguous so often that there is a technical term for the phenomenon: syncretism. For example, suffixing *–s* (with phonologically conditioned allomorphs) is used in English to mark third−person singular agreement in verbs, plural number on common nouns, and possessive forms of noun phrases. This can lead to ambiguous expressions like *the weapon(')s inspector*, whose two denotations are distinguished orthographically, but not phonologically. See Stump (2001), particularly Chapter 7, for many more examples.

There are also ambiguous multi−word expressions whose ambiguity cannot be attributed to the ambiguity of any morpheme or word in them. Many strings of words admit of more than one syntactic parse, and distinct parses are characteristically associated with distinct meanings. Classic examples are phrases like those in (2).

2 a. small dogs and cats
b. We saw the man with the telescope.

In (2a), *small* may modify just *dogs* or the conjunction *dogs and cats*; in (2b), *with the telescope* may modify either *saw* or *man*.

Other ambiguities do not appear to be attributable either to ambiguous morphemes or to multiple (surface) parses. These include scope ambiguities like (3), which can mean either: (i) there was no student who solved exactly two problems, or (ii) there were exactly two problems that no student solved.[2]

3. No student solved exactly two problems.

(4) is likewise ambiguous, since the chicken can be interpreted either as the eater or what is to be eaten. This ambiguity, like that in (3), cannot naturally be attributed to the ambiguity of any one word, nor to distinct (surface) parses.

4. The chicken is ready to eat.

---

[2] It has occasionally been suggested (see e.g., Kempson and Cormack, 1982) that examples like *Some student solved every problem* are not truly ambiguous, since once interpretation (in which a specific student solved every problem) entails the other (in which every problem was solved by a student, but not necessarily by the same student). Example (3) is constructed so that neither interpretation entails the other. This can be seen by considering a situation with three students (call them A, B, and C) and four problems (call them 1, 2, 3, and 4). If all three students solve 1 and 2 but fail to solve 3 and 4, then interpretation (i) is false and (ii) is true. If A solves only 1, B solves only 2, and C solves only 3, then (i) is true and (ii) is false.

There are ambiguous expressions that involve a combination of these types of ambiguity. For example, Chomsky's classic example (5) involves a subject/object ambiguity analogous to (4), as well as syncretism of the *–ing* suffix, resulting in syntactic ambiguity regarding the grammatical category of *flying*.

5. Flying planes can be dangerous.

On one interpretation, *flying* is a verb and *planes* is interpreted as its object; on the other, *flying* is an adjective and *planes* is interpreted as its subject. Similarly, *She gave the boys' animals toys* and *She gave the boys animals' toys* sound the same but mean different things, because of a combination of morphological and syntactic ambiguity.

2. Ambiguity is Pervasive in English

Lexical ambiguity is extremely common. A crude measure of just how common is provided by the number of definitions provided for words in standard dictionaries. To be sure, many of the fine meaning distinctions found in dictionaries reflect lexicographers' attempts to deal with the vagueness of most natural language words. But genuine polysemy is the rule, rather than the exception, particularly among frequently used words. (6) lists the ten most frequent nouns, verbs, and adjectives from the British National Corpus[3].

6.

NOUNS	VERBS	ADJECTIVES
time	say	other
year	get	good
people	make	new
way	go	old
man	see	great
day	know	high
thing	take	small
child	think	different
government	come	large
work	give	local

Many of these words (including at least *time, people, man, work, say, go, take, give, good, high, small, local*) have uses in a grammatical category other than the one they are listed in. Many also have clear ambiguities within the listed category. For example, *time* can denote an occurrence of an event (as in *one more time*) or a musical rhythm, among many other things; *way* can denote an abstract means or a road; *make a bed* denotes two entirely distinct acts, depending on whether the maker is a carpenter or a housekeeper; the denotations of *know* in knowing a

[3] We have excluded auxiliary verbs, pronouns, and titles (*Mr*). These are taken from a lemmatized list – that is, different inflected forms of the same word are counted together.

person, a proposition, and a language are sufficiently different that they are translated into three different verbs in German (*kennen, wissen,* and *können*); *old* can denote either aged or former (so that *her old young boyfriend* is not necessarily an oxymoron); and *local* contrasts with either *global* or *express*.

Even function words are often ambiguous. The uses of *to* as a preposition and an infinitive marker have no apparent semantic connection. *For* can be a preposition marking benefactives or a sentential connective indicating causation. A number of connectives (e.g. *then, since, while*) are ambiguous between temporal and conditional meanings. *That* can be a determiner, a demonstrative pronoun, or a complementizer.

It would be easy to continue citing many more examples, but it should not be necessary. A little reflection should make it evident that lexical ambiguity is extremely widespread in English.

Syntactic ambiguity is likewise very common, though perhaps not as obvious. In the 1970s and 1980s, when computers became powerful enough to support attempts to implement grammars with reasonably broad coverage, computational linguists were surprised to discover that their systems were assigning more parses to many sentences than they had expected. For example, Martin et al. (1987) reported that their system assigned 455 parses to the sentence *List sales of the products produced in 1973 with the products produced in 1972.*

Close examination of the extra parses typically showed that they were not the result of errors in the systems. Rather, well–motivated grammatical mechanisms would apply in unexpected ways, yielding unanticipated structures. Sometimes the extra parses were pragmatically bizarre, but hard to rule out without invoking extra–linguistic knowledge. In other cases, complex situations would need to be constructed to tease apart the truth conditions of different parses. But some constructions that can lead to massive ambiguity are syntactically and semantically quite straightforward.

One very productive source of syntactic ambiguity in English is the attachment of modifiers, especially prepositional phrases. Phrasal modifiers typically come at the ends of their constituents, and, since English is predominantly right–branching, this means that sentence–final modifiers often have numerous options for attachment. (2b) above is a simple example of this. Church and Patil (1982) discuss more complex examples and their implications for natural language processing. Starting with the simple sentence *Put the block in the box on the table*, which has two parses, they show how adding more prepositional phrases at the end (*in the kitchen*, etc.) rapidly increases the number of parses: with one more PP, there are five parses, and with two more, fourteen[4]. They sum up the situation as follows (p. 139):

[4] Church and Patil show that the same rapid increase in number of parses occurs in coordinate structures as the number of conjuncts increases, although it is less clear in that case that there is a corresponding semantic ambiguity.

> Sentences are far more ambiguous than one might have thought. Our experience…indicates that there may be hundreds, perhaps thousands, of syntactic parse trees for certain very natural sentences of English.

The situation is exacerbated by the interaction of independent ambiguities. Suppose a sentence has three ambiguous lexical items and two (disjoint) places with attachment ambiguities; even if each ambiguity allows only two possibilities, the sentence will have, in principle, $2^5 = 32$ interpretations. A simple example satisfying this description is given in (7); others would be easy to construct.

7. Old friends and acquaintances remembered Pat's last time in California.

Here *old* can mean aged or long−term (or former) and can modify either *friends and acquaintances* or just *friends*; *last* can mean final or previous; *time* can mean occurrence or duration (e.g. if Pat was a racer), and *in California* can modify *remembered* or *time*.

In short, the prevalence of independent morphological, lexical, syntactic , and scopal ambiguities can lead to a combinatorial explosion, making many English sentences massively ambiguous.

3. Ambiguity is Supposed to be Hard

Readers may wonder why the pervasiveness of ambiguity in English is of interest.

The central function of languages is presumably the representation of information, normally for communicative purposes. This is true of formal languages (that is, mathematical formalisms like predicate calculus and computer languages like C) as well as natural languages. In the case of computer languages, the communication is with a computer, rather than between people, but the language is still serving as a medium for conveying information from the programmer to the machine. Communication involves (at least) two participants, one that produces a signal and another that interprets it. Communication is successful to the extent that the interpreter assigns to the signal the denotation intended by its producer. If the language assigns a unique denotation to each signal, then the interpreter's task is relatively simple. If signals may have multiple denotations, then interpreters must try to disambiguate, so as to infer the producer's intention. This entails extra effort on the part of the interpreter. Moreover, if the methods of disambiguation are fallible, then ambiguity increases the chances of miscommunication.

For this reason, all formal languages have been designed to be unambiguous.

Where ambiguity might arise (for example, in the order of arithmetic or logical operations), either parentheses (or some similar notation) are employed, or an exceptionless convention is built into the syntax of the language (as in the case of Polish notation). Natural languages have not been designed. Rather, they have evolved in nature and exhibit the richness and complexity characteristic of natural systems. For that reason perhaps it should not surprise us if they are not maximally efficient media of communication. On the other hand, the extent of ambiguity in natural languages is such that they are not just less than maximally efficient, but strikingly inefficient.

Computational and psycholinguistic theories of language processing alike consistently assume or entail that ambiguity makes the task more difficult for the processor. This is a natural assumption: finding multiple interpretations of an expression seems *a priori* more difficult than finding just one, and choosing the intended one among many is an additional task. The first chapter of Manning and Schütze (1999) has a section entitled "The Ambiguity of Language: Why NLP Is Difficult", which includes the following (pp. 17–18):

> An NLP system needs to determine something of the structure of the text – normally at least enough that it can answer "Who did what to whom?" …Therefore, a practical NLP system must be good at making disambiguation decisions of word sense, word category, syntactic structure, and semantic scope.

They go on to discuss some of the difficulties that researchers have encountered trying to build such systems.

Their observation is far from new. Church and Patil (1982; 139) called a section of their paper, "Ambiguity is a practical problem." From a more psychological perspective, Frazier (1985; 135) wrote, "One source of processing complexity that has been recognized and studied for a relatively long time is ambiguity." In fact, the identification of ambiguity with increased processing complexity is widely taken as simply an obvious background assumption. And this makes sense: it is hard to imagine a procedure for assigning meanings to strings that would not require more steps to deal with ambiguous strings than unambiguous ones[5].

In light of this, it appears extremely puzzling that natural languages are so ambiguous. Languages are systems that constantly change. One would expect that at least some of the changes languages undergo would be functionally driven. That is, fitness as a vehicle of communication ought to influence how

[5] Church and Patil (1982) suggest one such procedure, at least for parsing. Roughly, their idea is that the role of the parser is to rule out trees that are **not** compatible with the string. In a suitably constrained universe of possible trees (for example, assuming strictly binary branching), this may actually be less complex than constructing the compatible trees. And it would obviously make parsing unambiguous expressions require more work than parsing ambiguous ones. However, it only deals with part of the task facing real language processors. Even with access to all possible parses, the comprehender must decide in any given situation which one the producer intended. And this is obviously easier if there is only one parse.

languages evolve. And if ambiguity significantly complicates the task facing the language comprehender, languages should evolve so as to reduce ambiguity.

We know of no evidence that language evolution has made languages less ambiguous. There have been suggestions in the literature (see, e.g. Frazier 1985; 145) that certain attested changes were motivated in part as a way to avoid excessive ambiguity – for example, the idea that the relatively fixed word order of modern English (as compared with earlier stages of the language) was triggered by the loss of case inflections. According to this reasoning, with neither case marking nor fixed word order, English would be too ambiguous, having no way to mark who did what to whom. Even if this reasoning is correct, however, the change in question did not reduce ambiguity; it only kept the language from becoming even more ambiguous.

Our central question, then, is this: Why are natural languages so ambiguous? We claim it is a fundamental question that linguists ought to be addressing more directly. We cannot at present answer this question. But in the next section, we will consider a range of possible answers that might be offered.

4. Possible Objections

Not everyone agrees that the pervasiveness of ambiguity in natural language is puzzling. In fact, over half a century ago, Zipf (1949; p. 27) asserted, "we may expect that at least some words must have multiple meanings". He arrived at this conclusion on the basis of a very simple argument, whose premise is what he called the Principle of Least Effort: "a person…will strive to solve his [sic] problems in such a way as to minimize the *total work* that he must expend in solving *both* his immediate problems *and* his probable future problems" (p. 1; italics in original). Zipf's argument went as follows. Suppose there are m possible meanings to be expressed. From the point of view of the speaker, work is minimized if there is only one word expressing all m meanings, since the speaker then doesn't need to think about what word to use to express any particular meaning. From the point of view of the hearer, work is minimized if there are m distinct words (one for each meaning), since this takes the guesswork out of determining the speaker's meaning. These two competing interests (Zipf called them "the opposing Forces of Unification and Diversification") would lead to a compromise in which the number of words is more than one but less than m. This entails that ambiguity should exist.

Even if we assume that the Principle of Least Effort is correct, there are grounds for questioning Zipf's argument. For one thing, any speaker who used the same word for every meaning will be faced with a great many "future problems", as a consequence of almost always being misunderstood. Hence, it is by no means clear that the speaker's effort would be minimized by having only one word. Zipf's formulation suggests that only listeners benefit from successful

communication, but that is patently false.

Moreover, the assumptions that the number of possible meanings is finite and that each meaning is expressed by exactly one word are both at odds with the facts of natural language. In order to assess whether Zipf's argument goes through in a more realistic scenario, one would have to determine how much effort was involved in composing meanings in multi−word expressions, as compared with storing and retrieving many individual words with associated meanings. More sophisticated modeling of this sort is now possible, thanks to the advent of computers. We return briefly to such work in the next section.[6]

Others have told us that they find the ubiquity of ambiguity unremarkable because in normal situations of language use, non−linguistic information usually suffices to disambiguate. Indeed, most people don't even notice the potential ambiguities in most utterances. For example, the sign (purportedly spotted in Heathrow airport) *Dogs must be carried* could, in principle, mean that only people carrying dogs are allowed in that location (analogous to *Shoes must be worn*) or, more naturally, that dogs are forbidden there except when they are carried[7]. But commonsense knowledge of the world keeps anyone but a linguist from noticing the former interpretation. If the silly interpretation is never entertained, then (according to the skeptics) its existence doesn't add to processing difficulty; hence we have no reason to expect language evolution to disfavor such ambiguities.

There is some psycholinguistic evidence suggesting that speakers do not try to avoid ambiguity in production (see Arnold, et al, ms, and Ferreira and Dell 2000). On the other hand, there is other evidence suggesting that at least some ambiguities do increase processing complexity for listeners. The most obvious is the existence of garden paths of various sorts (see, for example, Tanenhaus and Trueswell 1995), in which temporary misanalysis of an ambiguity causes noticeable disruption to comprehension. In the lexical domain, Swinney (1979) showed that all meanings of a word are initially accessed during comprehension; this means the processor has more work to do when it encounters an ambiguous word.

But even if the psycholinguistic evidence were unequivocal, an explanation would be called for. It is not enough to say that people use non−linguistic evidence to block spurious ambiguities. Nobody knows **how** we do it. Theories of language comprehension that are explicit enough to make predictions about

[6] Zipf's reasoning about the distribution of meanings actually went well beyond what we discuss here. Building on his best known result – that the product of the frequency and rank (with respect to frequency) of any word in a corpus is constant – he derived a prediction that the number of meanings of a word would be proportional to the square root of the word's frequency. His corpus investigations indicated that this prediction was correct. We do not know whether this prediction has survived further empirical scrutiny.

[7] In speech, the two interpretations would have different stress patterns, but the written sign is ambiguous.

what people will find easy or hard (for example, Gibson 1999) consistently predict that ambiguity should increase processing complexity. In the absence of a model of how we use context and commonsense knowledge to avoid ambiguities, saying that we do is simply to reformulate the question, not to answer it. Moreover, a full answer should provide insight into **why** we are so good at disambiguation. Is our talent for disambiguation a by−product of our general reasoning abilities, or did it develop in response to the ambiguity of language? If the latter, how did ambiguous languages emerge in the first place?

As a matter of fact, ambiguities are a source of real and potentially costly confusion in communication. We know of no study which systematically evaluates the degree to which theoretically present ambiguities cause confusion or otherwise hamper the process of communication. What we do know is that in our own experience, ambiguities present challenges to speakers, listeners, writers and readers. Who has not ever been faced with the clarification question "Left?" while giving directions, and then been momentarily phased by the inherent ambiguity of the natural confirmatory answer "Right!"? What writer has never struggled with how to reword a text so as to make it unambiguous? What teacher has never set a test only to find several students returning quite naively an answer to a completely different question than was intended? Or to take a case of speech act ambiguity, who has never mistakenly provided an answer to a rhetorical question?

Although we are unable to give a quantitative estimate of the degree to which ambiguity causes problems in communication, we can cite real world examples of ambiguities that have been problematic, taken from the domains of politics, religion, medicine, and law.

Example 1: Message from the President.
In an investment treaty between the US and Azerbaijan dated 9−12−2000, the President of the United States, or someone acting on his behalf, writes: "During a review of the Treaty in preparation for its submittal to the Senate for advice and consent to ratification, the Parties determined that there was an ambiguity in the Annex. This ambiguity reflected a misunderstanding regarding whether Azerbaijan had taken an exception from its national and MFN treatment obligation for insurance services. To resolve this ambiguity, the Parties agreed in an exchange of notes to amend the Treaty." Clearly an unintended ambiguity was deemed of such import that it deserved the attention not only of diplomats and lawyers, but also of the President and Congress.

Example 2: An anatomical ambiguity.
In a stimulating article on Hypokinesia and Hyperkinesia (in C. Goetz, E. Pappert and B. Schmitt, Eds., *Textbook of Clinical Neurology,* W B Saunders Co., 1999), Stanley Fahn writes: "Another problem with the term extrapyramidal is the different ways it is used, creating ambiguity and possible misunderstanding.

Although Wilson conceived of the term to refer to all types of hypokinetic and hyperkinetic disorders, psychiatrists today usually use the term extrapyramidal side effects to represent only drug−induced parkinsonism." This example is of course testament to the fact that even technical terms are not immune from developing multiple meanings.

Example 3: A good catholic upbringing
Theologian James Arthur (in *The Ebbing Tide. Policy and Principles of Catholic Education*, Leominster 1995) writes: 'The use of the phrase 'religious education' by the National Project in a non−confessional sense has led to ambiguity, and out of this ambiguity has grown misunderstanding and confusion about the aims of Catholic education." The question here is whether a religious education is one which involves religious practice, or merely one which involves instruction in religion.

Example 4: Who cleans up?
One of us once served as an expert witness in a legal case involving ambiguity of the expression *net of* in the following text from a formal contract: "Buyer is willing to buy the property at the Purchase Price, **net of** the full cost of remediating the Hazardous Materials to the satisfaction of whichever governing authority has final say on the remediation of the problem". While one party, the original land−owner, claimed that this text implied that the clause in question required the buyer to pay for the cost of cleaning up the site, the buyer claimed the reverse, or at least that there was an ambiguity present.

With the help of linguistic analysis, the buyer was able to argue successfully that there was an ambiguity present, citing amongst other argumentation a number of naturally occurring examples of the expression *net of*, some implying an increase in the amount paid by the buyer, and some implying a decrease in the amount received by the seller.

For example, in a form found in a web−search the following text was found, in which *net of* implies an increased payment by the buyer " The payments − **net of** bank charges − for registration fee, hotel deposit, social and tourist programmes should be made in US Dollars by...." But compare this to another document found online which states: "Respondent shall be bound to purchase the Offered Interest in accordance with such offer or contract, and Movant shall be bound to sell the Offered Interest on the terms and conditions set forth in Movant's written notice except that the purchase price paid by Respondent to Movant shall be **net of** any loans or other indebtedness, including accrued interest, created under the terms of this Agreement, owed by Movant to Respondent." Comprehending this remarkable sentence requires a great deal of effort, but the upshot appears to be that when Respondent pays Movant the financial goods in question, Respondent can deduct the money which Movant owes. So here "the purchase price ... net of X" can be paraphrased as "the purchase price minus the

amount of X".[8]

Confusing as this verbiage is, the point is clear: *net of* is ambiguous, and the choice of interpretation had expensive consequences. More generally, the four examples given are by no means atypical of the real problems for communication created by the ambiguity of English.

Other skeptics might question our claim that languages are not becoming less ambiguous. This part of our argument was rather stipulative, in that we presented no direct evidence for it. In fact, we didn't even present a metric to determine how ambiguous languages are. The sort of anecdotal argument we offered for the highly ambiguous character of English provides at best a rather coarse−grained measure. Suppose ambiguity is in fact diminishing, only very gradually. In the absence of more precise tools, we might well not notice.

This is certainly possible, but we do not consider it very plausible. The sorts of syntactic changes that are well attested in languages include some, such as changes in allowable word orders, which could in principle reduce ambiguity. For example, if the canonical ordering of postverbal NPs and PPs in English were reversed, some PP attachment ambiguities would be eliminated. But known word−order changes, such as the change of English from a verb−final to a verb−medial language, had no evident ambiguity−reducing effect. Moreover, some common types of language change – specifically loss of case, tense, and agreement inflections – tend to increase, rather than decrease ambiguity.

In the lexical domain, it seems particularly implausible to argue that ambiguity reduction is happening, but too slowly for us to notice. Words commonly change their meanings, as is evident from even a superficial examination of an etymological dictionary. In a language like English, whose history over many centuries is well documented, any systematic reduction of polysemy ought to be manifest.

5. Possible Answers and Directions for Future Research

[8] To exploit another ambiguity, the people who really cleaned up in this case were lawyers.

A very different type of answer to our question would be to argue that ambiguity actually serves some useful function in language. If this is the case, then whatever advantages it confers on a language might be sufficient to overcome the extra processing burden it entails. But how might ambiguity be useful?

The most obvious answer is that is saves memory. For example, if every word in a language had two meanings, then a speaker of the language would have to memorize only half as many phonological forms as if the vocabulary were unambiguous. Of course, the number of sound–meaning pairings would be the same in the two cases, but learning and remembering fewer forms might still be easier. Whether an ambiguous grammar takes less memory than an unambiguous one is less clear, but certainly not out of the question.

There are three obvious problems with this idea. First, the saving in memory comes at a substantial cost, namely, the extra processing required to disambiguate. Second, the probability of miscommunication is increased, because the comprehender might pick the wrong interpretation. Third, spoken communication is carried out under severe time constraints (see Clark, 1996, Chapter 9). Presumably, retrieving information from memory is generally faster than going through some inferential process to derive it. Moreover, available evidence suggests that the human brain has vast amounts of memory but performs computations relatively slowly. Hence, saving memory at the expense of extra processing looks like a false economy.

A related idea is that ambiguity emerges as a byproduct of something else. In particular, suppose for the moment that long morphemes are hard for people to learn and remember. Then, given restrictions on the number and sequencing of phonemes (dictated by human articulatory and perceptual capabilities), the number of short morphemes that could be constructed would be limited. Now the number of meanings expressible in natural languages is presumably unbounded, which is why they have the combinatory systems of morphology and syntax. These combinatory systems build up complex meanings from semantic atoms of some sort. If the number of such atoms is larger than the number of morphemes that can be readily learned and remembered, then some morphemes will be used ambiguously.

Why should longer morphemes be harder? The most obvious answer is that they take more effort to produce and are more subject to production errors. But the same could be said of long words and phrases, which are extremely common. A more plausible answer is that the component parts of morphemes – that is, phonemes – are not themselves inherently meaningful, and people are known to have severe limitations on their ability to recall strings of meaningless elements. Specifically, in the words of Miller (1956), the limit is "the magical number seven, plus or minus two". As Miller observed, we overcome this "informational bottleneck" by "organizing the stimulus input simultaneously into

several dimensions and successively into a sequence of chunks" (p. 95). That is, we build complex words and phrases out of short, simple morphemes, but the number of morphemes available in a language is constrained by the number of phonemes, restrictions on the sequencing of phonemes, and Miller's "magical number". Hence, if a languages needs a large inventory of atomic meanings, it will have ambiguous morphemes.

This explanation of ambiguity makes some testable predictions. In particular, it entails that morphemes longer than seven phonemes should be extremely rare in the languages of the world. This strikes us as plausible, though we know of no systematic study of morpheme length across languages. It also suggests that short morphemes should generally be more ambiguous than long ones, since the short ones impose the least burden on memory. Again, we know of no study of this exact question. Zipf (1949) claimed that more frequent words tend to be both short (p. 64) and ambiguous (p. 30), suggesting a likely correlation between word length and ambiguity. If the same is true of morphemes, it would lend plausibility to the explanation based on constraints on morpheme length.

The notion that ambiguity arises from a constraint favoring shorter morphemes, then, seems plausible as a possible explanation, but it is as yet untested. Furthermore, there is an obvious objection: if shorter morphemes are favored, then why are syllables that are not morphemes (e.g., *foo, wug, blick*) easy to think of? The force of this objection depends on how common such syllables are, but we will not pursue the matter here.

The idea of explaining syncretism and much lexical ambiguity on the basis of a constraint on morpheme length strikes us as a promising avenue for further research. We hasten to add that it cannot explain syntactic or scopal ambiguity.

A rather different tack is to consider whether ambiguity can serve some communicative function. We can envision two types of cases in which a speaker might benefit from using ambiguous language. For the first case, imagine a situation in which there are two distinct dialects spoken in neighboring geographical regions, and suppose there is some expression E that has different denotations in the two dialects. People living close to the border between these two regions will have occasion to interact with speakers of both dialects, so it will serve their communicative needs to allow E to have both denotations. The extra cost of having to disambiguate (presumably by identifying the dialect of one's interlocutor) could easily be offset by the value of being able to communicate successfully with a wider range of people.

This sort of situation is by no means unusual. Closely related dialects often do assign distinct denotations to a given expression (e.g. the term *potato chip* in American English denotes what British speakers call a *crisp*, whereas in British English it denotes what Americans call a *French fry*). And regions of overlapping dialects are the rule, rather than the exception, when one considers that many dialect differences are based on ethnic, racial, and class distinctions, rather than geography.

Thus, the heterogeneity of speech communities may provide a powerful motivation for the ambiguity of natural language. However, this argument seems far more persuasive as a reason for the existence of lexical ambiguity than for syntactic or scopal ambiguity. We know of no convincing evidence for dialect differences with respect to these latter types of ambiguity that could motivate their existence.

The second type of situation in which ambiguity might be useful is also one in which a speaker wants to be understood as saying one thing by one set of listeners and as saying something else by another set of listeners. These might be cases of negotiations, advertising, or simply ways of avoiding socially awkward situations. For example, a guest who finds a meal unpalatable but does not wish to offend her host could say *Nothing is better than your cooking*, knowing that it will be interpreted as praise, though really meaning that she would have preferred to go hungry. A real–life example is the standard use of the term *lifetime guarantee*. Most consumers probably assume this means that the product is guaranteed for their lifetimes, but most products carrying such a guarantee are covered only for the lifetime (or, in some cases, the normally expected lifetime) of the product[9]. The term *lifetime* is not vague, but it is ambiguous as to which lifetime is denoted.

Are these potential communicative advantages of ambiguity sufficient to outweigh the disadvantages we have discussed? This question can only be answered in the context of an explicit model of language evolution. The kind of informal weighing of hypothetical costs and benefits we have been engaging in is simply too imprecise to support any sophisticated predictions. Explaining why ambiguity remains pervasive will require better analytic tools.

In recent years, a number of scholars (e.g. Nowak, 1999, 2000; Briscoe, 1997; Kirby, 2001; Oliphant & Batali, 1997; see Kirby 2002 for an overview) have begun to investigate the evolution of language using modeling techniques borrowed from computational genetics. Almost none of this work addresses the question of ambiguity, and much of it adopts the simplifying assumption that expressions have unambiguous meanings. But the method is well suited to investigating the sorts of questions we have posed in this paper.

In collaboration with Aviv Bergman and Thorsten Brandts, we have begun to model highly simplified versions of the kinds of factors discussed above that might influence the survival of ambiguity as languages change. This work is still in its early stages, and it would be premature to report on it in any detail. Nevertheless, preliminary indications are promising.

Our models involve considering extremely simple "languages" consisting of only a few words, some of which are ambiguous and some of which are unambiguous. The degree of ambiguity differs for different types of language users. Fitness is linked to successful communication. The models involve

[9] Some states have adopted laws requiring this to be spelled out, but it is generally relegated to the fine print that few consumers read.

numerous (continuous) parameters, such as the degree of ambiguity of different words for different individuals, the initial distributions of various types of language users in the population, the relative frequency in the world of the objects or events taken as the denotations of the words, and the degree to which successful communication enhances fitness and miscommunication reduces it. By varying the initial values for these parameters and running our models through hundreds of generations, we can learn what combinations of settings lead to the reduction or disappearance of ambiguity, and what combinations support the maintenance of highly ambiguous language.

It appears that there are some combinations of initial parameter settings that preserve and even increase ambiguity. For instance, ambiguity is maintained in circumstances in which one possible denotation for a word is far more common than another. This finding is reflected in empirical data taken from a semantically tagged portion of the Brown corpus. We examined words that occurred at least ten times in the corpus and exhibited exactly two denotations in those occurrences. These words typically are used with one meaning much more than with the other. Almost 75% of all words with two possible meanings use one meaning 75% or more of the time; over 30% use one meaning 95% or more of the time. This suggests that the relative proportion of each denotation in ambiguous words may play a key role in preserving that ambiguity, possibly because of the effects prevalence has on interpretability.

Less suprisingly, heterogeneous language use in a population may also support the emergence and preservation of ambiguity. In some simulations, agents whose lexicon included a given signal performed better than agents who lacked it −− even if the signal was ambiguous. Why? As long there were at least a few individuals who use a particular signal differently from the rest of the population, tolerance of ambiguity in the interpretation of that signal (within reasonable limits) was advantageous. The agents who had ambiguous lexicons (i.e. who could interpret at least some signals they heard in more than one way) could understand others more often than could those who had only unambiguous interpretations available. The moral of this seems to be that since an ambiguous word is better than no word at all, ambiguous words will tend to enter the language. This may therefore be a route by which ambiguity emerges, since it is plausible that there will always be a few individuals who do not share the same signal/denotation assignment as the population in general (for instance, those who are just learning the language).

There are still many avenues to be investigated using these sorts of methods. It is our hope that they can inform future work in psycholinguistics and historical linguistics and ultimately lead to an explanation for the puzzle of ambiguity in natural language. In any event, we hope to have shown that this puzzle is worthy of more attention from linguists.

References

Arnold, J., T. Wasow, A. Asudeh, and P. Alrenga (ms) Avoiding Attachment Ambiguities: the role of Constituent Ordering.

Briscoe, E.J. (1997) Coevolution of language and of the language acquisition device. *Proceedings of the 35th Association for Computational Linguistics*. Morgan Kaufmann

Church, K. and R. Patil (1982) Coping with syntactic ambiguity or how to put the block in the box on the table. *American Journal of Computational Linguistics*, 8; 139–149, July–December.

Clark, H. H. (1996) *Using Language*. Cambridge: Cambridge University Press.

Ferreira, V.S., & Dell, G.S. (2000). Effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive Psychology*, 40, 296–340

Frazier , L. (1985) Syntactic complexity. *Natural Language Parsing*, ed. D. Dowty, L. Karttunen, and A. M. Zwicky, 129–189. Cambridge: Cambridge University Press.

Gibson E (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition* 68: 1–76.

Kempson, R. M. and A. Cormack (1982) Ambiguity and quantification. *Linguistics and Philosophy*, 4, 259–309.

Kirby, S., and J. Hurford (2001) The emergence of linguistic structure: An overview of the iterated learning mode. In A. Cangelosi & D. Parisi (eds), *Simulating the Evolution of Language*. Springer (6) 121–148

Kirby, S. (2002) Natural language from artificial life. *Artificial Life.*

Manning, C. D. and H. Schütze (1999) *Foundations of Statistical Natural Language Processing*. Cambridge, Massachusetts: MIT Press.

Martin, W. A., K. W. Church, and R. S. Patil (1987) Preliminary analysis of the breadth–first parson algorithm: theoretical and experimental results. In L. Bolc (ed), *Natural Language Parsing Systems*. Berlin: Springer Verlag.

Miller, G. (1956) The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *Psychological Review*. 63: 81–96.

Montague, Richard (1974) *Formal Philosophy*, R.H. Thomason (ed). New Haven: Yale University Press.

Nowak, M., and D. Krakauer (1999) The Evolution of Language. *Proceedings of the National Academy of Sciences*. 96:8028–8033

Nowak, M., J. Plotkin, V. Jansen. (2000) The evolution of syntactic communication. *Nature*, 404:495–498

Oliphant, M. & J. Batali. (1997) Learning and the emergence of coordinated communication. Center for Research on Language Newsletter, 11

Stump, G. (2001) *Inflectional Morphology: A Theory of Paradigm Structure*. Cambridge: Cambridge University Press.

Tanenhaus, M.K., & Trueswell, J.C. (1995). Sentence comprehension. In J. Miller and P. Eimas (Eds.), *Handbook of Perception and Cognition, Vol. 11: Speech, Language and Communication*, pp. 217–262. New York: Academic Press.

Zipf, G.K. (1949) *Human Behavior and the Principle of Least Effort*. New York: Addison–Wesley.

Zwicky, A., and J. Sadock. (1975). Ambiguity tests and how to fail them. In J. Kimball (ed), *Syntax and Semantics*, vol. 4. New York: Academic Press.