Contributions of Lexicography and Corpus Linguistics to a Theory of Language Performance

Patrick HANKS, Oxford, UK

Twentieth-Century Linguistic Theory

What would a theory of language performance be like? Linguistic theory in the 20th century has been dominated by the notion of competence. For forty years and more, following publication of Noam Chomsky's *Syntactic Structures* in 1957, attention was focused on competence: the underlying logical structure of language, and the astonishing ability of human beings to pick up a language in early childhood and to construct sentences in it that have never been uttered before, but which are nevertheless intuitively recognizable by other users of that language as syntactically well formed.

What is the relevance of transformational-generative linguistic theory to lexicography?

Chomsky has given us revolutionary insights into many aspects of the nature of language, and it is therefore with due diffidence that I suggest that at the root of Chomsky's work lies a claim which has been responsible for much confusion and wasted effort. On pages 16-17 of *Syntactic Structures* we read:

> Evidently, one's ability to produce and recognize grammatical utterances is not based on notions of statistical approximation and the like. The custom of calling grammatical sentences those that "can occur," or those that are "possible", has been responsible for some confusion here. It is natural to understand "possible" as meaning "highly probable" and to assume that the linguist's sharp distinction between grammatical and ungrammatical is motivated by a feeling that since the 'reality' of language is too complex to be described completely, he must content himself with a schematized version replacing "zero probability, and all extremely low probabilities, by *impossible*, and all higher probabilities by *possible*." We see, however, that this idea is quite incorrect. . . .Despite the undeniable interest and importance of semantic and statistical studies of language, they appear to have no direct relevance to the problem of determining or characterizing the set of grammatical utterances.

There are many things that could be said about this seminal passage, among them:

1. People somehow derived from Chomsky's early work the notion that the proper task, perhaps the only proper task, of linguistics is to devise a machine that, in theory at least, can generate all and only the grammatical utterances of a language. Why should this have been?

In the passage quoted Chomsky acknowledges, however grudgingly, that other kinds of linguistic studies may be interesting and important. Indeed, his own teacher, Zellig Harris, undertook statistical studies. Nevertheless, for 30 years after *Syntactic Structures* was published, linguistics, especially in America, placed great emphasis on syntax, while tending to neglect semantics, lexis, and other aspects of language study. When the Chomskyans arrived at the study of lexis, they brought with them a vast theoretical apparatus built up over decades to deal with issues in syntax. This seems to have interfered with the objective analysis of the actual behaviour of words in use. Lexis and semantics were processed as ancillaries of syntax. The problem was compounded by lack of evidence. In the absence of objective evidence, introspection was appealed to instead. But studies in corpus linguistics have shown that introspection is a very flawed technique. Corpus studies indicate that there is an inverse relationship between cognitive salience (what we can come up with by means of introspection) and social salience (what we find in corpora). We human beings are wired to register the unusual in our minds, generally in a way that is available to conscious recall. But we fail to pay any attention to the commonplace patterns of usage on which we rely so heavily in our everyday communications. If you do not know the term 'hermeneutics' and someone tells you about it, you may remember not only what it means, but also all the circumstances associated with your acquisition of that term. On the other hand, only a lexicographer would pause to ask what exactly you have to do to a photograph to take it, while ordinary English speakers asked to list the most common meanings of **take** never include expressions of time in their lists ("How long will it take", "It only took a few minutes"), although general English corpora show this to be extremely common. We register, and can recall, the unfamiliar new word, while passing over in silence the familiar.

2. More significant is Chomsky's insistence that the dividing line between grammatical and ungrammatical utterances is a sharp one. For the Chomsky of 1957, there are no intermediate cases. This was, and is, a controversial claim. There can be no dispute that some strings of words are totally ungrammatical, while others constitute well-formed sentences. But this does not entail that there must be a sharp dividing line between grammatical and ungrammatical. Indeed, a contrary claim has been coming out of corpus linguistics in recent years, to the effect that some sentences are more grammatical than others. Grammaticality, according to corpus linguists, is a gradable. Gradability cannot be argued away by appealing to a distinction between competence and performance. Many utterances in the grey area of grammaticality cannot be classified as performance errors, but must rather be viewed as exuberant *exploitations* of the conventions that constitute our linguistic competence. This implies that a theory of exploitations, alongside a theory of convention, is needed to explain human linguistic behaviour.

3. Chomsky makes it clear here and elsewhere that he is concerned literally with what is possible in a language, no matter how unlikely the possibilities may be. It is time to shift our focus from what is possible in language to what is probable, and to look at the theoretical consequences of such a shift.

4. Also buried in the quote from Chomsky is an assumption that linguistic theory represents some psychological reality, i.e. that it represents processes that really do go on in-

> side the human head when sentences are generated and deconstructed. Chomsky couches his explanation in psycholinguistic terms: a person's "ability to produce and recognize grammatical utterances." But the psychological reality of Chomsky's model has been frequently questioned. One such challenge comes from the field of cognitive linguistics, in particular the work of Ronald Langacker.

Cognitive Linguistics

In *Concept, Image, and Symbol: the Cognitive Basis of Grammar* (1990), Langacker asserts that:

> The ultimate goal of linguistic description is to characterize, in a cognitively realistic fashion, those structures and abilities that constitute a speaker's grasp of linguistic convention. A speaker's knowledge is procedural rather than declarative, and the internalized grammar representing this knowledge is simply a structured inventory of conventional linguistic units.

A dictionary, too, is a form of linguistic description. A monolingual dictionary is, at its simplest, a structured inventory of a set of linguistic units, namely words: units that are conventional in form, are used in conventional ways, and have conventional meanings. We should look to theory, therefore, to tell us something about the relationship between units (in our case, words) and meanings.

Some questions arising from Langacker's formulation are:

1. What is the relationship between the dictionary in people's heads and the dictionary on the page?
2. What is the relationship between word meaning and words in use (i.e. between words and the procedures of making meanings)?
3. How are we to regard the "definitions" in a dictionary, if knowledge of a language is procedural rather than declarative?

I will discuss each of these points in turn, in a little more detail.

1. What is the relationship between the dictionary on the page and the mental lexicon? Each term in the mental lexicon serves as a node or focus for a variety of memories, beliefs, perceptions, and conceptions. The connectivity of a dictionary entry is necessarily restricted to other words, but the connectivity of the mental lexicon is not so restricted. No doubt the "meaning" of each term in the mental lexicon is subtly different for each user of the language, but we have no way of knowing precisely what each term means to each person. In *Word and Object* the American logician W.V.O. Quine comments:

> Different persons growing up in the same language are like different bushes trimmed and trained to take the shape of identical elephants. The anatomical details of twigs and branches will fulfil the elephantine form differently from bush to bush, but the overall outward results are alike.

A dictionary on the page represents the meaning of words only in terms of other words. There is no place in a dictionary for stored memories of physical sensations, no place for remembered sights, sounds, smells, or emotions all of which are associated with the mental lexicon but only for words and, occasionally, pictures. The dictionary attempts to represent in words only the common convention which speakers rely on in order to communicate with one another. In the case of the great historical dictionaries, it is assumed that everybody knows the common convention of what words mean (or else that conventional meaning is unknowable or inexpressible). So the great historical dictionaries of the world's languages focus on saying where the modern meaning of a word came from, how it developed, without ever being very explicit about what the modern meaning is.

2. What is the relationship between word meaning and words in use? It came as a shock to some people in the Natural Language Processing community that many dictionaries, especially those first used in NLP laboratories, have nothing to say about word use. They list many sense of a word, but rarely explain how a user is supposed to distinguish one meaning of a word from another. It is, of course, not the case that word senses are freely interchangeable. To take a time-worn example, it is possible, but preposterous, to say, "John swam to the bank" and mean "John visited a financial institution by swimming." Linguists have got into the habit of deploying immense ingenuity in constructing scenarios in which preposterous interpretations become plausibleons become plausible for example, if the High Street were flooded and under 6 feet of water, John might have swum to the bank but such ingenuity does not make the interpretation any less preposterous if it is seriously presented as representing usage. It is possible, though preposterous, to talk of snails galloping: this is a grammatically well-formed sentence of English, however odd its meaning might be. And, of course, it was Chomsky himself who first pointed out that the preposterous sentence "Colorless green ideas sleep furiously" is syntactically perfectly well formed. In recent years, it has become customary in linguistics to talk about "selectional restrictions", e.g. there is a selectional restriction on the verb **gallop** such that the subject must be a horse or inflation. But the term *selectional restriction* is a rich potential source of further confusion. Given that there is a literally infinite number of sentences that are possible but unlikely, a more accurate term is *selectional preferences*. The verb **gallop** prefers, but is not restricted to, subjects that are horses or inflation. Even galloping snails are grammatically possible, though in practice unlikely. Corpus linguistics has increased our awareness, not only of the overwhelming frequency of some preferences, but also of the uneven nature of the middle ground, and of the exuberance with which speakers occasionally utter, for rhetorical effect, non-preferred options sentences that lie deep in the grey area between grammatical and ungrammatical. This apparently trivial point lies at the heart of the distinction between linguistic performance theory and competence theory.

3. How are we to regard the "definitions" in a dictionary? A practical example may help to illustrate the problem. The *New Shorter Oxford English Dictionary* says that a spider is

"an arachnid . . . having a narrow-waisted body and eight jointed legs." Now, this statement does not constitute a necessary condition: a fat spider with only seven legs is still a spider. Nor does it constitute a sufficient condition: a narrow-waisted creature with eight legs might be a thin octopus. Of course, by using the word *arachnid* as the genus word in the definiens, the editors of the Shorter were committing themselves to the principle that a definition should define, i.e. that it should constitute a set of a necessary and sufficient conditions, a decision procedure for determining what is and what is not a spider. But in fact, what they committed themselves to was a tautology. A spider is an arachnid and with a few exceptions such as mites, ticks, and scorpions an arachnid is a spider. The really informative part of the definition, "having a narrow-waisted body and eight jointed legs" constitutes a sort of appendix to the defining term "arachnid". If the purpose of saying "with eight jointed legs", is to distinguish spiders from other kinds of arachnids, it fails, for all arachnids have eight jointed legs. But of course that is not really the purpose at all: the true purpose is to inform. Saying that a spider has eight jointed legs is descriptive, informative, and helpful. Saying that it is an arachnid is a dutiful nod in the direction of zoological taxonomy which conveys no information about spiders to anybody. Those who know and care what an arachnid is will already know that arachnids have eight jointed legs. Those who don't, for the most part won't care. There may, of course, be a few readers specifically interested in taxonomic hierarchies in the life sciences, who will be glad to know that a spider is classified as an arachnid, but they are in a tiny minority.

My comments are not intended to be critical of the *New Shorter* a dictionary which comes from a stable for which I was formerly responsible, and which in my view represents the finest available example of traditional historical lexicography. It behoves us to be full of admiration for the ingenuity of the *New Shorter* lexicographers in applying traditional criteria of classification and substitutability to definition writing. But should they have done it at all? The point of the discussion is, not to criticize the execution of any one dictionary, but to open a debate about the underlying theoretical assumptions on which European lexicography is based.

Three main points emerge from all this:

1. Classifying objects in the world (e.g. classifying spiders as arachnids) is not the same as explaining what words (e.g. 'spider') mean. Saying that a spider is an arachnid does not explain anything, and it may not be helpful to anyone. We should not imagine that, by classifying a term, we have explained anything.

2. If knowledge of a language is procedural rather than declarative, then meanings are events, not objects, and dictionary definitions are not statements of meaning, but rather organized lists of 'meanings potentials'. They represent an idealized and partial verbalization of something that is available in our heads, ready to be drawn on by speakers and writers to make meanings. This account of definitions and meanings goes a long way to explain the difficulties encountered by the Hector project, the Senseval lexicographers, and others who have attempted to map real examples of language in use onto dictionary definitions. An important constraint of such projects has been to account for *all* uses of the given word in the given sample, not just those which best suit the lexicographer's

purpose. Some perfectly standard-seeming uses do not match the dictionaries very well, and yet they do not always provide sufficient motivation for rewriting or adding to the dictionary.

3. To this extent at least, the dictionary functions in a similar way to the mental lexicon. Both contain inventories of conventional linguistic units which are available for use in making meanings, but utterers do not follow the conventions slavishly; rather, they exploit them, to say new and interesting things in new and interesting ways.

In the words of Dwight Bolinger:

> A dictionary is a frozen pantomime, . . . a nosegay of faded metaphors.
>
> Dictionaries do not exist to define, but to help people grasp meanings, and for this purpose their main task is to supply a series of hints and associations that will relate the unknown to something known. Dictionaries do not exist to define, but rather to provide a series of hints and associations connecting the unknown with the known.

Much modern monolingual lexicography, especially lexicography for foreign learners, is concerned with identifying and describing linguistic conventions (or at least a very large subset of them). and distinguishing them from the accidental outcomes of the procedures of using language (i.e. non-conventional uses of a kind which can readily be found in large corpora and are eagerly seized on by citation readers).

An example:

> **Ligger** is defined by slang dictionaries as "a freeloader in the music industry" This has given rise to the term **liggerati** in some circles, denoting a freeloader who is also a celebrity or member of High Society. Should we add **liggerati** to the inventory of conventional units of English? Probably not, unless evidence is also adduced that it is now in conventional use, whether in slang or in more formal registers of English. Caution, in contradistinction to the wishes of marketing managers the world over, urges us to classify **liggerati**, for the time being at least, as an exploitation of the established term **glitterati**, denoting fashionable people in general, and the less well-established slang term **ligger**. Exploitations like this are commonplace; they are everywhere about us, if we care to look. Exploiting conventions is part of normal human linguistic behaviour, posing endless challenges for lexicographers.

Langacker also has something to say about definitions:

> Cognitive grammar ... assumes that a frequently used morpheme or lexical item has a variety of interrelated senses. They can be thought of as forming a network, where some senses are prototypical, others constitute either extensions or specializations of a prototypical value or of one another.
>
> Cognitive grammar assumes that meanings are always characterized relative to cognitive domains, i.e. knowledge structures or conceptual complexes of some kind.

This raises another theoretical issue of profound importance, to dictionary making, to corpus linguistics, and to performance theory alike. How much knowledge can a dictionary writer, or indeed any other human being trying to explain something to others, expect the reader or hearer to have? A dictionary definition of a term in cricket, for example **googly**, cannot be written in such a way that it explains the meaning accurately to someone who has no knowledge at all of what goes on in cricket. It is legitimate to assume that a reader looking up the term **googly** has at least some idea of what cricket is and in particular what bowling is in cricket. It is not only legitimate, but also unavoidable, to use the verb **bowl** in defining cricketing terms and to expect the reader to know what it means or, if not, to find out by looking it up. **Bowl** is a more general term than **googly**, so it must be explained in language that is more accessible to laypeople. And when we come to look for evidence of how the word **googly** is used in English, we need to be able to look in corpora of writings about cricket before we confuse ourselves with the metaphorical uses of **googly** that can be found in, say, reports of proceedings in parliament or writings about business transactions. Large corpora should not really be thought of as homogeneous wholes, but rather as sets of overlapping subcorpora. For human language in use is very domain-specific.

I conclude this discussion of psychologically real linguistic theory with two more quotes from Langacker. The first is taken from a discussion of the nature of meaning:

> It is common for linguists to assume (often tacitly) that all the meanings of a lexical item must be predictable from a single basic sense, and that separate lexical items must be posited when no such meaning can be found. This is an unwarranted assumption that creates more problems than it solves. The network model is far more realistic and descriptively adequate, for it permits and indeed requires all of the following:
>
> (i) a statement of the full array of conventionally established uses;
>
> (ii) a characterization of the relations between individual senses;
>
> (iii) a description (in the form of schemas) of whatever generalizations can be extracted from sets of particular senses.

The next quote lends support to those who, like myself, argue that prototype theory is of immense importance to lexicography and to corpus analysis:

> Traditionally dominant has been the view that a category is defined by a set of criterial attributes, i.e. necessary and sufficient conditions for class membership. ... In fact, recent findings by cognitive psychologists strongly favor an alternative conception: categorization by prototypes, where membership in a category is determined by perceived resemblance to typical instances.

This is a far cry from determining all and only the grammatical utterances of a language, or indeed, regarding dictionary definitions as decision procedures for identifying all and only those creatures which are spiders.

Social Theory

Chomsky was interested in the relations between language and logic, and like most great Western thinkers before him, he assumed that logic underlies language. It is not entirely clear why we should accept that the relationship is this way round. What would it be like if we worked on the hypothesis that logic is a construct just one among many of natural language? Or rather, since there are many logics, that logics are constructs of natural language. I will not pursue this point in any detail here, but one possible benefit of turning the language - logic relationship on its head is that it would free up the study of natural language from the constraint of assuming that linguistic behaviour is necessarily logical (in particular, that it is governed by a particular kind of logic), and that if it isn't some performance error must be involved.

If, instead of seeking the underlying logical structure of sentences, we look at linguistic behaviour as a form of social interaction, then we can link language performance to social theory.

In *Foundations of Social Theory* (1990), James Coleman observes that fashions and tastes are collective processes. Stanley Lieberson summarizes the argument as follows:

> One's choice is affected by the choices that others make, and since this is the case for all others, "there is some kind of dependency among the actions; individuals are not acting independently."

Human language users are not acting independently, and the choice of words to make meanings is determined by collective processes. The selectional preferences of words that are so striking when we look at language en masse, as recorded in a corpus, are as much a matter of fashion as anything else, but fashion with an utterly serious purpose, namely to communicate with, and interact with, other people. When language users flout convention, by exploiting some norm of meaning or belief, they do so for rhetorical effect, in order to get the attention of an audience, or to make a point in a way that will impinge on the audience's consciousness and be noticed and remembered.

Social theorists such as Lieberson and Schelling also account for the rapidity with which social conventions can change. Schelling studied the process by which a racially mixed area can suddenly lose its equilibrium and become segregated:

1. A small number of people from a new ethnic group moves into a neighbourhood.
2. Their presence increases the propensity of other members of that group to move in.
3. There is a decline in the propensity of members of other groups to move in.
4. The propensity of members of other ethnic groups to move out increases.

A very similar mechanism governs the adoption of new linguistic conventions. A currently topical example is the adoption of rising intonation in English declarative sentences among the young, a new convention which causes older English speakers like myself to constantly mistake statements for questions:

1. A small group of English-speaking teenagers, who are perceived by their peers as "cool" adopt rising declarative intonation.

2. This increases the propensity of other teenagers who want to be identified as "cool" to use rising intonation.

3. There is a decline in the propensity of teenagers to continue using falling or flat intonation.

4. Tension arises among users of flat and falling intonation. Some older speakers, finding themselves in a quandary or feeling isolated, begin to use rising intonation.

5. Rather than move out (i.e. give up speaking English altogether), by 2010 everyone will be using rising intonation for declarative sentences. Or maybe not. Maybe the whole process will go into reverse, and disappear as rapidly as it arrived. Predictions in matters of language are hostages to fortune of the most vulnerable kind.

The same model can be applied to almost all forms of linguistic change. For example, the well-known change in the meaning of the word **gay**:

1. A small group of homosexuals use **gay** to refer to themselves and other homosexuals.
2. This increases the propensity of other homosexuals to use **gay** to mean 'homosexual'.
3. There is a decline in the propensity of other English speakers to continue using **gay** to mean 'bright and cheerful'.
4. Very soon, it becomes impossible to use **gay** to mean 'bright and cheerful' without causing a snigger or other comment.

The model also applies to fashions in phraseology, for example the rise of the expression **between you and I**, which is anathema to the few surviving English speakers who have any awareness of traditional grammatical case, but which is now well established in standard English and probably not only impossible to dislodge, but will very soon have driven out **between you and me** completely, except perhaps as a pedantic curiosity used only by old-fashioned purists.

When we look at a corpus and are astonished by the overwhelming and often unsuspected frequency of conventional phraseology, we are looking at traces of thousands of instances of fashionable linguistic behaviour. If we then turn to a historical corpus, we can see how rapidly the conventions of meaning and use can change. The equilibrium of word meaning and phraseological norms is very unstable. In fact, it is constantly changing. It is social theory, not logic, that explains how these changes come about.

Corpus Evidence and Performance Theory

During the past fifteen years, as very large electronic corpora became more and more widely available, corpus researchers began to notice an uncomfortably wide nay, yawning gap between the predictions of linguistic competence theory and the evidence for what actually happens when language is used. Examples are encountered by corpus lexicographers every day, but there is not yet an established theoretical apparatus that enables them to deal with the dichotomy comfortably.

In order to account satisfactorily for language in use, a theory of language performance will be needed, a theory that is statistical and probabilistic, rather than certain and cut and dried. Of

all the many words, uses, and structures that are *possible* in a language, it will show us how to pick out just those that are *normal*, and it will relate other uses to the norms by a theory of exploitations: a set of exploitation rules that will say how a normal use may be exploited to form metaphors and other unusual uses, and what the constraints are. (Norms, of course, may be genre-specific, as well as general.)

Until the advent of large corpora in the 1980s, there was simply no way of analysing the characteristic behaviour of each word in the language. Now we have large corpora, it is time to revisit theory from a lexical point of view, taking account of what can be learned from corpora.

In pursuit of definitions that accurately summarize the unique contributions of words to the meaning of sentences in which they occur, modern lexicographers can now study concordance lines from a corpus. What they find is interesting, and not always expected, even though, in all too many cases, what they find is be determined by what they expect to find. Some lexicographers and linguists have treated the corpus merely as a quarry, a source of examples for what they already 'know'. And very often the corpus obliges. If you look long enough and hard enough, and if you have a large enough corpus, or enough texts of the right kind, you will find what you are looking for. For example, a large historical corpus may yet be found that contains an example or two supporting the notion that the verb **fan** means 'to winnow (grain)'. But that does not mean that this is part of the meaning of the modern word **fan**. In fact, to use a corpus in this way, i.e. to make self-fulfilling prophecies, is precisely what corpus linguistics is *not* about. (This does not prevent lexicographers from doing it, however.) Corpus linguistics, if it is about anything, is about observing the conventions of language in use, and then observing the great variety of ways in which these conventions are exploited. (It is perhaps worth mentioning in passing that a corpus does not, of course, provide direct evidence for meaning; it consists of a record of traces of linguistic behaviour, from which meanings can be inferred.)

Some grammarians have used corpus evidence in a similarly supplementary way. Beth Levin, for example, in compiling her (partial) inventory of English verb classes and alternations, first consulted her intuitions, then (with the help of colleagues) checked the corpus to see if she had missed anything. The result was undoubtedly an improvement on intuition alone, but nevertheless some of the verbs in Levin's classification rarely if ever behave in the way that the classification predicts. The corpus, evidently, was used to supplement intuitions rather than to motivate the analysis, and examples which satisfied intuition but for which no corpus evidence was available were not rejected. But Levin might ask, why should they be? For we must beware of the failure-to-find fallacy: the fact that we have failed to find something does not mean that it does not exist. Against this must be set the line of argument that says that if something does not occur in a corpus of 100 million words equivalent to half a dozen years of hard uninterrupted reading for a normal person then it cannot be very important.

Another example is the COMLEX project (Grishman, McLeod et al.), which describes in detail the *possible* complementation patterns of English verbs. Because the focus of COMLEX is on the possible, not the probable, it is perhaps a less useful tool than it might have been. And I think the COMLEX people recognize this. It is surely no accident that one of the driving forces behind the American National Corpus initiative is Catherine McLeod, who was also one of the prime movers in COMLEX. Her experience on COMLEX was not dissimilar to that of many British lexicographers in the 1970s and 80s. Using their intuitions, she and her colleagues compiled

a work of reference in which no distinction was made between structures for which there was only intuitional evidence and structures for which there is plentiful evidence of actual usage.

Conclusion

Is a theory of language performance in competition with a theory of language competence? It might seem as if the two are necessarily in competition, but in fact they are complementary, and they are further complemented by cognitive theory. Language impinges on every aspect of our being and is central to the vast majority of our everyday activities, so it is hardly surprising that an exceptionally wide range of theories is called for to account for so wide-ranging a phenomenon as human language. For all their desire to present transformational-generative theory as psychologically real, much of the work of the Chomskyans is really about the relationship between language and logic, and in this regard it raises fundamental unanswered questions. To return to the question with which I started, it is hard to see how transformational-generative linguistic theory could be of much interest to lexicography. In the first place, its great insights are focused on the clause or sentence as a unit, not on the word or phrase, which are the ranks of unit that are of interest to lexicography. In the second place, the logical relations which permeate traditional lexicography are inherited from the traditional logic of Leibniz and others, going back to the medieval schoolmen and beyond. I have been arguing that these traditional notions need to be replaced, but by the analytical concepts inherited from modern prototype theory and theories of social convention, not by the logical relations accounted for in transformational-generative theory.

Observation of the psychological realities suggests that human beings have simultaneously both digital and analogical reasoning powers. We all have the ability to calculate (though some people are better at it than others), and we all have the ability to draw analogies (though some analogies are more imaginative and informative or should we say far-fetched than others). Somehow these two abilities coexist in a single human skull, and both are invoked as we interact with other members of our own species.

When all is said and done, human beings are social animals, and language is the instrument of their sociability. A satisfactory theory of language performance, therefore, must be pursued as a subset of social theory, explaining the preferences of linguistic units in terms of the forces governing collective behaviour and the vagaries of fashion, rather than in terms of logical structures.

A language consists of sets of units and structures, structures which are activated by people behaving linguistically. Units at the word level more strictly speaking, at the lexeme level are fired up and pressed into service by speakers or writers seeking to make meanings, Treating meanings as events rather than objects yields a more satisfactory explanation of the dynamic nature of language than treating them as objects.