

# Semantics and Linguistic Explanation

**Nirmalangshu Mukherji**

I will assume that the study of language essentially falls under what Chomsky (1986) has called “Plato’s problem”: how do organisms form such rich cognitive structures from so little exposure to the environment? I take this to be the original and fairly classical motivation for cognitive science although, arguably, not everyone who currently works in cognitive science shares the motivation.<sup>1</sup> Plato’s problem arises from what have come to be known as the “poverty-of-stimulus arguments” which show that there is not enough information in the environment for the rich systems constructed by organisms (Chomsky 1957; Piatelli-Palmarini 1980; Wexler 1991; Crain and Pietrosky 2001, etc.). As Chomsky puts it,

“We can check the experience available; we can look at it and see what it is. It’s immediately obvious that it’s just much too limited and fragmentary to do anything more than shape an already existing common form in limited fashions” (Chomsky et al 2000a:6).

Accordingly, the specific goal of Bilingualistics is to explain how sound-meaning correlations are established in languages within the constraints imposed by the poverty of the stimulus. The Government-Binding theory (Chomsky 1981) sought to explain those correlations from universal principles alone. The Minimalist Program went several steps ahead to show that such correlations are established from universal principles just as legibility conditions at the interfaces are optimally met:

“If human language is perfect in a very strong sense, then the sound-meaning relations will follow from an optimal solution to the legibility conditions.” (Chomsky et al 2000a: 20)

But the descriptive goal remains the same: establish sound-meaning correlations.

## 1. “Incompleteness” of Grammar

The problem is that, in some widely accepted sense of “meaning”, Bilingualistics fails to establish sound-meaning correlations. Grammar will not distinguish between (the wide differences in) the meanings of (1) and (2).

(1) John decided to attend college.

(2) Bill tried to attend church.

(1) and (2) have nearly-identical LF representations although they differ widely in sound. So if LF is to be viewed as restricted to the grammatical sense of “meaning”, then *two* PF representations have been correlated with *one* LF representation. We will see in a moment why this happens: any linguist will spot it immediately.

---

<sup>1</sup>See Leiber (1991) for a similar conception of cognitive science. See Seidenberg (1995) for a somewhat different conception.

Let me be very clear about exactly what the problem is. The task for Biolinguistics is not to establish sound-meaning correlations *per se*, but to establish them in accordance with the native speaker's intuition which is based on the knowledge of language internalized by her. So, it is not the *properties* of English sentences which are the focus of inquiry; rather, it is the *knowledge* of English under usual abstractions. If the native speaker takes a single sound to have different meanings (consider ambiguity, as in, *flying planes*), the theory explains that by postulating mechanisms for displacement. If the native speaker takes two sounds to have the same meaning (synonymy, as in active-passive pairs), the theory explains that as well, again by postulating mechanisms for displacement. By parity, if the native speaker takes two sounds to have two different meanings, as in (1) and (2), the theory ought to explain that too. This, the theory fails to do.

The basic idea is that a linguistic token<sup>2</sup>, in its full complexity, is an event in the external world which the mind perceives. As the percept is presented to the mind, various capacities of the grammatical mind are activated to analyse the percept. Each capacity analyses specific aspects of the percept with an array of types such that, when all the capacities to which the percept supplies information have finished analysing the percept, every property of the percept has been examined. The (net) resulting representation may be thought of as the representation of the percept/token. As more types are invoked by each capacity, finer partitioning of possible representations, compatible with that capacity, takes place. A token is reached when the output-sets of each capacity have interacted, in some order, to factor out uniquely all other representations except that of the token, much like the intersections in a Venn diagram. Before a token is reached, every type-theoretic characterization will enumerate a large, perhaps infinite, class of representations.

Consider the category of verbs which take clausal complements. Among these, some form a separate class since they have the complex property of deleting the S-bar node in their complement structure (in the G-B formulation). This class includes verbs such as *believe, know, seem, appear*. An even finer classification is possible among these four verbs. We can have *it seems / appears that ...* but we can not have *it believes / knows that ...* with pleonastic *it*. Among verbs which take infinitival clause, some are "control" verbs (*decide, try, persuade*). Control verbs must have thematic subjects, among other interesting properties. So, grammatical theory does take us tantalisingly close to a unique characterization of individual lexical items. Yet, we are still left with at least two verbs (*believe, know*), (*seem, appear*), and (*decide, try*) in each subclass. Each category is certainly "open-ended" in that new verbal labels are constantly introduced as speakers of English experience new processes and events; some of these labels might enlarge the subclasses just listed.

Thus, even though theoretical classification cuts it pretty fine, it seems that an exhaustive application of grammatical theory inevitably leaves at least a pair, the members of which can not be distinguished any further within the theory. Thus, given a pair of grammatical representations (PF, LF) for the strings *John tried to attend college* and *John decided to attend college*, no further partitioning of representations is possible in grammatical theory. These two sentences obviously differ in meaning; hence, no unique sound-meaning connection has been established.

---

<sup>2</sup> See Bromberger (1989) for type/token distinction in linguistics.

The problem is compounded for items in the [+N] category: (*John, Bill*), (*college, church*), etc. This much is pretty obvious (Marconi 1996:Chapter 1). What is not so obvious is the lesson to be drawn from this.

If grammatical theory indeed had the ability to uniquely identify each lexical item non-phonetically, then the demand for further *linguistic* explanations would have looked somewhat questionable. In that case, grammatical theory would have analyzed the linguistic material - phonetic and non-phonetic - contained in a token in its entirety; so the rest of the material, whatever their characteristic, would have fallen out of the agenda of a general theory of language. However, this option is not available as far as current understanding of language-theory is concerned: we want a theory of language to uniquely correlate a sound with its meaning. Grammatical theory apparently fails to do that.

I said *apparently* because the conclusion depends entirely on how we construe “meanings”. If we had a concept of meaning under which (1) and (2) are viewed as synonymous, then unique correlation would have been achieved in grammar itself. I return to this option in a moment. Perhaps a more radical alternative will be to find some natural way to give up the requirement of unique correlation itself. A still more radical idea will be to give up the very requirement of correlation as long as a solution to “Plato’s problem” has been found. After all, a “final” account of utterance-tokens must deal with many physiological factors that relate to production and hearing of sound. Yet we do not expect a linguistic theory to include a subtheory of ear- and throat-mechanisms. So why not drop the axe even earlier at the natural joint where grammar ends and other things begin? Essentially, this last option will mean abstracting away from the allegedly “linguistic” nature of the system. Currently, we do not have such concepts of meaning. Hence, to that extent, it is reasonable to view grammatical theory as an incomplete theory of language.<sup>3</sup>

This problem does not arise in physics. In principle, if two things differ in the space-time framework, then, *ceteris paribus*, the magnitudes of the forces acting on the things will differ as well. Thus, there is a complete *physical* characterization of anything insofar as it has physical properties at all. No doubt, physics can not fully characterize all the properties of those objects which happen to have physical properties as well. For example, physics can not furnish a complete account of what makes something a tree or an elephant. The additional properties that are required to distinguish trees from elephants thus belong to other disciplines; biology, in this case. Therefore, unless one is a “reductionist” in the sense that one believes that all the properties of anything at all must have a physicalist account, no clear sense of incompleteness attaches to physics. A similar picture attaches, via genetic theory, to the individual objects of biology

To get some hold on the issue just raised, consider again the option mentioned above. I suggested that the problem of incompleteness of grammar disappears *if* we entertain a concept of meaning

---

<sup>3</sup>Just to keep the record straight, there are other notions of incompleteness which might render even physics incomplete. For example, Albert Einstein thought that quantum physics is incomplete in a certain sense. Such notions are not at issue here. The concept of completeness of physics sketched here also differs from that of, say, David Papineau’s (Papineau 1993; also Spurrett and Papineau 1999). As far as I can see, Papineau thinks of physics as complete because for him application of physics is essentially domain-free. This enables him to formulate various theses of ‘supervenience’ for mental phenomena with physics as the base. I am suggesting the completeness of physics in *its* domain.

such that (1) and (2) are viewed as synonymous. I am surprised to learn that a similar view has sometimes been officially aired. Jackendoff cites Grimshaw (1993) as follows, [highlight needed?]

“Linguistically speaking, pairs like [*break* and *shatter*] are synonyms, because they have the same structure. The differences between them are not visible to the language.” (Jackendoff 2002: 338)

The option under discussion proposes that Grimshaw’s idea be extended from her “thin” cases to “thick” cases such as *decide* and *try*, *John* and *Bill*, *church* and *college*, etc. In effect, we view the pair (1)-(2) on a par with, say, a pair of active-passive sentences.

The trouble is that our current conception of meaning is such that the synonymy of a pair of active-passive sentences forms a crucial data for linguistic theory. But then, by parity of conception, the pair (1)-(2) cannot be viewed as synonymous at the same time. In fact, if we are to give an account of the sameness of meaning for active-passive pairs, then giving an account of the *difference* in meaning between (1) and (2) becomes part of the agenda for linguistic theory, as noted. Until a satisfactory account is reached, the theory remains incomplete. Such is the pressure of the current largely commonsensical concept of “meaning”. This leaves us with the following options, in order of immediate concern, as far as I can see.

1. We continue with the current thick concept of meaning, attempt to disentangle its parts and attach accounts of these parts, arranged in some order, to achieve growing completeness.

2. Failing (1), we try to dissociate the scope of grammatical theory from the current putative scope of (broad) language-theory. Whether we continue to call grammatical theory “linguistic theory” becomes a verbal issue. We think of grammatical theory as defining, not unlike physical theory, its own domain internally. Some initial data, such as active-passive pairs, no doubt *triggered* the search for this theory. But once a certain internally-coherent model has been reached, we simply learn to ignore some of the things which are currently thought to be crucial features of the data. This is a routine practice even in linguistic research anyway (Larson and Segal 1995:8-9), not to speak of the more advanced sciences. I am suggesting that we push this practice to its logical end. Clearly, the major task at that point is to come up with some conception of the domain so defined.<sup>4</sup>

3. Failing (2), we conclude that grammatical theory is indeed an inherently incomplete theory and we try to find some philosophical justification for this conclusion. I think a number of philosophers profess this option without really considering (2) first. Notice that we could settle on this option immediately after the failure of (1). That means we could have taken the current conception of language-theory for granted to conclude that the theory cannot be realized. I believe this is how the philosophical/sceptical literature usually works: despite its apparently radical stance, the sceptical literature is essentially conservative.

---

<sup>4</sup> See Mukherji (2002) for a discussion of the domain of language-theory from this direction; see Mukherji (2003a) for an attempt to enlarge the domain from a very different, non-linguistic, direction. Also, Mukherji (2003b) for the effect of this enlargement on a general theory of interpretation.

In the tradition as a whole then, we really have two options: either pursue (1), or, give up. Since giving up is not a cherished option in the non-sceptical tradition, tremendous effort is made to pursue (1). In any case, these remarks carry some significance only when option (1) fails. We haven't seen that happen yet.

## 2. Semantic Data

The demand for non-grammatical (post-LF) linguistic explanations is forced on us by the alleged incompleteness of grammatical theory. Once the demand is seen, it is natural to take the pair of strings (1)-(2) a bit closer to the respective tokens by distinguishing between the specific meanings of, say, *deciding* and *trying*. Popular examples include (subtle) differences between *murder*, *kill* and *assassinate*, and, between *lobbing*, *tossing* and *throwing*. Each of *bank*, *kite*, and *bat* has at least two readings such that the interpretation of a token will vary according to the reading chosen; we do not expect grammatical theory to distinguish between these interpretations. The reading of the indexical element *I* varies as the speaker varies; again, grammatical theory leaves the token-interpretation undecided. A vast range of data thus opens up.

Therefore, a necessary, but certainly not sufficient, condition for generating a representation of a token is to invoke enough non-grammatical types to capture specific meanings of words. A natural first step in that direction is to attach *selectional features* to lexical items to account for lexical ambiguity. In Chomsky (1965:85), lexical items belonging to the category [N] were assumed to have features such as  $\pm$ Common,  $\pm$ Count,  $\pm$ Animate,  $\pm$ Human etc., arranged in the order just stated. These features were then used to display more fully some of the subcategorization properties of verbs; for example, subcategorization frames were now supposed to mention types such as [[+Abstract]-Subject], [[-Animate]-Object], and the like. This generates an elaborate system of agreements in which subcategorization frames of verbs will be checked to see whether they match the selectional features of arguments. Resources such as these may now be used, in fairly obvious ways, to throw out 'deviant' strings such as *colourless green ideas sleep furiously*, *golf plays John* and *misery loves company*, while admitting strings such as *revolutionary new ideas appear frequently*, *John plays golf* and *John loves company* (Chomsky 1965:149).

An obvious problem with these ideas is that one does not know where to stop. For example, Jackendoff (1990:51-52) points out that the verb *drink* takes (names of) liquids as internal arguments. Should we, therefore, include [+Liquid] in the subcategorization frame of *drink*, and as a selectional feature of whatever happens to be in the internal argument position? In Bengali, the verb *khaawaa* applies to any of solids, liquids and gases, among various other things: *bhaat* 'rice' *khaawaa*, *jol* 'water' *khaawaa*, *haawaa* 'air' *khaawaa*, *cigarette khaawaa*, *gaal* 'abuses' *khaawaa*, *chumu* 'kisses' *khaawaa*, and so on. How is the subcategorization frame of *khaawaa* structured?

Selectional features, which attach to the category [N], do not immediately explain other varieties of meaning-relationships which give rise to synonymous expressions between pairs of verbs. Additional mechanisms are needed to account for these facts. Consider the following pairs of sentences (Chomsky 1965:62).

- (3) John strikes me as pompous / I regard John as pompous
- (4) John bought the book from Bill / Bill sold the book to John
- (5) John struck Bill / Bill received a blow from John

In each case, roughly, the relation that the first verb establishes between two NPs (*John, I*) and (*John, Bill*) respectively is maintained by the second verb, although the relative positions of NPs vary. It is natural to express these constancies in terms of thematic roles of NPs. Thus, in (5), the two sentences are related by the fact that, in each case, John is the agent and Bill the recipient/patient. Similar is the case in (3) and (4).

A somewhat different “meaning” relationship – not of synonymy, but of entailment – obtains between the verbs *persuade* and *intend*. Chomsky (1991:35) suggested that *John persuaded Bill to attend college* implies, apparently independently of world-knowledge, that Bill decided or intended to attend college. Further, there is some sort of presuppositional link between *John is proud of what Bill did* and *John has some responsibility for Bill’s actions*, that needs to be explained in terms of the universal concepts of PRIDE and RESPONSIBILITY (Chomsky 1972:60). A variety of concepts play crucial roles throughout these examples, although their specific roles vary according to the relationships that are established between words. These form a small sample of the data that a putative lexical semantics needs to give an account of.

Chomsky did pay some attention to these facts and to the issues that arise from them in his *Aspects of the Theory of Syntax* (1965:4.1), as noted. There is some discussion in this book regarding the form a semantic theory might take to account for the facts listed above. In any case, the general conclusion reached in this book was: [is the following highlight required?]

“(s)electional rules play a rather marginal role in the grammar.” (Chomsky 1965:153)

Thus, “(o)ne might propose ... that selectional rules be dropped from the syntax and that their function be taken over by the semantic component”. Since Chomsky has repeatedly urged the need to develop “a more general theory of language” that is concerned with syntax, semantics, and pragmatics and their points of connection (Chomsky 1957:108; 1999), it is interesting to ask what, according to Chomsky, are the prospects for attaining such a general theory.

Chomsky’s writings on language, both technical and informal, have been torrential in the decades that followed *Aspects*. Yet it will be surprising if his constructive remarks on the issues just raised exceed a couple of dozen pages in all.<sup>5</sup> In these pages, he basically repeats the examples cited above to suggest that a universal theory of concepts is urgently needed, without suggesting how this theory is supposed to get off the ground.

---

<sup>5</sup>I am ignoring his (often lengthy) critical remarks on theories of meaning proposed by others, especially by philosophers (Chomsky 1975; 1980; 1986; 1993; 2000b). Recently, he has even questioned the validity of such deeply-entrenched notions as “intentionality” and “representations” which dominate much research in this area (Chomsky 1994; 1995b; 1999; 2000a; 2000b; 2003; forthcoming).

Even though the scope of grammatical theory has been enlarged since *Aspects* to include some aspects of semantics (LF), selectional rules *still* do not play any role in this theory. It is routinely said (Chomsky 1993; 1995a) that lexical items, say *book*, carry semantic features, such as ARTIFACT, along with the usual phonological and grammatical features. However, *nothing* happens to these features after lexical insertion except that they are simply passed on to the “external systems.” In line with the current conception that language interfaces with the conceptual-intentional systems, it is reasonable to conclude then that problems of lexical semantics will be given a general account of, if at all, in a *non-linguistic* theory of concepts. Chomsky’s prolonged silence on this issue could be interpreted as his basic reservation about this enterprise.

### 3. Uncertain Intuitions

Given the limited scope of a paper, I will restrict the discussion to just one of the sides of this issue, namely, whether the sort of data lexical semanticists typically cite to explore the semantic intuitions of native speakers is helpful for reaching an explanatory theory of those intuitions. More specifically, the query is whether such data have the desired form to allow progressive theoretical abstractions once we have described the intuitions informally. I will basically focus on lexical semantics of nouns although I will touch upon the semantics of verbal frames and verbal content.

As a starter, it seems that the data cited above is not as salient as the typical data for grammatical theory. The core data for grammatical theory carry a sense of immediacy and irrefutability. Confronted with the paradigmatic cases of unacceptable strings, it is hard to see how the cases may be “saved”. \**John appeared to the boys to like each other* is flat wrong and tinkering with the meanings of *appear* or *each other* does not help; in fact, such tinkering is not even attempted by the native speaker since, by the time she reaches *each other*, all possible interpretations have collapsed.

More interestingly, even when judgements of acceptability are relatively uncertain, the uncertainty, typically, can neither be removed nor enhanced on reflection. Thus, in an attempt to extend their coverage of data, linguists depend not only on sharp judgements of acceptability, but also on uncertain judgements, where native judgements are uncertain, for example, with respect to whether a given string is okay. The data is listed in some order of increasing uncertainty, and an attempt is made to explain the uncertainties according to the order in which they arise. Consider the following (Chomsky 1986:76-7).

- (6) \*the man to whom I wonder what to give
- (7) \*the man whom I wonder what to give to
- (8) \*the man to whom I wonder what he gave
- (9) \*the man whom I wonder what he gave to

Although each of these sentences is marked as unacceptable, it is clear that they are not all unacceptable to the same degree. Thus, (6) is perhaps the most acceptable, (9) the most hopeless, and the rest fall somewhere in between. We will expect a grammatical theory to order these sentences as they occur because, although our judgements are uncertain, the degree to which a judgement is uncertain, for most speakers of a language, is largely invariant and is not likely to change with further thought.

Semantic judgements, in contrast, are typically open to further thought, even if such thoughts might, on occasion, confirm our initial semantic intuitions. As noted, “deviant” strings illustrate this point directly. In class, I need to shake my head rather vigorously when students attempt to attach coherent interpretations to the string *colourless green ideas sleep furiously*. The point is, I am never allowed to just mention the string, and proceed. We can throw this string out by invoking selection restrictions, as noted. But we may as well get it in by *relaxing* some of them; that’s what the students want. As for *golf plays John*, John’s addiction to golf might lead to the point where it’s golf which takes over.<sup>6</sup> *Misery loves company* sounds fine and is frequently in use since misery is infectious - it spreads. I am obviously stretching things here, but how could I allow myself to do so if the relations are supposed to be “wired-in”?

Consider the relation between *persuade* and *intend*, Chomsky’s favourite example. If I have persuaded X to do Y, does it always follow that X intends to do Y? Is it meaningless to say, *I have been persuading John to attend college, but so far he hasn’t agreed to?* Persuasion seems to be an act that is stretched over time; hence the continuous form. John’s intentions, on the other hand, are not actions, but states of John, which he either has or does not have. It is not obvious that I have failed to act at all just because John failed to attain the relevant state. It will be said that my action was really not of persuading John, but of *trying* to persuade John; the try goes through even if my persuasion fails. So the suggested entailment does hold, confirming Chomsky’s intuitions. But it took some persuasion to make a fairly competent user of English to agree.

Even when we confirm our initial intuitions after analysis in the case just cited, another worry begins to grow. In the reading of *persuade* just confirmed, *pursuade* essentially amounts to a causative, say, *make-intend*. Now we know the source of the initial intuition: the relationship between *persuade* and *intend* is trivially analytic. X and make-X are obviously related, the latter being a function of the former.

This is not to suggest that the observed relation between *persuade* and *intend* is without theoretical interest. As Chomsky has shown recently in a (rare) extensive discussion of the issue (Chomsky 2003), the lexical items under consideration occur in fairly restrictive syntactic contexts. For example, *persuade* typically occurs in the syntactic frame,

- (10)            Nominal -- V -- Nominal -- [Infinitival Clause]  
                  |            |            |            |  
(11)            John persuaded Mary            to buy the book.

<sup>6</sup>See Fodor and Lepore (1991:146-47) for remarks on *the answer calculated John*.

Chomsky points out that (11) entails something about Mary's intentions, namely that she intends to buy the book, but it entails nothing about John's intentions. Interestingly, as with a host of other verbal items, the lexical item *expect* also appears in the syntactic frame (10), without entailing anything about Mary's intentions.

(12) John expected Mary to buy the book.

The parallel breaks down further as *persuade* can not appear in (13), but *expect* can appear in (14):

(13) \*John persuaded there to be a hurricane tomorrow

(14) John expected there to be a hurricane tomorrow.

These facts certainly show that specific aspects of meanings of verbs have syntactic consequences, that is, these aspects enter into computation. Chomsky calls these aspects of meaning "I-meanings." In other words, these aspects of meaning now fall under grammatical theory. So far, these observations are just listed as properties of lexical items; for example, *persuade* is listed as an object-control verb. Hopefully, a more principled and abstract grammatical description will be found in future as the understanding of the structure of the lexicon improves.

The point can be extended to cover examples (3)-(5) above. Thus, *sell* is another causative amounting to *make-buy*, *strike* amounts to *make-receive-blow*, and so on. The burden of explanation thus shifts to the individual concepts INTENDING, BUYING, BLOW-RECEIVING, and the like, plus some general cognitive account of formation of causatives.<sup>7</sup> Given Plato's Problem, it was always obvious that we need an explanatory account of the system of concepts. The data cited above was expected to supply some point of entry to this complex system. The worry is that they don't. They simply state what we knew before; namely, that an account of individual concepts is needed.

The alleged relationship between *proud of Bill* and *responsible for Bill's actions*, another of Chomsky's favourite examples, is even more problematic. I am proud of various sportsmen who have brought laurels for the country. It is not clear at all that I am responsible in any way for what Sachin Tendulkar or Viswanathan Anand achieved. The suggested relationship is possibly restricted to small units, such as a family or a college-class, where each member is likely to be responsible in some ways for the actions of others. Even then it is not clear that an estranged father can not be proud of what his daughter did.

Just as grammatically unacceptable strings are marked with a '\*', it is a standard practice in the literature to mark semantically anomalous strings with a '?'. I will now present a list of such

---

<sup>7</sup>See Levin and Rappaport Hovav (1994) for some interesting data and attempts at theory.

anomalous strings picked virtually at random from an influential book on semantics (Jackendoff 1990) and leave it to the reader to see if they find them to be *clearly* bad; I don't.

- (15) ?For hours, Bill ran into the house.
- (16) ?The snake climbed down the tree.
- (17) ?Bill showed Mary herself (in the mirror).
- (18) ?Some people were streaming into the room.
- (19) ?What Bill did to the wall was smear paint on it.
- (20) ?He paid Bill some money for the book.

The qualification “clearly” is important. We are discussing the nature and quality of semantic data *vis-à-vis* grammatical data. Native speakers may well find (10)-(15) to be bad *after* reflection. That will not defeat the point. If we disagree about the alleged anomalous character of the data, or if we are unsure that the anomaly is a stable feature of the data, then it will be unreasonable to take the next step to theory.

I am not suggesting that strings can not be non-grammatically anomalous. Asking of certain questions is moot: for example, *Can you hear me?*. We are likely to be puzzled by strings such as *Bill admires the Pope's wife*, *John's subordinate is John's boss*, *The mat sat on the cat*, and the like. Any popular piece on newspaper howlers, political speeches, student scripts etc. contains dozens of examples of this sort. However, examples from language-use, unless they are properly controlled, admit of any number of factors not all of which pertain directly to a theory of language; clearly, the anomaly of *Bill admires the Pope's wife* has nothing to do with English (or Italian, for that matter). When such popular, and often hilarious, examples are closely scrutinized, it turns out that they involve any of grammatical failure, performance failure, pragmatic failure and insufficient world-knowledge, among many other things (Pinker 1995). I *am* suggesting, however, that when it comes to data which is supposed to be theoretically interesting, the quality of the data leaves much to be desired. So the chances are that troubles with the data may lead up to troubles with the theories they spawn.

Let me review the situation once again. I attempted to raise two points simultaneously: an explicit point and an implicit one. The explicit point was to cast some preliminary doubt as to whether the alleged data for lexical semantics is salient enough for us to feel interested in the search for deep theory. As we are looking for a theory in this domain from LF-up, as it were, we will expect to find some initial data that has some promise of throwing further light on the human genetic make-up: that is the only concept of an explanatory (cognitive) theory under discussion here. I am personally not convinced that the popular examples cited in this regard hold the promise.

The implicit point was to show that what is called “semantics” in this area essentially reduces to giving at least a systematic account of the organization of a vast network of concepts used by

humans. In each case, the data demanded that we know more about individual concepts and their relations to one another, as they attach to words, whether a noun or a verb. Apparently then, the two points are opposed to one another: the second urges a research program which the first casts doubt upon. Yet they are not really in opposition. *If* we could pursue the program to some degree of success, then it is quite possible that new data will show up, and doubts about the quality of initial data will be gradually set aside.

Perhaps we can already see that the resulting enterprise will have a very different flavour than that of grammatical research, in which a general research program took off from *very specific* (and incontrovertible) data of linguistic behaviour. The same, it could be argued, is the case with physics. It is hard to imagine that physics could have reached the depths it soon achieved if some ancient scientists were to proclaim, “Let us try to understand the physical universe.” It wouldn’t have been clear what is specifically there to understand. Even though the general query might well have been entertained in classical philosophy, what led to physical theory were some sharp and deep facts which demanded sustained explanation: day and night, tides, eclipses, motion of pendulums and projectiles, angle of shadow cast by Sun, and so on.

Such facts abound in grammatical theory at various levels of generality: the fact of rapid acquisition of language, the ambiguity of *flying planes can be dangerous*, the missing lexical object in *John is eager to please*. I am not convinced that *this* sort of motivation exists for lexical semantics; all we have is some general motivation, not unlike the proclamation cited above. For example, a recent work on lexical semantics begins with the following project:

“it would be at least useful to investigate our semantic competence, that is, to wonder what kind of knowledge and abilities *we* possess that make it possible for us to understand language” (Marconi 1996:1).

With these sober expectations in mind, let us return to the second (implicit) point mentioned above.

#### 4. Semantic Decomposition

The task then is to come up with some account of how individual concepts are organized to lend, so to speak, meanings to individual words. Following poverty-of-stimulus arguments, it is natural to think that most concepts are somehow built out of primitive concepts which must be just a handful. In this picture, the meanings of most lexical items, perhaps all, will be captured “decompositionally”: that is, the total meaning of an item will be broken down until the conceptual primitives are reached. Once the system is made available to the child, some of the nodes of the system (that is, the individual concepts) will be associated to the sounds that the child hears: each sound-node association will count as a word. This is a fairly obvious and standard assumption:

“(T)wo things are involved in knowing the meaning of a word – having the concept and mapping the concept onto the right form. This is the sense of ‘knowing the meaning of a word’ implicit in most discussions of language development, both scientific and informal” (Bloom 2000:17).

The child has to learn these associations individually in any case. I said the project looks “natural,” although it is far from clear how it is going to be executed.

Since the aim of the present discussion is to examine the prospects for a unified theory of language, and since we have already included Biolinguistics as the first step towards that end, only those theories count which have some chance of getting attached to the LF level of linguistic representation. After all, we are trying to find out if the alleged incompleteness of Biolinguistics can be overcome in stages; so frameworks that reject this theory in the first place naturally fall out of consideration, assuming that their rejection of Biolinguistics is in fact reflected in the structure of the semantic theories they propose. To take just one example out of many, Peter Gärdenfors offers the thesis that,

“semantics is primary to syntax and partly determines it; this thesis is anathema to the Chomskyan tradition within linguistics” (Gärdenfors 1996:164-65).

As a result, we will not discuss his “Conceptual Spaces” Semantics. Very similar remarks apply to Anna Wierzbicka’s work (1996). In general, I am aware that the criteria proposed here leave a massive body of work on semantics untouched, for example, the entire body of research that falls under “Cognitive Linguistics”. I am assuming that they are not relevant for the issues at hand.

The classic work of Jerrold Katz (Katz and Fodor 1964; Katz 1972) is the most explicit articulation of the goal of semantic decomposition. Katz proposed his theory several decades ago, and everyone in the field knows that it has fatal problems. Yet, it is unclear if these problems pertain to the specific formulations in Katz’s theory, or whether the stated goal of lexical semantics itself is fundamentally flawed. In order to focus on this issue, I will ignore several aspects of Katz’s work which he thought to be central to his theory, and which have been strongly criticized in the subsequent literature. First, I will ignore the suggestion that

“the meaning of a word can be exhaustively decomposed into a finite set of conditions that are collectively necessary and sufficient to determine the reference of a word” (Jackendoff 1983:112).

In other words, with respect to a proposed decomposition, we will not ask if the decomposition is exhaustive; we will just ask if it tallies with the intuitions that are generally associated with the use of the word.

Second, I will ignore Katz’s distinction between semantic markers and “distinguishers” (Katz and Fodor 1964), or between markers and encyclopedic knowledge (Katz 1972). Following the philosophical tradition initiated by Gottlob Frege (1892), Katz made the plausible assumption that the meaning of a word *per se* needs to be distinguished from what (else) we know about the things to which the word applies. For example, ARTIFACT seems to be a part of the meaning of *computer*; the knowledge of its make or its colour does not seem to belong to the meaning. As a result, giving up the distinction between markers and world knowledge means that everything that is truly correlated with the use of a lexical item will now count as part of the word-meaning itself. If it is true that politicians are commonly corrupt, we might find *honest politician* to be anomalous, and we proceed to explain the anomaly by introducing DISHONEST in the path of *politician*. Continuing, we distinguish between *chair* and *stool* by invoking  $\pm$ HAS A BACK, and between *duck* and *goose* by  $\pm$ LONG NECK (Jackendoff 1992:44). So there is much in favour of the distinction. In practice, however, it is extremely difficult to keep semantic and encyclopedic knowledge separate.

Thus, we will simply pick a noun and ask what concepts can be listed, in some order, to specify its meaning, that is, we will treat *all* concepts that enter into the decomposition of a complex concept as semantic markers. To emphasize, we will not ask if (some of) these concepts are in fact primitive, nor will we ask how they distribute between semantic and encyclopedic knowledge. The idea is that once we have a sufficient body of decompositions for a variety of nouns, some patterns, uniformities etc. are likely to emerge to lead us towards a finite set of general concepts. Following this intuition, the only (very natural) constraint we will impose on the theory is that of *superordination*: a more general concept must dominate a less general concept in a semantic tree. This will ensure that the *most* general concepts dominate the rest.

Furthermore, it is advisable that we begin with nouns whose decompositional character is something of common knowledge - nouns which have generally agreed definitions for some of their central uses. It is well-known that explicit decomposition of meaning is a rare phenomenon in any case. It seems to work, if at all, for fairly restricted classes of lexical items (Fodor et al. 1980): these include jargon vocabularies (*ketch, highball*), terms in axiomatized systems (*triangle*), and kinship vocabularies (*grandmother, bachelor*). As Jerry Fodor (1998:70-2), following Putnam (1983), observes, it is difficult to deny the “conservative” intuition that *bachelor* and *unmarried* have an “intrinsic conceptual connection” such that *bachelors are unmarried* is “boringly analytic.” We will focus on *bachelor* to judge the theory on its strongest ground.

Katz thought that a semantic theory ought to describe the semantic competence of native speakers by specifying rules and definitions of semantic properties and relations, stated in terms of the output, that is, the relevant cognitive structures attained by the speaker. Furthermore, continuing the analogy with syntactic theory, he thought that a semantic theory is required to provide a scheme of representation for meanings, which he took to be “a universal theory of concepts” (Katz 1972:32). This part of the program thus covers universal semantics on analogy with Universal Grammar, UG.

As with concepts, we represent a semantic marker in the uppercase. Markers such as HUMAN and COLOR are *not* supposed to be English words. They are supposed to be “theoretical constructs” lexicalized in English as *human* and *color*. In a language different from English, these constructs will be either lexicalized differently, or borrowed from some other language, or not lexicalized at all. In other words, the English word *human* is supposed to invoke, among other things, the conceptual information HUMAN; similarly for other common nouns such as *male, adult, animal*, and the like. Thinking of common nouns as “predicates” that take (a set of) semantic markers as “arguments,” these markers may now be used to decompose the meaning of a complex predicate such as *bachelor*. Following the constraint of superordination, the dictionary entry for *bachelor* will then be listed as a path with markers HUMAN MALE ADULT UNMARRIED, *in that order* of increasing specificity. Despite the high rhetoric borrowed from grammatical research, the theory is beset with serious internal problems from the beginning.

The lexical item *bachelor* typically means an unmarried male, as noted; but it could also mean a baby fur-seal, among other things. It is important that these two readings apply to the same word *bachelor* since it is a part of the native user’s grasp of *bachelor* that, unlike *kite* and *bat*, it is not ambiguous, although it has multiple application. How do we capture this in a single lexical

entry? The problem is that the *differences* between these two readings of *bachelor* begin at the top. Since *bachelor* could cover either people or fur seals, the markerese paths branch right at the top, HUMAN or ANIMAL. So, the semantic decompositions for the two readings will consist of, say, HUMAN MALE ADULT UNMARRIED and ANIMAL MALE YOUNG WITHOUT-A-MATE respectively. Notice that we must use MALE after HUMAN (or, ANIMAL) since HUMAN is a higher category. This means that the marker MALE will have to be listed *twice*: once after HUMAN to terminate in UNMARRIED, and again after ANIMAL to terminate in WITHOUT-A-MATE. The generalization that MALE is common to the two readings of *bachelor* is missed. The result is that we are forced to view *bachelor* as lexicalizing two entirely different conceptual paths with nothing in common. But in fact, apart from MALE, the two readings have a central feature in common – something like WITHOUT-A-MATE; how do we insert this item?

Notice that we can not begin with WITHOUT-A-MATE at the top to solve the problem because this category is subordinate to both HUMAN and ANIMAL: everything that is without a mate is either a human or an animal (or a robot), but not all humans or animals (or robots) are without a mate. If we insert WITHOUT-A-MATE after HUMAN or ANIMAL in separate paths, the commonality will not be represented. Suppose, we simply declare that *bachelor* is indeed hopelessly ambiguous in that it applies either to humans or animals. The philosopher of science, Norwood Russell Hanson once opened a very high-profile conference on meaning of theoretical terms in science with the observation that *electron* means little things that wiggle. That is not much of a semantics, but that is all we have.

However, as Jackendoff (1983:14) following Bolinger points out, such an austere semantics can not be granted. One of the tasks facing a semantic theory is to explain semantic anomalies; for example, we need some distinguishing feature to explain the anomaly of *some green things are blue*. But, it is unclear how we distinguish between *blue* and *green* after we introduce COLOR unless we introduce BLUE and GREEN themselves as markers (Jackendoff 1983:117); to do so, of course, is to confuse between data and its analysis.

Similarly, we need to explain the anomaly of *the bachelor's legitimate daughter*. This can not be done unless the feature UNMARRIED is taken into account. Moreover, we would want a semantic theory to explain why *bachelor* and *spinster* are closer in meaning than, say, *bachelor* and *professor*. Here again the feature UNMARRIED plays the crucial role *before* the path reaches MALE. But we can not list UNMARRIED before MALE since there is no heirarchy: not everything unmarried is a male, not everything male is unmarried. So, UNMARRIED will have to be entered separately in the paths for *bachelor* and *spinster*; thus, missing the desired generalization.

Suppose we give up everything else and settle for UNMARRIED, as Fodor (above) suggested, as the meaning of *bachelor*, since we failed to accomodate this crucial item in tandem with other markers except HUMAN and MALE. As noted, there is also an intuition that the sense in which *bachelor* applies to people has something in common with its application to fur-seals. That intuition has to be given up now since UNMARRIED will not apply to anything outside humans. In this picture, HUMAN immediately dominates the sequence MALE ADULT UNMARRIED where the path terminates.

The composite concept BACHELOR now applies to each unmarried adult human male. A common counterexample, due to George Lakoff (1987), is the Pope; the Pope is unmarried, is he a bachelor? Suppose we introduce the item CAN-MARRY to exclude the Pope. All human males of (marriageable age) are thus subjected to a binary classification, UNMARRIED and MARRIED. Under the first, we have another binary branching into CAN-MARRY and CANNOT-MARRY. So, a bachelor is UNMARRIED CAN-MARRY, the Pope is the other one. But then a muslim (another popular counterexample), who is allowed upto four wives but in fact has two wives, can not be placed in the scheme: he is both married and can marry. If we now allow the second binary branching under MARRIED as well, it generates the contradictory path MARRIED CANNOT-MARRY. To solve the problem, we go back and revise the item CAN-MARRY to CAN-MARRY-AGAIN so that we have a consistent path MARRIED CAN-MARRY-AGAIN which accomodates the muslim. But now, a bachelor is a human male who is unmarried, but who can marry again. But CAN-MARRY-AGAIN applies only to someone who is already married which conflicts with the item UNMARRIED. We can not have it both ways.

Given problems with MARRIED, suppose we go back to the wider WITHOUT-A-MATE; this gives the fur-seal (and the Pope) another chance. Intuitively, this is “really” what *bachelor* means: a bachelor is essentially some sort of a solitary individual, someone without a mate. Marriage is not really part of the concept BACHELOR; the “conservative” intuition that bachelors are unmarried is in fact conservative because we expect every male of a certain age to be with a mate, and in an orthodox setup the only way to find one is to get married.

Suppose all this and consider MATE. Suppose we decide that as long as someone is with a mate, he is not a bachelor. In this reading, something like a heterosexual partnership with some sex-life is assumed for both humans and fur-seals. The assumption can be questioned at least for humans:<sup>8</sup> (a) one can live together with a mate who belongs to the same gender, (b) one can live together with a mate without having a sex-life. With such options in hand, it is difficult to decide who counts as a mate (are R2D2 and 3CPO mates?). If a person decides to spend his life with his unmarried brother, then, other things being equal, both are bachelors, but are they without a mate? We can go on like this. All we are left with is that bachelors are some sort of solitary individuals. But *solitary individual* and *bachelor* are not only not synonymous, they do not even have the same extension.

Most of these problems are well-known in the literature, as noted. Ray Jackendoff for one, and there are many others, certainly acknowledges the problems just described. Regarding the attempt to give an “exhaustive decomposition” of the “necessary and sufficient conditions” governing the meaning of *bachelor*, Jackendoff dismisses the entire project because

“the meaning of *bachelor* is inseparable from the understanding of a complicated social framework in which it is embedded” (Jackendoff 2002:375).

---

<sup>8</sup>Recent evidence suggests that homosexuality is prevalent in non-human species as well such as bonobos, formerly called “pigmy chimpanzees” (de Waal 1989, cited in Leiber 2001:90-1). If we wish to extend the use of *bachelor* from fur-seals to bonobos, the concept WITHOUT-A-MATE will have to be delinked from the assumption of heterosexuality there as well.

But then he suggests immediately that “*someone* has to study these more complex aspects of meaning eventually” (376). This raises the prospect that someone, may be Jackendoff himself, is prepared to enter into an understanding of the “complicated social framework” to explain how the meaning of *bachelor* is decomposed, even if the effort may not immediately lead to some conceptual primitives that we hope to reach eventually.

The fact that the project is never launched by Jackendoff, or by anyone else to my knowledge, deserves examination. In our analysis of *bachelor*, we did not ask for “exhaustive” decomposition, neither did we look for all the necessary and sufficient conditions that govern the meaning of *bachelor*. Also, we did not insist on a sharp demarcation between semantics and world knowledge. We just asked for some non-trivial addition to the grammatical story that tells us how to represent the meaning of *bachelor* in terms of concepts that distinguish bachelors, from, say, spinsters. The basic problem seems to be quite overwhelming, namely, that in order to analyze the meaning of a common noun such as *bachelor* just to find a distinguishing concept, we have to jump into a vast sea of world knowledge fairly blindly; this sea doesn’t seem to part *in accordance with* our semantic intuitions. Let me explain.

What is the data for the semantics of *bachelor*? It is a central part of the data for *bachelor* that *bachelors are unmarried males* is “boringly analytic,” as noted. To that extent, there is no problem in representing this intuition in terms of concepts ADULT, MALE, HUMAN, UNMARRIED in some order or other without inquiring which of these are markers and which belong to world knowledge. It is a legitimate question whether this notational move has any explanatory value or whether it is just representation of data. I am setting such problems aside for now because the issue of the (explanatory) salience of a notational scheme arises only after a coherent scheme has been found. It is unclear if that stage has been reached for *bachelor* because the data for *bachelor* also includes intuitions such as baby fur-seals are bachelors, *bachelor* is closer in meaning to *spinster* than to *professor*, the Pope though unmarried is not a bachelor, bachelors are without a mate, two unmarried adult male shipmates are both bachelors, etc.

The complexity of the organization of these intuitions has little to do with the “complicated social framework” involved in MARRIED without denying that marriage is complicated. For example, we will readily agree that nuns are spinsters, and will be reluctant to seriously entertain the suggestion that they may be married to God. Also, we have no problem in observing that two heterosexual adult humans living together are unmarried. So, the problem is not just that whatever decomposes BACHELOR is itself complicated, the problem is that *our intuitions seem to wander across these complicated decomposers*.

These intuitions concern any of marriage, mates, adulthood, heterosexual activity, loneliness, and so on, on a *piecemeal basis* apparently without any concern for global coherence with respect to the decomposers just listed. I said “apparently” because there must be some coherence somewhere as the fluency and convergence of these intuitions on *each* occasion of use amply testify; yet, these intuitions do not seem to converge with respect to each other. The fluency of these intuitions begins to falter when we question them from the “outside,” as it were, that is, with respect to intuitions displayed on other occasions of use. Obviously, the problem multiplies as the original term is used with progressively more world knowledge.

The preceding diagnosis of the problem might be seen as a step towards a more comprehensive lexical semantics that takes into account many more dimensions of world knowledge. Thus James Pustejovsky (1995) proposes that we enlarge the scope of semantics to include aspects of world knowledge such as the origin, material constitution, layout, function and future course of things. It is hard to see how the notions of origin, material constitution, future course etc. might throw significant light on the semantics of abstract terms such as *bachelor*. But maybe a systematic application of these dimensions allows a more comprehensive semantics for concrete nouns (Moravcsik 1981). It is unclear if the complexity of semantic intuitions described above can be interestingly accommodated even with such rich resources.

Chomsky's analysis of *river* is a dramatic illustration of the suspicion just raised (Chomsky forthcoming). Elaborating on Hobbes' definition that we call something the *same river* if it comes from the same source (=origin), Chomsky points out that if the river changes its course by several miles, we will continue to call it the *same river* as long as it has the same source. But then suppose the course of the river is reversed, that is, it does not have the same source any more and we know about it, even then it would be the same river. Next, we can have a river that is artificially broken into tributaries so that it ends up somewhere else; it will be the same river. Suppose it is filled with waste containing 99% arsenic from some chemical plant upstream thus changing material constitution drastically; it is still the same river. Similarly, if the river dries up completely the material constitution changes, but it is the same river.

In contrast, suppose we make a minuscule quantum-theoretic change so tiny that nobody can even detect it, and the river hardens into a glassy substance. Suppose we sprinkle something on it to add friction, draw a line in the middle, and cars start going up and down. It is the same object at the same place with virtually the same material constitution. But it will not be called a *river*, it will be called a *highway*. To add to Chomsky, suppose further that the river remains dry for many years, the bed fills up with erosions from the banks, and people start cultivation. Suppose this state of affairs continues for several generations, and the name of the river disappears from the local dialect. Then, one day, torrential rains start, the sand and the mud is washed away, and water starts flowing. It will now be called a *flooded field*.

The net result is that even if we enlarge the scope of semantics in many dimensions to include the origin, material constitution, layout, function and future course of things as well as social expectations, conventions and psychological needs, it is unclear how these additional dimensions fit the organization of semantic intuitions. It seems that semantic decomposition basically boils down to making of lists of current uses of a term, as above, to be supplemented in the future as and when new uses are detected. One could add some theoretical flavor to such lists by introducing technical terms such as polysemy, family resemblance, cluster concept, and the like. But all they do is to highlight the fact that we have a list, nothing more.

## 5. Concluding remarks

Ray Jackendoff raises a similar complaint with much influential work in lexical semantics: building "an industry on the endless details of a single word" is not "properly systematic", he says. Thus, he is unhappy with the work of Steven Pinker and Anna Wierzbicka because "the result is all too often a tiring list, impossible for any but the most dedicated reader to assimilate"

(Jackendoff 2002:377). Jackendoff's complaint is that the activity of making "unsystematic" lists is not even a first step for the solution of Plato's Problem.

To that end, Jackendoff's suggestion is that we stick to the original agenda of semantic decomposition, but aim for a "far richer notion of lexical decomposition" (2002:xvi). Thus, he wishes to reformulate the agenda with the explicit recognition that word meanings constitute a "richly structured system" along a variety of dimensions. These dimensions include "conditions that shade away from focal values, conditions that operate in a cluster system, features that cut across semantic fields," and so on. We note that the cited suggestion was made in the context of the complexity of *bachelor* and proper names, among other things (2002:377).

This formulation of semantic complexity is laden with heavy theoretical expectations. What is the "focal value" in the "richly structured" system of the meanings of *bachelor* and *river*? Are UNMARRIED ADULT MALE and SAME SOURCE the focal values from which the meanings of *bachelor* and *river* respectively "shade away"? For *bachelor* at least, we tried our best to formulate paths that begin with UNMARRIED ADULT MALE and then branch off – "shade away" – in different directions to include other concepts required by other intuitions accompanying the same term. The attempt failed because, for one, it clashed with the intuitions that baby fur-seals are bachelors and the Pope is not a bachelor. Does the giving up of most or even all of the semantic properties of a "focal value" count as "shading away"? What is the condition that operates in the "cluster system" mentioned by Jackendoff? If anything, the system seems to be flatly contradictory, unless by "cluster" Jackendoff means just the list of conditions that govern particular uses. What are the "features" involved and what are the "semantic fields" in the meanings of *bachelor* such that the former can be viewed as "cutting across" the latter?

However, there is another way of looking at Jackendoff's suggestion that does seem to offer, perhaps subliminally, a theoretical handle on the issue of complexity. Jackendoff could be suggesting that notions such as focal value and semantic field already form a part of the theoretical vocabulary in the field of lexical semantics. But the domains in which these notions seem to have some theoretical value do not currently include the sort of complexity of nouns under discussion. Chomsky (forthcoming, emphasis added) makes a similar point as follows:

The internal structure of nouns is not studied very much. Lexical semantics, which is a very rich study, has dealt mostly with verbs. The reason is because you can discover things; *you can find the primitive elements that seem to be rearranged in different ways*. When you start looking at the structure of the simplest nouns, their complexity is overwhelming.

From this perspective, the study of common nouns with "richly structured systems" of meanings becomes a research program already implemented in less complex cases.

Thus it is not surprising that most of the work, including Jackendoff's own, on semantic fields and related aspects of lexical meaning is concentrated on verbs and prepositions rather than on nouns. In our terms, we may not yet know what distinguishes *college* from *church* in a systematic study of nouns, but there is some promise of a theoretical framework that will ultimately make a systematic post-grammatical distinction between *decide* and *try*. The hope is that once the semantics of verbs is adequately understood, it might lead to a better understanding

of the semantics of nouns. A variety of theoretical ideas and their applications were treated at length in Jackendoff (1983; 1990), and summarized in Jackendoff (2002).

I have no space here to study the proposals in detail. However, two quick reflections suggest that a high degree of optimism about semantics of verbs may not be warranted. First, as we saw briefly, much of the data displaying the “richly structured system” of verbs pertains to “I-meanings” of verb-frames; in that the research essentially falls under syntax, as Chomsky stresses, throwing little light on the conceptual organization of verbs. My own feeling is that Chomsky’s remark, cited above, points only to this part of the research. As to the study of verbal content, rather than verbal form, verbs and nouns seem to share very similar problems of complexity. Instead of presenting data, I will close the discussion by citing Chomsky’s view on this matter. In a revealing remark, Chomsky in fact groups three verbs and two nouns to raise a single problem:

“There is overwhelming reason to believe that concepts like, say, CLIMB, CHASE, RUN, TREE and BOOK and so on are fundamentally fixed. They have extremely complex properties when you look at them. This is not recognised in traditional lexicography. When you read the huge *Oxford English Dictionary* (the one you read with a magnifying glass), you may think that you are getting the definition of a word but you’re not. All you are getting is a few hints and then your innate knowledge is filling in all the details and you end up knowing what the word means. As soon as you try to spell out what’s taken for granted in the lexicon, you find that these concepts are incredibly complex (Chomsky et al 2000:75).

## References

1. Bloom, P. 2000. *How Children Learn the Meanings of Words*. Cambridge: MIT Press.
2. Bromberger, S. 1989. Types and tokens in linguistics. In A. George (Ed.) *Reflections on Chomsky*. Oxford: Basil Blackwell.
3. Chomsky, N. 1957. *Syntactic Structures*. The Hague: Mouton.
4. Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge: MIT Press.
5. Chomsky, N. 1972. *Language and Mind*. New York: Harcourt Brace Jovanovich (expanded Edition).
6. Chomsky, N. 1975. *Reflections on Language*. New York: Pantheon.
7. Chomsky, N. 1981. *Lectures on Government and Binding*. Dordrecht: Foris.
8. Chomsky, N. 1986. *Knowledge of Language*. New York: Praeger.
9. Chomsky, N. 1991. Linguistics and cognitive science: problems and mysteries. In A. Kasher (Ed.) *The Chomskyan Turn*. Oxford: Basil Blackwell.
10. Chomsky, N. 1993. A minimalist program for linguistic theory. *MIT Working Papers in Linguistics*. Reprinted as Chapter 3 in Chomsky (1995a).
11. Chomsky, N. 1994. Naturalism and dualism in the study of language and mind. *International Journal of Philosophical Studies* 2, 181-209.
12. Chomsky, N. 1995a. *The Minimalist Program*. Cambridge: MIT Press.
13. Chomsky, N. 1995b. Language and nature. *Mind* 104, 1-61.
14. Chomsky, N. 1999. On the nature of pragmatics and related issues: an on-line interview with Noam Chomsky. Conducted by Brigitte Stemmer. *Brain and Language*, 68.3, 393-401
15. Chomsky, N. 2000a. *The Architecture of Language*. The Delhi Lecture, 1996. N. Mukherji, B. N. Patnaik and R. Agnihotri (Eds.). New Delhi: Oxford University Press.
16. Chomsky, N. 2000b. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
17. Chomsky, N. 2003. Reply to Horwich. In L. M. Antony and N. Hornstein (Eds.) *Chomsky and His Critics*. London: Blackwell Publishers.
18. Chomsky, N. Forthcoming. *Language and Naturalistic Inquiry*. Compiled and edited by N. Mukherji.
19. de Waal, F. B. M. 1989. *Peacemaking among Primates*. Cambridge: Cambridge University Press.
20. Crane, S. and P. Pietroski 2001. Nature, nurture and universal grammar. *Linguistics and Philosophy* 24, 139–186.
21. Fodor, J. A. 1998. *Concepts: Where Cognitive Science went Wrong*. Oxford: Clarendon Press.

22. Fodor, J. A., M. Garrett, E. Walker and C. Parkes 1980. Against definitions. *Cognition* **8**.
23. Fodor, J. A. and E. LePore 1994. Why meaning (probably) isn't conceptual role. In S. Stich and T. Warfield (Eds.) *Mental Representation*. London: Basil Blackwell, 142-56.
24. Frege, G. 1892. On sense and reference. In P. Geach and M. Black (Eds.) *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Basil Blackwell, 1966.
25. Gärdenfors, P. 1996. Conceptual spaces as a basis for cognitive semantics. In A. Clarke (Ed.) *Philosophy and Cognitive Science*. The Netherlands: Kluwer Academic Publishers, 159-180.
26. Grimshaw, J. 1993. Semantic structure and semantic content in lexical representation. Ms, Center for Cognitive Science, Rutgers University
27. Jackendoff, R. 1983. *Semantics and Cognition*. Cambridge: MIT Press.
28. Jackendoff, R. 1990. *Semantic Structures*. Cambridge: MIT Press.
29. Jackendoff, R. 1992. *Languages of the Mind: Essays in Mental Representation*. Cambridge: The MIT Press.
30. Jackendoff, R. 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.
31. Katz, J. J. 1972. *Semantic Theory*. New York: Harper and Row.
32. Katz, J. J. and J. Fodor 1964. The structure of a semantic theory. In J. J. Katz and J. A. Fodor (Eds.) *The Structure of Language*. Englewood Cliffs: Prentice Hall.
33. Lakoff, G. 1987. *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.
34. Larson, R. and G. Segal 1995. *Knowledge of Meaning*. Cambridge: MIT Press.
35. Leiber, J. 1991. *An Invitation to Cognitive Science*. Oxford: Basil Blackwell.
36. Leiber, J. 2001. Turing and the fragility and insubstantiality of evolutionary explanations: a puzzle about the unity of Alan Turing's work with some larger implications. *Philosophical Psychology*, **14.1**, 83-94.
37. Levin, B. and Rappaport-Hovav, M. 1994. A preliminary analysis of causative verbs in English. In L. Gleitman and B. Landau (Eds.) *The Acquisition of the Lexicon*. Cambridge: MIT Press.
38. Marconi, D. 1996. *Lexical Competence*. Cambridge: MIT Press.
39. Moravcsik, J.M. 1981. How do words get their meanings?. *Journal of Philosophy*, **78**, 5-24
40. Mukherji, N. 2002. Shifting domains. In M. Chadha and A. Raina (Eds.) *Basic Objects: Case Studies in Theoretical Primitives*. Shimla: Indian Institute of Advanced Study, 91-106.
41. Mukherji, N. 2003a. Is  $C_{HL}$  linguistically specific? *Philosophical Psychology*, **16.2**, 289-308.

42. Mukherji, N. 2003b. Is there a general notion of interpretation? In A. D. Ritivoi (Ed.) *Interpretation and its Objects: Studies in the Philosophy of Michael Krausz*. Amsterdam: Rodopi Publishers, 39-54.
43. Papineau, D. 1993. *Philosophical Naturalism*. Oxford: Basil Blackwell.
44. Piatelli-Palmarini, M. (Ed.) 1980. *Language and Learning: The debate between Jean Piaget and Noam Chomsky*. Cambridge: Harvard University Press.
45. Pinker, S. 1995. *The Language Instinct: How the Mind Creates Language*. New York: HarperCollins.
46. Pustejovsky, J. 1995. *The Generative Lexicon*. Cambridge: MIT Press.
47. Putnam, H. 1983. 'Two Dogmas' revisited. In H. Putnam *Realism and Reason: Philosophical Papers*, Vol. III. Cambridge: Cambridge University Press.
48. Seidenberg, M. S. 1995. Language and connectionism: the developing interface. In J. Mehler, J. and S. Franck (Eds.) *COGNITION on Cognition*. Cambridge: MIT Press, 415-32.
49. Spurrett, D. and Papineau, D. 1999. A note on the completeness of 'physics'. *Analysis* **59** (1), 25-29.
50. Weirzbicka, A. 1996. *Semantics: Primes and Universals*. Oxford: Oxford University Press.
51. Wexler, K. 1991. On the argument from the poverty of the stimulus. In A. Kasher (Ed.) *The Chomskyan Turn*. Oxford: Basil Blackwell.