

RAISING IN SYNTAX, SEMANTICS AND COGNITION

Richard Hudson

Abstract

The 'raising' pattern is the one found in sentences where the subject of a subordinate clause is 'raised' to act as the subject or object of the higher verb. I argue that this pattern can also be identified outside syntax, and specifically both in semantic structures and also in general cognitive structures. Moreover the same pattern underlies some of the most complex syntactic patterns, such as extraction, so these similarities between syntax and general cognition may explain some of the most interesting features of syntax as special cases of more general cognitive phenomena.

1 Introduction and overview

How similar are syntactic and semantic structures? Some linguists have found, or assumed, considerable formal differences - for example, according to Chomsky (1981/82:35) relational categories are basic in semantics, where they are called 'thematic relations' or 'theta-roles', but their counterpart in syntax, 'grammatical relations', are derivative and not part of the basic structure. Again, it has been suggested (e.g. Heringer 1970) that syntax is organised in terms of constituent structure whereas semantic structure is based on dependency; indeed, any theory which combines a function-argument semantic structure with a syntactic structure where heads are not marked as such must, at least implicitly, follow this pattern - one thinks here in particular of GPSG (Gazdar et al 1985).

These theories contrast with other theories which assume formal uniformity across at least these two levels (and possibly across morphology and phonology as well). The extreme case is of course the kind of theory which denies any distinction between syntactic and semantic structures, notably Generative Semantics (though these ideas keep reappearing, as in a recent paper by Pollard, 1989), but a wide range of other theories recognise syntax and semantics as distinct levels of analysis, while at the same time allowing the same range of formal devices for each level - e.g. Stratificational Grammar (Lamb 1966), Functional Grammar (Dik 1978, Siewierska forthcoming), Categorical Grammar (Wood, forthcoming), Word Grammar (Hudson 1984, 1990). The main difference in these theories between syntax and semantics is in their analytical categories. The representational vocabulary for syntax includes categories like 'noun', 'nominative', 'subject', 'clause', and HORSE (the lexical item HORSE), none of which is in the vocabulary for semantics, while that of semantics contains categories like 'person', 'agent', 'action' and 'horse' (the concept 'horse').

The question is, of course, who is right. The optimistic assumption, which I shall make, is that the question is about the facts of language, and does not boil down to a matter of taste, prejudice or other personal attributes of the linguists who hold these different views. I shall try to present a small amount of factual evidence in favour of the second view, that the formal

properties of syntactic and semantic representations are in fact remarkably similar. The evidence will consist in showing that 'raising' structures are found in semantics which are formally just like those familiar from syntax - the kind of thing the Generative Semanticists did so well, but with the difference mentioned above, that I assume an absolutely clear distinction between syntactic and semantic structures. This has the effect of allowing the mapping relations between syntactic and semantic structures to be freer than they would have been had I been obliged to justify a transformational route between them.

I have not yet explained what I mean by 'semantic' and 'syntactic', but I should do so before turning to the third kind of structure, which I call 'cognitive' in the title of the paper. By 'syntactic structure' I mean a structure whose only elements are words, so that all the information contained in the structure is information about these words - about how they are classified and how they are related to one another. In contrast, the elements of a semantic structure are all meanings, and not the words which have those meanings. But we must distinguish semantic structures from other cognitive structures as well, by limiting them just to what is defined in the grammar, so excluding any information derived from context. For example, the grammar tells us that the phrase *the horse* refers to a horse, and that this horse is known to the addressee, but it does not tell us which horse; and so on for all the familiar territory of 'pragmatics'.

The first point of this paper, then, is to argue that syntactic and semantic structures, in the sense just defined, are formally similar. The second point concerns the relations between semantic and cognitive structures - the kinds of structures that we construct by applying contextual information to semantic structures. These include the structures which have a truth value but which only have the minimum of extra contextually-supplied information - for example, the structure for *The horse kicked her* where the particular horse has been identified though no further inferences have yet been drawn (Kempson 1988). But to be realistic it seems unlikely that it will ever be possible to distinguish this level of structure from one where far more inferencing has been done than mere identification of referents - e.g. the level that would allow someone to use *The horse kicked her* as an answer to the question *How did she get that bruise?*. I shall use the term 'cognitive structure' indiscriminately for any level of analysis where contextual and real-world information has been exploited.

Given this very general definition of 'cognitive structure', how similar are semantic and cognitive structures from the point of view of their formal properties? To my knowledge there are no theories which claim that they are different, and I think it would be fair to say that everyone who works in this area assumes that semantic structure has the same formal properties as the structures that can be derived from it - though there is of course considerable disagreement over what these formal properties are (e.g. one thinks of feature structures (Halliday 1985), predicate-argument structures (most semantic and pragmatic theories), semantic dependencies (Schank 1975), frames (Minsky 1974) and mental models (Johnson-Laird 1983)).

This consensus over the formal similarities between semantic and cognitive structures should, however, be seen in the context of our other question, about syntactic and semantic structures. In this case we found some support for the view that syntactic and semantic structures are formally similar, so if everyone agrees that the same is true of syntactic and semantic structures does this mean that those who support the former view also think that syntactic structures are

similar to cognitive structures? I think we cannot draw this inference, though it may in fact be true, for the simple reason that the people who write about the syntax/semantics interface typically do not express opinions about the semantics/pragmatics interface, and vice versa. Nevertheless it is clearly important to take some position on the relations among syntactic, semantic and cognitive structures.

My aim in this paper is to show that syntactic, semantic and cognitive structures are all formally similar, my evidence being the existence in all three structures of a similar structural pattern, the 'raising' pattern.

2 Raising patterns in syntax

In a sentence like *It stopped raining*, there are very good reasons for saying that *it* is the syntactic subject of both the verbs. It must be the subject of *stopped* because any tensed verb has to have a subject and *it* is the only candidate. And it is the subject of *raining* because RAIN requires *it* as its subject. Notice that *it* is required even when RAIN is non-finite, as in **(It) raining all day* was a nuisance. Similarly, in *There seems to be a fly in my soup*, we have to take *there* as subject of *to be*, as well as of *seems*, because its presence has precisely the same explanation as that of *there* in *There is a fly in my soup*. The familiar, and uncontroversial, conclusion is that the subject of STOP or SEEM is also (at some level) the subject of its complement.

I shall follow the transformational tradition in calling this pattern 'raising', because the name is evocative and established, although there are now a number of theories where there is no 'raising' transformation since both the subject relations are shown in a single structure - LFG (Bresnan 1982a), HPSG (Pollard and Sag 1987:20) and Word Grammar (Hudson 1984: 83ff). I shall assume a single-level analysis of this type, in which the grammatical relations are shown directly; the notation and theoretical assumptions are from Word Grammar (where the relations are dependencies), but similar ideas could be expressed in terms of any theory that can be expressed in terms of slots which can share fillers.

According to Word Grammar a verb has a subject slot and a number of complement slots, all of which are dependents of the verb (which is therefore their head). Take a simple example like *Fred kissed Mary*. The structure can be displayed as in (1), where each arrow points from one word towards its dependent.


- (1) <—subject—|—object—>
 Fred kissed Mary

If we now use the same system for showing the raising relation, we get (2).

- (2) <—subject—|—
 <—subject—|—complement—> |
 It stopped raining.

The raising pattern is extremely common in syntax - for example, every


auxiliary verb has a raised subject, not to mention all the well-known 'raising verbs' like SEEM and 'raising adjectives' such as LIKELY. However, what is particularly important about it is the abstract geometry of the pattern:




- i. one element, A, which is related as a dependent to two others, B and C; of which
- ii. B is also the head of C.

This is what I shall call the 'raising pattern', although some of the structures where I shall identify it are not described as 'raising' by all linguists.

One such structure is what used to be called 'subject-to-object raising', as in *I expect it to rain* or (without the complicating *to*) *The wind stopped it raining*. Here everyone agrees that *it* is the subject of *rain(ing)*, but some people are sufficiently impressed by the arguments surveyed in Postal (1974) to believe that *it* is also the object of *expect* or *stopped*. Even in GB it is accepted that *it* is governed by the higher verb (Chomsky 1981/82: 66), which is tantamount to recognising a dependency relation between them. Here the order of elements is different, but the basic geometry is the same:



Once again we have the three elements of a raising pattern, in the correct arrangement: A shared as a dependent by B and C, where B is also the head of C. For convenience I juxtapose this pattern (5) with the one in (3), where the order of elements is different.



From a syntactic point of view, though not semantically, these examples are all indistinguishable from traditional 'Equi', or 'Control', structures, such as *Fred tried to kiss Mary* or *We persuaded Fred to kiss Mary*. Here there is no dispute about the syntactic relation between *Fred* and the first verb, but what is disputed is whether *Fred* is also the subject of the second verb, *kiss*. Given the clear (and undisputed) semantic links between *Fred* and *kiss*, and the existing apparatus for handling raising patterns, the easiest way to analyse these structures is to use the raising pattern, leaving the semantics to distinguish control structures from raising structures in the original sense. For example, the difference between *He seems to like it* and *He tries to like it* is that there is a semantic link between the meanings of *he* and of *tries*, but not between the meanings of *he* and of *seems*.


We have identified the raising pattern in traditional raising and control

structures, each of which may identify a lower subject either with a higher subject or with a higher object. We could go on to show the same pattern in other obviously related structures, such as 'tough-movement' (e.g. *Fred is easy to tease*), but it is more interesting to look further afield, in extraction (Hudson 1988, 1990).

Here we seem to be concerned with a completely different kind of pattern, where an element has simply been moved to an earlier position in the sentence (e.g. *Apples I know you think I ought to like, but ...*). However it is worth noticing, first, that this movement always involves 'raising' in the sense of moving into a higher phrase. Moreover, it always requires the moved element to take its position from the head of that phrase (its verb), a position which is commonly labelled 'Comp' (a term that we can continue with here, though I tend to use a completely different one in my own work, 'visitor'; for present purposes there is no important difference between 'Comp' and 'specifier of Comp').

Comp is always just before the subject of a finite verb (the 'anchor verb'), and there is no obvious reason why it should not be treated as a grammatical relation, like 'subject'. It is true of course that the filler of Comp is not semantically related to the anchor verb - Comp is not a 'theta-position' - but the same is true of the relation between traditional raising verbs and their subject - e.g. in *It stopped raining*, there can be no semantic relation between *It* and *stopped* since (at least arguably) *it* is semantically empty.


Suppose, then, that we assume that there is a grammatical relation, comparable to the subject relation, between the filler of Comp and the anchor verb. We already have the beginnings of a raising pattern, because we have one item, the extracted one, which is dependent on two other items: the anchor verb, and the word on which it 'really' depends, and from which it receives its semantic relation (e.g. *like* in *Apples I know ...I should like*). Let us call the latter word its 'real head', for short. What is needed to complete the pattern is a direct link between the anchor verb and the real head, but of course this we do not have so far. Indeed, the interest of extraction is precisely that the dependencies involved are unbounded - an unlimited number of dependency links may separate the anchor verb from the real head. Diagram (6) shows the relevant dependencies without labelling them, with the dependency between *apples* and its anchor verb written below the line for convenience:



It is generally accepted nowadays that extraction needs a 'hopping' analysis - e.g. from one Comp position to the next, as in GB, or from a phrase to one of its daughters, as in GPSG - rather than a 'swooping' analysis in the tradition of classical LFG (Kaplan and Bresnan 1982). Perhaps the most important evidence for this comes from the Island Constraints (e.g. the rule that nothing can be extracted from inside a relative clause) which can cut the chain of links at any point, including the middle - e.g. the badness of an example like **Apples I know someone who thinks one shouldn't eat would be hard to explain*

if the displacement of *apples* involved only this word, its anchor verb (*know*) and its real head (*eat*).

Suppose, then, that we assume the dependency equivalent of a hopping analysis, an analysis in which the displaced word is linked in turn to each word in the dependency chain between the anchor verb and its real head. This gives (7), an enriched version of the structure in (6).




Inspection of this structure reveals that it is actually made up of a series of raising structures. The first involves *apples*, *know* and *think*, the second *apples*, *think* and *should*, and the third *apples*, *should* and *like*. In each case, *apples* is shared, as a dependent, by the other two words; and in each case the first verb is the head of the second. (*Apples* actually has two separate dependency links to *like*, but that is another story, which is also interesting because it shows a similarity between long-distance dependencies and passivization, explained in Hudson 1989, 1990.)

What we have seen in this section is that a single abstract pattern, the raising pattern, can be identified in several different 'difficult' areas of syntax. This in itself is worth pointing out, but it is not the end of the matter, as we shall see when we try to broaden the search for raising patterns to include semantics.

3 Raising patterns in semantics

Raising patterns are obvious in semantics if we look at syntactic control structures. For example, take a sentence like *Fred tried to kiss Mary*, whose syntactic structure is (8).




I have used the LFG term 'xcomp' (Bresnan 1982a), rather than my own usual term ('incomplement'), to help readers familiar with LFG. It means a complement whose subject is shared with its head. Again ignoring *to*, we find that the word *Fred* is the subject of two verbs, *tried* and *kiss*, so we should expect that the semantic element corresponding to *Fred* (which for the time being we can call just 'Fred') will also be directly related to both the semantic elements which match these two verbs. If so, this will be an example of a

semantic raising pattern, by our definition, provided the verbs' semantic elements are directly linked to one another, and provided all the links are dependencies.

First, then, are semantic links dependencies? A dependency is an asymmetrical relation between two elements, one of which provides the only connection between these two elements and any structure within which they are embedded. For example, in syntax we say that attributive adjectives depend on the noun, rather than vice versa, because the combination [adjective + noun] occurs where the grammar allows nouns to occur (e.g. in the direct object position), rather than where it allows adjectives. Similarly, we can be sure that 'Fred' depends on 'tried' because if we embedded the whole in a larger structure, such as *I think ...*, the external links would involve 'tried', not 'Fred' (because the object of THINK has to refer to a thought, not a person - cf. **I think Fred*).

The conclusion, then, is that at least some relations in semantics are dependencies - but this is so widely believed that it hardly seems necessary to argue the point. As I noted earlier, any analysis in terms of functors and arguments implies a dependency analysis, because it is the functor that provides the links to larger structures. What is not so widely recognised, of course, is that this constitutes an important formal similarity between semantic and syntactic structures; but this point is overlooked because most linguists do not think that syntax is structured in terms of dependencies.


What about the link between *tried* and *kiss* in *Fred tried to kiss Mary*? As far as the syntax is concerned, *kiss* is not a direct dependent of *tried*, because *to* separates them (in terms of dependencies as well as word order). However, *to* (presumably) has no meaning of its own, so we can make it 'borrow' its meaning from the following infinitive; whatever *Fred kiss Mary* means, *to* means the same. (This is the same treatment that we give to any semantically empty words, such as auxiliary DO or prepositions like the *on* of DEPEND ON which are lexically selected.) The semantic structure for our sentence is thus as in (9), which includes the syntactic structure of (8).



This semantic structure begs a very large number of questions (Hudson 1990), but I hope the main features will be comprehensible without detailed explanation. The dotted vertical lines connect each word to the part of the semantic structure that corresponds to it, which I shall simply call its 'meaning'; so the meaning of *Fred* is 'c1' (read as 'concept 1'), that of *tried* is 'c2', and so on.

The main point about this structure is the configuration in the top left corner, involving c1, c2 and c3, whose geometry can be seen more clearly in (10). This is precisely the same as one of the two raising patterns that we found in syntax, (3): the first element depends on both the other elements, and

the second is the head of the third.




This, then, is our first example of a raising pattern in semantics, but there are more to follow.

We must pause briefly to consider the nature of these semantic relations. I have provided labels for the dependents (i.e. arguments) of c3, the meaning of *kiss*: 'kiss-er' and 'kiss-ee'. These are in fact meant to be taken more seriously than might be expected, but that again is another story. The question is what we should call the relations between the meaning of *tried*, c2, and its arguments, which are shown by '?'. One possibility, of course, is 'try-er' and 'try-ee', but I think we can do better in this case, if we consider what TRY actually means. If I try to do something, what do I do? We know that I must do something, but the verb gives no indication whatsoever of the nature of my behaviour, which might range from writing a letter to becoming a monk. What the verb TRY does specify, however, is my purpose in doing it - my answer to a question containing WHAT FOR? So if someone sees me standing on my head, and says *What are you doing that for?*, a perfectly direct answer would be *I'm trying to make my hair grow*, because this specifies the purpose directly. Similarly with our sentence *Fred tried to kiss Mary*: what this means is that Fred did something with the purpose of kissing Mary. Assuming this analysis, the infinitival argument of the meaning of TRY is its purpose.


Before we look at the other argument of TRY it is worth considering another very similar argument-type, result. We can see the difference between purposes and results very clearly if we change the sentence to *Fred managed to kiss Mary*. From this sentence it follows that Fred did kiss Mary, whereas this does not follow from the sentence with *tried*. The two sentences form a minimal pair both in syntax and in semantics, because their structures are identical apart from containing TRY or MANAGE, and having the label 'purpose' or 'result' on the infinitival argument. We shall see that between them these two argument-types account for a surprisingly large number of 'syntactically obscure' semantic relations such as the ones we have been considering. Of course, they can also be indicated by their own separate syntactic realisations, adjuncts introduced by words like SO AS TO, SO THAT, IN ORDER TO or just TO - which leads into yet another general topic which we cannot develop here, about semantic relations which are sometimes realised as adjuncts and sometimes as complements or even internally to a word.

The first argument, realised in our examples by *Fred*, is easier to label. According to our paraphrase, Fred did something, though we do not know what; an unspecified 'doing' is an action, and so Fred is the actor. We can now complete the semantic structure for *Fred tried to kiss Mary*, concentrating on the raising pattern shown in (10); this is given in (11a), and (11b) gives the structure for *Fred managed to kiss Mary*.




It is easy to see that the pattern in (11b) occurs in the semantic analysis needed for a sentence like *Fred persuaded Mary to kiss him*, which means that Fred did something to Mary whose result was that Mary kissed him. There does not seem to be a minimal pair contrasting 'result' with 'purpose' for PERSUADE, but TELL (or ASK) is similar: *Fred told Mary to kiss him* means that Fred did something to Mary, namely talking to her, whose purpose was that she should kiss him. In both these cases we again find the raising pattern, because the person affected by Fred's action, whom we can call the patient, is also the actor of the purpose or result, which in turn is directly dependent on the action.

In the examples considered so far the purpose or result has been realised by a separate phrase, based on an infinitive, but this is not always so. Sometimes it is built into the semantic structure of the verb, as in another use of the verb TRY which we can call TRY/trans, in contrast with TRY/infin: *Fred tried the apple / bed / film*. In spite of its syntax, this is quite similar in meaning to the TRY/infin examples, because it means that Fred did something to the apple (etc), whose purpose was that he should find out what it was like. What he did to it is left completely unspecified (just as with TRY/infin), and clearly varies from case to case (e.g. he ate part of the apple, he lay on the bed, he watched part of the film). And as with TRY/infin, the actor is also an argument of the purpose; so if Fred tries the apple, the purpose is that he, not someone else, should find out what the apple is like. This produces one raising pattern, but another one arises because the patient (the apple) is also an argument of the purpose - it is the quality of the apple, not of something else, that Fred wants to investigate. These comments suggest a semantic structure along the following lines for *Fred tried it*:



The raising patterns are shown separately in (13) and (14).




What is striking about this example is how much semantic structure is built into the meaning of one word, TRY/trans. But this is not at all uncommon

in semantics, so let us consider some other densely packed words, in the search for raising patterns. Take the notorious KILL, for example. We all know that in some sense KILL means 'cause to die', and yet we also know, following arguments given by Fodor (1970), that the semantic structure of a sentence containing KILL cannot be the same as one containing CAUSE TO DIE, because (inter alia) KILL means that the 'causing' action and the dying must be simultaneous, in contrast with CAUSE TO DIE. For example, suppose that Fred set a trap on Sunday which killed a mouse on Wednesday; we could (perhaps) describe this as *On Sunday Fred caused a mouse to die on Wednesday*, but not as *On Sunday Fred killed a mouse on Wednesday*. (Analytical points like this were at the heart of the debate about Generative Semantics summarised in Newmeyer (1980: 112ff, 164ff).)



The explanation for this difference between KILL and CAUSE TO DIE is easy, provided the semantic structure is clearly distinct from the syntactic structure. Roughly speaking, the reason for the difference lies in the lexical entry for KILL, whose semantic structure is indeed similar to the structure of CAUSE TO DIE, but with the important difference that the time of the causing action and the time of the dying have to be the same. As we shall see, this is another example of a raising pattern.

What we have already said provides an analysis for all causatives as well, because if x causes y, x does something whose result is y. The sentence *Fred killed a mouse* tells us nothing about the nature of Fred's activity (though this can be made more specific - e.g. by using a verb like SHOOT or POISON), but does tell us what its effect was; and the same is true of *Fred made a mouse die*, *Fred got Mary angry*, *Fred got Mary to kill a mouse*, and a host of other causative constructions. I shall therefore not treat 'cause' as a primitive (contrary to most other analyses), but as analysable into action and result - which has the desirable consequence, incidentally, that the causer has the same semantic role as any other actor. There seems to be a general rule that when an action and its result are both built into the semantic structure of a single word, they must be simultaneous. To take another example, 'x SEND y to z' means that x does something whose result is that y goes to z, where the action and the start of the going are (at least notionally) simultaneous.

Returning to our example of KILL, then, we find that its semantic structure must contain the pattern in (15).



In this structure we can isolate two raising patterns:



The analyses that we have given for these examples are easy to generalise to other lexical items with causative, resultative or purposive meanings - examples like PUT, SEND, ROUGHEN, AVOID, REACH, TEMPT and LOOK FOR, not


to mention the more obvious ones like GET and INTEND, where the semantic structure can be read more directly off the syntactic structure. In sum, there can be no doubt that raising patterns exist in semantics even when there is no corresponding raising pattern in the syntax.

4 The function of raising


We can now ask a very general question about these raising patterns: why are they so common? The short answer is, of course, that we do not know and perhaps cannot know, but we can at least hazard a guess.

Let us assume that our knowledge consists of a network of concepts connected by relations of various types, at least some of which are dependencies (in the sense defined earlier, where a dependent is linked to the larger structure only via its head). Now suppose one of the concepts is 'kill' - a reasonable assumption, since otherwise there would be no concept to serve as the meaning for KILL. Similarly we presumably have a concept 'die', but these must be connected to each other, as action and result; and 'die' in turn is connected, again through a result link, to 'alive', and so on.

But these result links are not the only dependencies which the concepts have - each one may also have an actor, a patient, a time, and so on, depending on what kind of situation it refers to. What connections are there between the dependents of one concept and those of its head? Suppose there were none, and we had a network which could be displayed as in (18).




This network would be quite useless, because a complete list of the dependents of 'kill' would not tell us anything about the dying, beyond the fact that it was a result of killing. In order to say who died and when, it would be necessary to specify this information separately, in terms of a complex relation like 'die-er of result of kill', or 'time of result of kill'. How much more useful, and simple, if the 'higher' dependencies were bound to the 'lower' dependencies, as in (19).



This is a raising pattern - more precisely, it is two raising patterns - and my suggestion is that raising patterns exist in order to improve the flow of information through a network. If higher dependents are bound in this way to dependents of lower ones (e.g. to time and die-er of a result), then the lower concepts can be used to define the higher ones - which means that the dependency asymmetry can also be linked to the relative 'complexity' of the concepts concerned, a head concept always being defined in terms of its dependents, and therefore more complex than them. Without the binding, the lower concepts contribute very little information to the higher one; but each binding increases the information contribution.



The raising pattern has two characteristics, of which we have so far explained only one: why there is one dependent of the higher element which is shared with (i.e. bound to) the lower element. What about the restriction that the relation between the higher and the lower elements must be direct - i.e. that a raising pattern is strictly a three-element structure, to which no intermediate nodes can be added, rather than one like (20)?



One explanation for the restriction is that it is simply a matter of definition: a structure like (20) just does not satisfy our definition of 'raising structure'. The question would still arise, however, why there are so many examples that do satisfy our definition; and again the explanation may be a rather uninteresting one, namely that structures with three elements are bound to be more common than structures with four elements, on the assumption that simpler structures are more common than more complex ones.

But a more interesting explanation comes from the rather natural assumption that complex structures are built out of simpler ones - e.g. 'kill' contains 'die', and 'die' contains 'alive'. If this is so, then structures like (20) are unlikely to arise, because the relation 'X of D' is likely to have already been bound to some argument of C before the latter was embedded in the larger structure whose root is B. For example, the concept 'die' has the argument 'die-er', which is bound to the 'be-er' of 'not alive', which is the result of 'die'; so when 'die' is embedded in the concept 'kill', there is no need to relate the latter's patient directly to the 'be-er' of 'not alive', because the same effect could be achieved more simply by referring to the 'die-er' of 'die'. Moreover, the argument 'die-er' can be applied to other concepts that contain 'die', such as 'resurrect', so overall there is a great saving in the number and complexity of links in the system, even though the individual structures are more complex.

Considerations of both storage and processing thus seem to favour a structure in which all the relations are strictly local - an extension of the principle of 'hopping' rather than 'swooping' that syntacticians have discovered. Putting it abstractly, this means that structure (21a) is preferred to (21b) as a way of relating A to Z - although (21b) looks simpler, because it involves fewer connections.



These arguments seem to suggest, then, that raising structures have a very natural and important place in any large informational structure such as a language. Not surprisingly, then, we have been able to find evidence for them both in syntactic structure and in semantic structure.

5 Raising patterns in cognition

We now broaden our horizons even further, to consider examples of raising patterns whose elements are neither words nor semantic elements as defined above. I shall call these elements 'concepts' and the structures containing them 'conceptual structures'. According to Word Grammar, both words and their meanings are particular kinds of concept, so our task is to find evidence for raising structures whose elements are concepts which are neither words nor word-meanings.

A background assumption which I should make explicit is that an utterance's 'content' is not exhausted by its semantic structure, the structure which is directly related to the words. Our knowledge of the world and our ability to draw inferences allow us to build much fuller and more precise structures on top of the semantic structure by supplying referents, defining illocutionary forces, and so on. For example, the semantic structure for *He ran* is neutral as to who 'he' is and when the running happened, but these facts are supplied by the hearer in a successful interpretation. This assumption in itself is quite uncontroversial, but it carries with it the implication that the two kinds of structure are formally similar, and differ only in the amount of detail that they contain. In other words, at least some parts of the 'message' conveyed by an utterance are formally similar to the utterance's semantic structure - which we have just argued is formally similar to its syntactic structure (Fodor 1975).

My main source of evidence for raising structures in general cognition is the expansion of utterance-meanings, though no doubt other sources could be tapped, such as the acquisition of new word-meanings. We shall see that the cognitive structures that exist independently of language must include some raising structures.

Consider the meanings of verbs like DO and USE, as in (22).

(22a) We did the beds.

(22b) We used the beds.

In both cases the verb's meaning is an action which affects the meaning of its object, so we can assume a simple semantic structure in which 'we' is the actor and 'the beds' is the patient. (In this discussion I shall concentrate on cases where the object is concrete, ignoring examples like *do a job* or *do a dance*, whose semantic analysis is rather different.) However there is clearly a great deal more than this to the total message conveyed by these sentences, because we all know that the action described by (22a) is quite different from the one described by e.g. *We did the spades*, although the verb is the same in both cases, and similarly for (22b) compared with *We used the spades*.

There are two reasons for being certain that this information comes from our knowledge of the world, and not from the meanings of the verbs DO and USE. One is that the different objects which imply different actions can easily be coordinated:

- (23a) We did the beds and the spades.
 (23b) We used the beds and the spades.

This fact shows clearly that DO and USE each have a single general meaning which is equally compatible with beds or with spades as object. We can therefore eliminate the possibility that these verbs are polysemous, with different meanings for different types of object.


The other reason for being certain that the action is defined by knowledge of the world and not in the grammar is that the action may vary from context to context. Indeed, this must be the case for *We did the spades*, because so far as I know there is no standard action which is performed on spades, so whatever this sentence means must be defined by the local context. In contrast, of course, *We did the beds* normally means 'We tidied the beds', though in context it could refer to almost any kind of action that can be performed on a bed - painting, burning, making, etc.

This added information about actions must be represented in terms of some kind of structure, as part of our world knowledge; and it is reasonable to assume that we acquire much of it without any help from language (by watching people acting on different objects). For each type of object we must know what kind of action involves it, first, as a 'do-ee', and second, as a 'use-ee'. (It should be clear from the examples given so far that these two action-types are quite distinct; e.g. using a bed involves a complete different action from doing a bed.) And although we can only guess at the structures in terms of which this knowledge is represented, our earlier assumptions suggest that these structures must be very similar to semantic structures. This must be so if the message conveyed by an utterance is a more fully specified version of its semantic structure. Let us call the structures which define these action types 'action-schemas'.


It should be noticed that the semantic relations play a crucial role in this analysis, because both DO and USE refer to 'an action', without further differentiation. I assume that their objects have the semantic roles 'patient' and 'instrument' respectively - the latter an exceptional instrument which is realised by a noun-phrase, not by WITH. These two roles can be combined in the same sentence, as in *He did the beds with the spades*, which again refers to some action which is unspecified except for its actor, its time, its patient and its instrument.

The model of processing, then, is one in which an utterance is assigned a semantic structure on the basis of purely grammatical information, and then this structure is expanded by consulting (inter alia) the relevant action-schemas. Thus we understand *He did the beds* by first constructing a semantic structure in which 'the beds' is patient of an action which is unspecified (except for its actor and its time), and then finding an action-schema relevant to beds as patients of actions. The procedure for *He used the beds* is similar, except that the semantic relation between 'the beds' and the action is 'instrument', not 'patient', so we are directed to a different action-schema for extra information.

What I shall show is that the action-schema for patients contains a raising pattern, but we start with the simpler case of action-schemas for instruments. We need an action-schema which will tell us what kind of action involves a spade as instrument. It seems fair to assume a very simple structure, in which some kind of action (e.g. digging) is linked directly to the concept 'spade' by an 'instrument' relation:




Once this little structure is in place, we can use it to interpret a structure like (25), which is part of the semantic structure of *He used a spade*, where c1 is an action, and c2 is a spade.




Since a spade is typically an instrument of digging, we infer that c1 is an instance of digging.

Now consider the more complex case of DO, where the relevant relation is 'patient'. DO is different from USE because what is common to all the different uses of DO with some particular kind of object is not that they all involve the same kind of action, but rather that they all lead to the same resulting state of the object. For example, in *He did the beans* we can infer that when he had finished the beans were 'done', but we do not know how he produced this result - e.g. whether by using a machine or by hand. So one crucial bit of information about beans is what beans are like when they are considered 'done' - which of course may well vary with context. For simplicity let us assume that beans are clean when they are done. Part of this information is contained in (26).




But this structure misses a crucial fact: beans are clean only when they are 'done', and not otherwise. In other words, this state is always the result of being acted upon. The only way in which this extra fact can be stated is by means of a raising structure, (27).



This raising analysis is supported by another fact, which is that in some cases the action-type can also be specified more precisely. For example, I know that the standard tool for 'doing' potatoes is a potato-peeler, so the schema for doing potatoes must have room for an instrument, which in turn presupposes an action. The structure in this case would be the following, which incidentally also indicates the typical application of a potato-peeler.

(28)



Similarly, we can imagine a situation where 'doing the chairs' means bringing the chairs in from the garden, where the action is not just an indefinite 'something whose result is that the chairs come in from the garden', but an example of carrying. What all these examples have in common is that they all require an action node as well as the object node itself and the result-state node, and that these three nodes are in a raising pattern. And of course the structure is part of general cognition, without any sort of dependence on linguistic structures.

6 Conclusion

We have seen that the raising patterns which are now well-known to grammarians (and in some cases uncontroversial) also crop up not only in semantic structure, where they could possibly be explained as the effects of syntactic structuring of experience, but even in general-knowledge structures which seem to be completely independent of language. This conclusion admittedly stands at the end of a fairly long chain of assumptions any one of which is open to question, but it seems reasonably well supported by evidence. Let us now extend the chain by two more links, to stress the importance of this conclusion (if it is true).

First, it suggests that the formal properties of the structures that contain syntactic, semantic and general knowledge cannot be very different if the same formal pattern plays as prominent a role in all three as I have suggested. Even if the pattern had been a very simple one its occurrence in all three kinds of knowledge might have counted as a significant finding - e.g. in my opinion it is quite significant that it is possible to analyse all three in terms of dependencies. But the raising pattern is fairly complex, allowing as it does for three nodes to be related to one another by means of dependencies. Its occurrence across different types of knowledge is therefore all the more significant for that. Nor is the raising pattern so vague that it could be imagined in any formal pattern; provided we can decide whether or not a relation is a dependency it is always crystal clear whether, formally speaking, some pattern counts as a raising pattern.

The discovery of this rather peculiar pattern in all these types of knowledge would be hard to reconcile with any claim that they have fundamentally different kinds of structure. And yet this claim appears to be widely accepted in linguistics and cognitive science.

The second conclusion is that whatever explanation we find for the existence of these patterns in one kind of knowledge should generalise to all other kinds. In section 4 I tried to build an explanation in terms of the needs of a processor and storer of information, stressing the advantages of raising and of locality, the two components of the raising pattern. But I also showed, in the section on syntax, that raising is the basis for some of the most difficult syntactic structures (including unbounded dependencies). If we put these two

points together, we find we have a ready-made explanation for at least some of the properties of these structures, as special cases of the 'pan-cognitive' raising pattern. I take it that an explanation in such terms is to be preferred to one in terms of arbitrary innate structures specific to syntax.

Acknowledgements

I owe a particularly large debt to John Fletcher and Norman Fraser for detailed comments and discussion of an earlier draft. I also had helpful comments from And Rosta and Michael Brody.

References

- Bresnan, J. (1982a) Control and complementation. In Bresnan (1982b: 282-390).
 Bresnan, J. (ed.) (1982b) *The Mental Representation of Grammatical Relations*. Cambridge, MA: MIT Press.
 Chomsky, N. (1981/82) *Lectures on Government and Binding. The Pisa Lectures*. Dordrecht: Foris.
 Dik, S. (1978) *Functional Grammar*. Amsterdam: North Holland.
 Fodor, J. A. (1970) Three reasons for not deriving 'Kill' from 'Cause to die'. *Linguistic Inquiry*, 1. 429-438.
 Fodor, J. A. (1975) *The Language of Thought*. New York: Crowell.
 Gazdar, G., Klein, E., Pullum, G. and Sag, I. (1985) *Generalised Phrase Structure Grammar*. Oxford: Blackwell.
 Halliday, M. A. K. (1985) *An Introduction to Functional Grammar*. London: Arnold.
 Heringer, H.-J. (1970) *Theorie der Deutschen Syntax*. München: Max Hueber.
 Hudson, R. (1984) *Word Grammar*. Oxford: Blackwell.
 Hudson, R. (1988) Extraction and grammatical relations. *Lingua*, 76. 177-208.
 Hudson, R. (1989) English passives, grammatical relations and default inheritance. *Lingua*, 79. 17-48.
 Hudson, R. (1990). *A Word Grammar of English*. Oxford: Blackwell.
 Johnson-Laird, P. (1983) *Mental Models*. Cambridge: Cambridge University Press.
 Kaplan, R. and Bresnan, J. (1982) Lexical-functional grammar: a formal system for grammatical representation. In Bresnan (1982c: 173-281).
 Kempson, R. (1988) Grammar and conversational principles. In F. Newmeyer (ed.) *Linguistics: The Cambridge Survey II. Linguistic Theory: Extensions and implications*. Cambridge: Cambridge University Press.
 Lamb, S. (1966) *Outline of Stratificational Grammar*. Washington, D.C.: Georgetown University Press.
 Minsky, M. (1974) A framework for representing knowledge. MIT AI Laboratory, Memo 306.
 Newmeyer, F. (1980) *Linguistic Theory in America. The first quarter century of Transformational Generative Grammar*. New York: Academic Press.
 Pollard, C. (1989) The syntax-semantics interface in a unification-based Phrase-structure grammar. mimeo.

- Pollard, C. and Sag, I. (1987) *Information-based Syntax and Semantics, I: Fundamentals*. Stanford: CSLI.
- Postal, M. (1974) *On Raising*. Cambridge, MA: MIT Press.
- Schank, R. (1975) *Conceptual Information Processing*. Amsterdam: North Holland.
- Siewierska, A. (forthcoming) *Functional Grammar*. London: Routledge.
- Wood, M. (forthcoming) *Categorial Grammar*. London: Routledge.