The difficulty of (so-called) self-embedded structures[1]

Richard Hudson

Abstract

Why is sentence (1) impossible to process?

(1) The rat the cat the dog chased ate died.

The standard explanations all focus on its syntactic structure, but the present paper offers an alternative explanation in terms of semantic structure. The syntactic account cannot explain why some sentences which are syntactically similar to (1) are much easier to process. The difficult sentences seem to be those in which an embedded clause which modifies a noun has its own subject modified by another clause whose subject is a common noun (not a pronoun). The explanation offered here is that the senses of the modified nouns are not sufficiently distinct to be maintained as separate concepts in short-term memory, so the various concepts stored `interfere' with one another.

[1]I should like to acknowledge help from Sue Blackwell, Annabel Cormack, Jennifer Ganger, Ted Gibson, Chris Golston, So Hiranuma, Steve Harlow, Richard Lewis, Caroline Liberg, John Limber, Bruce Nevin, Michael Niv, Neal Pearlmutter, Colin Phillips, Karin Stromswold, Patrick Sturt, David Wharton.

Why is sentence 0 so hard to process?

(1) The rat the cat the dog chased ate died.

The example was used by Miller and Chomsky (1963:286) to illustrate the effects of limited working memory on the processing of complex syntactic structures. Their conclusion relates the difficulties to the `parenthetical embedding' in the syntactic structure:

(2)There can be little doubt that natural languages permit ... parenthetical embedding ... For example, [our 0] is surely confusing and improbable but it is perfectly grammatical and has a clear and unambiguous meaning. (Miller and Chomsky 1963:286)

Since then the consensus has been that the problem lies in the syntactic structure: if a sentence has a structure like 0, then it must be very hard to process. The following quotations are typical:

(3)Miller and Chomsky 1963 noted that sentences with multiple center-embeddings are invariably unacceptable .... (Newmeyer 1988:7)

(4)Following Miller and Chomsky (1963) and Chomsky and Miller (1963), all current theories of center-embedding propose that the difficulty that people have in understanding sentences like [*The frightened child who the old woman who the rescue worker looked for had comforted greatly survived the crash*.] is rooted in a limitation on memory. The memory requirements associated with a partial parse are a function of 1) the number of unattached structures that the parser is retaining; and 2) the number of categories that have been predicted and that must be completed at that point. (Babyonyshev and Gibson 1985)

(5)What boggles the human parser is not the amount of memory needed

but the *kind* of memory: keeping a particular kind of phrase in memory, intending to get back to it, at the same time as it is analyzing another example of *that very same kind of phrase*. Examples of these `recursive' structures include a relative clause in the middle of the same kind of relative clause, or an *if ... then* sentence inside another *if ... then* sentence. (Pinker 1994:207)

The aim of this paper is to question the received wisdom.

One of the problems in discussing these matters is the great uncertainty about definition and terminology: both the above quotations refer to `center-embedding' (for simplicity I shall use the American spelling), but should it not be `self-embedding', or `multiple self-embedding'? Or perhaps it should be even more precise, `multiple self-embedded object-relative clauses' or even `multiple self-embedded object-relative clauses modifying a subject'? (De Roeck et al 1982 is a helpful survey of the confusion.) Chomsky and Miller cast their net wide; having defined self-embedding as a variety of center-embedding (Chomsky and Miller 1963:290) they conclude that self-embedding is not the only source of processing problems:

(6).. nested [i.e. center-embedded] structures even without self-embedding quickly become difficult or impossible to understand. ... sentences of natural languages containing nested dependencies or self-embedding beyond a certain point should be impossible for (unaided) native speakers to understand. This is indeed the case ... There are many syntactic devices available in English ... for the construction of sentences with nested dependencies. These devices, if permitted to operate freely, will quickly generate sentences that exceed the perceptual capacities (i.e., in this case, the short-term

memory) of the native speakers of the language. (Miller and Chomsky 1963:471)

Nevertheless, virtually all the subsequent discussions have focussed on examples like 0, which are in fact self-embedding (relative clause inside relative clause, or at least clause within clause), even while using the term `center-embedding' (as in the quotation in 0). Clearly these questions matter a great deal in the hunt for an explanation, because different definitions of the problem will inevitably lead to different solutions. The discussion will proceed by gradually clarifying the focus until we reach a point where we can say what it is that leads to the complete breakdown of processing that we all experience with sentences like 0. Needless to say, sentence 0 will itself belong to the category that we define, but we shall see that it qualifies by virtue of its semantics as well as its syntax.

1. Multiple center- and self-embedding?

Is it multiple center-embedding or self-embedding as such that causes breakdown? These notions are often used in an unclear way, and are sometimes even treated as synonyms (Crystal 1980/1991:309), so let's make the strongest assumption possible by applying them together: maybe breakdown comes from structures which are both center-embedded and self-embedded?

This cannot be generally true because there are languages where such structures are regularly used (and, presumably, understood) in possessive patterns. One such is Ancient Greek. According to Chris Golston (pc), possessive genitive NPs regularly occur centre-embedded in NP, which could in turn be genitive as in the following example from

Plato (Politicus 281a):

(7)[to[tees[touksainontos]tekhnees]ergon]

the the the wool-carder artwork

ACC GEN GENGEN GENACC

The innermost NPs in this example must, surely, be both self-embedded and center-embedded, and yet the phrase is presumably acceptable. (Of course we don't in fact know whether Ancient Greeks could process such examples, but if they occurred regularly in the corpus of classical writing this does seem a reasonable assumption.) Another language which allows this pattern is Moru (Hagège 1976, quoted in Trask 1993: 40).

(8)kokyE[toko[odrupi[ma ro]ro]ri]drate.

dog wife brother me ofofofis-dead

`My brother's wife's dog is dead.'

However, what we are concerned with is certainly an example of both self-embedding and centre-embedding, with one relative clause centre-embedded in another. What we need is a more precise notion than either of these. There are two ways to make it more precise: by relating the centre-embedding to the position of the head, and by relating it to the category of the embedded constituents. The next section will make both of these moves at the same time.

2. Pre-head subordinate clauses?

Maybe the processor collapses whenever the self-embedded constituent is a clause (in contrast with the possessive NPs discussed above)? And maybe the collapse happens only when the innermost clause precedes the head of its `host' clause? This is still not specific enough, because no

collapse seems to happen with adverbial clauses such as the following. (In case readers doubt this claim, the Appendix gives some experimental evidence in support of it; the examples quoted are based on those used in the experiment[2].)

(9)a If [although [you try hard] you still don't understand], tell me.

b Although [when [people drink vodka] they get headaches] some people forget.

It is easy to find textual examples like the following:

(10)For a moment she felt almost frightened, as if, [although [this was what she thought an impossible hope], it had turned, when realised, into something she did not altogether like]. (Larkin 1947/1975: 135)

De Roeck et al 1982 and Sampson 1996 give a great many more textual examples of `center-embedding'.

Is it just relative clauses that cause these problems? According to Gibson et al 1995, when students rate sentences for `goodness', they rate sentences like the following much higher than those like 0:

(11)The confession [that the fact [that the noble huntsman was kissing the sleeping princess] might lift the curse] upset the evil witch-

---

[2]It is a matter of debate whether or not these embedded clauses are center-embedded. It all depends on where we think the clause `proper' begins: are if and although part of the clause which they introduce? If they are, the bracketing is as follows:
a. [If [although you try hard] you still don't understand], tell me.
b. [Although [when people drink vodka] they get headaches] some people forget.
Under this analysis, neither of these sentences is a true example of double center-embedding. However I think it is quite reasonable, and quite widely accepted, that the introductory word if or although is like a preposition; in GB terms, it is not a complementiser but takes a CP as its complement. According to this analysis, both the subordinate clauses are truly center-embedded.

queen.

In this sentence one `content clause' is embedded inside another, and precedes the latter's verb, but 48 students gave such examples an average rating of 2.22 on a scale from 1 (good) to 5 (bad). In contrast, they rated examples with relative clauses at 3.07, a considerably worse mark.

I found something similar in the small-scale investigation which is reported in the Appendix, and which used a different method from Gibson's - asking students to write down a sentence immediately after hearing it read out. In one of these experiments the students had to hold 11 words in memory for each sentence, which is virtually impossible if the words constitute a random list, but very easy if they make up an easily processed sentence. Only one student had any trouble with sentence 0, which contains one content clause embedded in another.

(12)The question [whether the fact [that he sniffs] annoys us] is interesting.

On the other hand, when the content clauses were similar they seemed to be rather hard to process: 36% (out of 36) failed to remember the last ten words accurately in 0, and 30% failed on 0.

(13)The fact [that the fact [that he snores] surprised us] is interesting.

(14)The idea [that the fact [he lied] was obvious] worried me.

As we shall see below, this performance is far worse than for some examples with relative clauses, so it should be taken seriously.

Furthermore, Gibson et al showed that content clauses do cause problems when they occur inside a relative clause in its subject position, as in the next two examples:

(15)a The old woman [who the information [that the frightened child

survived the crash] had comforted] looked for the rescue worker.

b The old woman [who [that the frightened child survived the crash] had comforted] looked for the rescue worker.

The average rating for examples like (b), where the content clause is the relative clause's whole subject, was worse than for any other sentence-type: 3.98. This very poor rating may well reflect the general difficulty of content clauses used as subjects. After all, a tensed content clause is hard enough to deal with when it is the subject even in a main clause (hence the need for extraposition), and seems to be equally bad in any kind of subordinate clause. Even adverbial clauses, which are generally quite easy to process when nested as we noted above, become very hard when their subject is a content clause. Examples like the following strike me as comparable in difficulty to Gibson's relative clause example:

(16)[Although [that the frightened child survived the crash] comforted her], the old woman looked for the rescue worker.

In the absence of further evidence I shall assume that we should ignore examples like these, in which a content clause acts as subject; they are very difficult to process, but for reasons that have nothing to do with the self-embedded relative-clause patterns that we are exploring here.

In contrast, content clauses that complement a noun like *information* in 0a are highly relevant. The rating reported by Gibson et al for examples like this was 3.38, which was roughly the same as for the cases with two relative clauses; and in general content clauses are no harder to process than relative clauses when they complement nouns. For example, the addition of *the fact* to the last example improves it considerably:

(17)[Although [the fact that the frightened child survived the crash] comforted her], the old woman looked for the rescue worker.

I conclude, therefore, that content clauses cause processing problems for much the same reason as relative clauses.

What do these two types of clause have in common? In all the relevant examples they both modify nouns - for instance, in 0a the relative clause *who the information ... had comforted* modifies *woman* (which woman?), and the content clause *that the frightened child survived the crash* modifies *information* (which information?). In both cases, the subordinate clause is needed in order to identify the noun's full meaning, so in principle the processor cannot complete the noun's semantics until after it has completed that of the modifying clause. We shall look at this process of modification in more detail in the last section, but for the present the only further fact that we should notice is a curious asymmetry between content and relative clauses. As noticed above, a content clause is as hard to process as a relative clause when it is inside a relative clause, as witness the difficulty of 0; but when the relative is inside the content clause, the difficulty vanishes (Gibson, Thomas and Babyonyshev 1995).

(18)The information [that the frightened child [who the rescue worker looked for] had survived the crash] comforted the woman.

I shall suggest a rather tentative solution to this problem in the last section.

The conclusion so far, then, is that special processing difficulty (which I have been calling a processing `collapse') is caused whenever one clause modifies the subject of another clause which in turn modifies a noun. It will be helpful to be able to refer to these various elements by simple names, so here is an abstract formula:

(19)N_1 [N_2 [N_3 .. V_3 ..] .. V_2 ..]

The innermost clause, which contains N_3 and V_3, modifies N_2, which is

(in) the subject of the clause headed by the verb V_2; and the clause [..N_2 ... V_2 ..] in turn modifies N_1. Here are some examples to illustrate the structure:

(20)a The rat_{N1} [the cat_{N2} [the dog_{N3} $chased_{V3}$] ate_{V2}] died.

b The old $woman_{N1}$ [who the $information_{N2}$ [that the frightened $child_{N3}$ $survived_{V3}$ the crash] had_{V2} comforted] looked for the rescue worker.

c The $information_{N1}$ [that the frightened $child_{N2}$ [who the rescue $worker_{N3}$ $looked_{V3}$ for] had_{V2} survived the crash] comforted the woman.

Where does this conclusion leave the notion of `center-embedding' with which we started? Would it help to describe the structure in 0 as one in which the lower relative clause, [N_3 .. V_3 ..], is center-embedded in the higher one? Or is the crucial fact about it that the lower one precedes the head (V_2) of the latter? There is some evidence that center-embedding is not, in fact, relevant. According to Hagège (1976), there are languages in which center-embedded relative clauses are normal and not particularly difficult to process; but these are languages in which relative clauses are signalled by markers at both ends. In Mbum, for example, these enclosing markers are formally identical to deictics: *àí* ... *núí*. `It is not possible to use *àí* at the beginning without *núí* at the end.' The following is an example with a single relative clause:

(21)úì àí mì zàŋzáŋ nú bèlbél

woman I met is beautiful.

`The woman I have met is beautiful.'

If center-embedding as such causes processing problems, sentences should be problematic if they contain center-embedded relative clauses:

... *àí* ...*àí* ... nú nú. But according to Hagège, they are not: `Sequences of three nú are not absent, but less frequent in conversation, although they occur in oral literature.' His example is the following:

(22)úì àí mì zàŋzáŋ bí pàk àí mù màmá nú nú bèlbél
woman I met in house you know is beautiful
`The woman that I have met in the house that you know is beautiful.'

It seems, then, that `center-embedding' as such is not problematic. What makes this example different from the unprocessable English ones is that the lower relative clause here follows the head of the higher one, the verb zàŋzáŋ, `met'; so we can probably conclude that processing problems arise only when [.. N_3 .. V_3 ..] precedes V_2 - and perhaps even more specifically, only when this clause also follows the modified noun N_2. It seems we can at least reject center-embedding as the explanation (a conclusion reached independently by Lewis 1996)

There is one major gap in the discussion so far: what about the position of N_1 in the top clause? In all the examples discussed so far N_1 has been the subject of the main verb; should we not include this as a condition for processing collapse? After all, it seems obvious that a long NP containing a relative clause is easier to process at the end of the sentence than in subject position, and we even have a special process - extraposition - for shifting the relative clause alone to the end of the sentence. We even have experimental evidence that single relative clauses are harder to process when they modify (sentence-initial) subjects than when they modify (sentence-final) objects (Stromswold et al forthcoming). However, rather surprisingly the position of N_1 turns out to be irrelevant in more complex sentences. Gibson, Thomas and

Babyonyshev report that nested relative clauses are judged just as bad at the end of the main clause as when they are in the subject position. Pairs like the following were given much the same ratings.

(23)a The child$_{N1}$ [who the old woman$_{N2}$ [who the rescue worker$_{N3}$ looked$_{V3}$ for] had$_{V2}$ comforted] survived the crash.

b The crash was survived by the child$_{N1}$ [who the old woman$_{N2}$ [who the rescue worker$_{N3}$ looked$_{V3}$ for] had$_{V2}$ comforted].

It seems, therefore, that we do not need to say anything at all about the relation between N_1 and V_1, the verb of the main clause.

3. Modifying clauses with modified subjects?

Is it enough, then, to say that processing is threatened by any clause which modifies a noun and whose own subject is modified by another clause? This definition says nothing about the internal structure of the innermost clause, the one headed by V_3; is this right? It seems clear that in at least one respect it is too general: if V_3 is a participle, processing is relatively easy. My evidence for this is the one sentence 0, which defeated only 6 out of 36 students.

(24)People$_{N1}$ [that footballers$_{N2}$ [hired$_{V3}$ by managers] attack$_{V2}$] get plenty of excitement.

Even easier examples can be constructed, such as the following:

(25)a Books$_{N1}$ [that people$_{N2}$ [studying$_{V3}$ linguistics] buy$_{V2}$] are expensive.

b Cars$_{N1}$ [that people$_{N2}$ [living$_{V3}$ in our street] own$_{V2}$] clutter the place up.

We hardly need experimental evidence to prove that these are much easier to process than the standard examples with finite relative clauses.

Why should this be so? Two possible answers suggest themselves. One is that a participle eliminates one of the three nouns, N_3. The other is that a participle is itself the head of the relative clause, so the word which tells the processor that N_2 is modified also supplies the most important part of the modifier's meaning.

Which of these explanations is right? Maybe they both are, but the first leads to a further prediction: that any other way of eliminating N_3 will lead to the same ease of processing. What if the innermost relative clause is a subject-relative? This will certainly eliminate N_3, because the relative pronoun or <u>that</u> occupies the subject position; but are such sentences as easy to process as those in 0 and 0? According to Gibson, Thomas and Babyonyshev (1995), they are not. They compared pairs like the following:

(26)a The child$_{N1}$ [who the old woman$_{N2}$ [who the rescue worker$_{N3}$ looked$_{V3}$ for] had$_{V2}$ comforted] survived the crash.

b The child$_{N1}$ [who the old woman$_{N2}$ [who looked$_{V3}$ for the rescue worker] had$_{V2}$ comforted] survived the crash.

What they found was that sentences like (b), with a subject-relative clause, were almost as hard as those with an object-relative clause. (The average rating for subject-relative examples was 3.17, contrasting with 3.40 for object-relative examples and 1.92 for easy examples with a relative clause inside a content clause, like 0.) The difference between the subject- and object-relative examples was significant, but only at the 0.05 level, in contrast with the 0.001-level significance of the earlier contrasts. These results suggest that the elimination of N_3 in itself does not make the sentence easy to process, because the subject-relative examples have no N_3, but are still relatively hard to process. On the

other hand, it is perhaps also important to recognise that the loss of N_3 does help a great deal. The evidence just quoted shows that we have less difficulty with subject-relatives than with object-relatives, even if the difference is less dramatic than in other cases. Three other pieces of evidence support this view. First, Sampson (1996) quotes two published examples where the innermost clause is a subject-relative:

(27)aThe only thing [that the words [that can lose -d] have in common]
is, apparently, that they are all quite common words.

bWhen the pain, [which nobody [who has not experienced it] can
imagine], finally arrives, they can be taken aback by its severity.

Second, Carpenter, Miyake and Just (1994:1081) survey four experimental projects on single relative clauses which showed that they are harder to process if it is the object that is relativised than if it is the subject; for example, it is possible to measure the time a reader devotes to each word, and using this method it turns out that object-relatives make readers spend more time on the lower and higher verbs than subject-relatives do.

The third piece of evidence is that my students had far less difficulty in remembering the eleven words of the subject-relative examples (a) and (b) in the following than in remembering the object-relative (c). The numbers show how many of the 14 students who were involved in this particular test failed to recall the sentence accurately.

(28)a The farmer$_{N1}$ [that the cow$_{N2}$ [that gave$_{V3}$ bad milk] kicked$_{V2}$] died. (4/14)

b Everything$_{N1}$ [the people$_{N2}$ [who came$_{V3}$] wanted$_{V2}$] was waiting in their rooms. (3/14)

c Beer$_{N1}$ [students$_{N2}$ [policemen$_{N3}$ follow$_{V3}$] buy$_{V2}$] comes from very many different places. (12/14)

In other words, 4 out of 14 students failed to remember (a) correctly, in contrast with 12 who failed on (c). This difference is particularly important because example 0a is from Hawkins (1994:5), where it is quoted as an example of a sentence which is impossible to process. In contrast, Hawkins claims, a comparable German sentence is unproblematic:

(29)Der Bauer, [der die Kuh, [die schlechte Milch gab], schlachtete], ist krank.

`The farmer [who killed the cow [which gave bad milk]] is sick.'

Contrary to what Hawkins claims, therefore, the goodness of this sentence can not be taken as evidence that German and English have different thresholds for self-embedding. The English example is itself much easier than the standard object-relative examples.

A conclusion that could reasonably be drawn from all this is that the elimination of N_3 helps, but does not explain in full why participial relative clauses are so much easier to process than those with finite relative clauses; accordingly, we must assume that the finiteness of the verb is also relevant. It would be reassuring to have direct experimental support for this conclusion, but meanwhile the indirect support is fairly strong. In any case, we can be sure that a participial clause does not make the processor crash, so we have to refine our definition of the

structures which do:

(30)The processor cannot handle a sentence in which a finite clause modifies a noun and has its own subject modified by another finite clause.

We could probably refine this definition further by referring to the different kinds of modifying clause. On the one hand, it would be good to know whether the processor is affected differently by non-restrictive relative clauses, as in the following examples.

(31)a Bill, [who Sam, [who Mary hates], admires], is a linguist.

b Rome, [where the Colisseum, [which Caesar built], is found], is hilly.

The following example is quoted by Sampson (1996) from a British tabloid newspaper, the News of the World (dated 1961):

(32)And yet a widow, [whose pension, [for which her husband paid], is wiped out because she works for a living wage], will now have to pay 12s 6d. for each lens in her spectacles, and 17s. 8d. for the frames.

And on the other hand, we should probably distinguish ordinary restrictive relatives from free relatives, which seem comparatively easy to understand:

(33)Everything$_{N1}$ [that what$_{N2}$ [Mary$_{N3}$ earned] bought] is in that brown suitcase.

Unfortunately we are very short of facts on these matters; all I can contribute is that my students handled the last example quite easily, with only 3/14 failures.

4. Finite noun-modifiers with finitely-modified subjects?

The generalisation in 0 takes the syntactic definition of problematic

sentences as far as our state of ignorance allows. But is it the case that any sentence that fits the definition is equally hard to process? It is not. Consider the following example, due to Robert Frank (1992):

(34)A book [that some Italian [I've never heard of] wrote] will be published soon by MIT Press.

I agree with Michael Niv (who gave me the example) that this is in fact very easy to process; and yet it fits the formula perfectly (in fact, if anything it is more complex than the formula requires because V_3 and V_1 are auxiliary verbs which support other verbs).

(35)..book [..Italian [..I've ..] ..wrote ..] .. will ..

N_1 N_2N_3 V_3 V_2 V_1

The following examples come from the literature on center-embedding.

(36)The game [those boys [I met] invented] resembles chess.

This is quoted in Smith (1989:56), where it is described as `inelegant but comprehensible and acceptable', although the next page seems to describe it as typical of sentences which are `notoriously difficult to understand'. I agree with the first verdict. Again, consider the following examples from Pinker (1994:206):

(37) aThe cheese [that some rats [I saw] were trying to eat] turned out to be rancid.

bThe policies [that the students [I know] object to most strenuously] are those pertaining to smoking.

Pinker comments that these are `a little hard because of the memory load but are not nearly as opaque as the [example 0]'.

(38)The rapidity [that the motion [that the wing [that the hummingbird has] has] has] is remarkable.

Pinker's examples are not really comparable, of course, because the unprocessable 0 has one more embedding than the other two, but his point is surely true even if we take a simpler example as a point of comparison:

(39) The rapidity [that the motion [that the wing has] has] is remarkable.

According to Pinker the crucial difference between this example and the easier ones is that the relative clauses are of the same type (both introduced by that), so it is truly self-embedded, but it should be noticed that Smith's easy example is also self-embedded in this sense, because the relative clauses are both of the same type (zero object-relatives). We shall suggest an easier explanation for Pinker's difference. Yet another example is reported in de Roeck et al 1982 (and Sampson 1996). This one was actually used in speech - in fact, during the discussion after a seminar in which Sampson had said that such sentences were always unprocessable!

(40).. but don't you find that sentences [that people [you know] produce] are easier to understand.

Interestingly, Sampson understood the example so easily that he did not at first notice its relevance. Lewis (1996) quotes two other examples:

(41)a The guy [whom the secretary [we fired] slept with] is a really lucky dog. (Kimball 1975)

b The rat [that the cat [I saw] chased] squeaked. (Kac 1981)

To complete our collection, Gibson (p.c.) attributes the following to Bever:

(42)The pictures [that the photographer [I met at the party] took] turned out very well.

It is quite easy to construct processable examples; here are some that I invented for use in my memory test, described in the Appendix; once again, the numbers show how many students failed to repeat them accurately:

(43)a The shells [that the children [we$_{N3}$ watched] collected] were piled up. (2/27)

b The apples [that the girls [Bill$_{N3}$ watched] collected] were heaped up. (2/14)

c Sentences [that students [I$_{N3}$ teach] memorise] can be quite short. (4/27)

d People [that things [I$_{N3}$ say] annoy] tend to tell me off. (1/36)

e People [that things [which I$_{N3}$ do] upset] always complain about the effect. (3/14)

These scores contrast sharply with those for the really hard sentences which I included, listed below:

(44)a People [that politicians [who journalists$_{N3}$ interview] criticise] can't defend themselves well. (22/27)

b People [that politicians [who journalists$_{N3}$ interview] criticise] have no legal protection. (27/36)

c People [that criminals [who the police$_{N3}$ recognise] attack] claim legal compensation. (14/25)

d Beer [students [policemen$_{N3}$ follow] buy] comes from very many different places. (12/14)

Why should these two lists of sentences affect processors in such different ways? Since the sentences were all presented in the same way (orally, with expressive intonation), the difference must lie in their internal structure, but it cannot be a matter of syntactic geometry

because the general syntactic patterns are the same throughout (except for the relative pronoun in the innermost relative clause, which is both present and absent in both groups). The most obvious difference lies in the identity of N_3, the subject of the innermost clause. In all the easy examples it is either a pronoun (I or we) or a name (Bill), in contrast with the common noun found in the difficult examples: journalists, the police, policemen. It is noticeable that all the other examples that have been found hard to process, including those cited by Gibson and colleagues, N_3 has been a common noun such as child, woman or rescue worker. These other examples are helpful because the NPs all have specific referents in contrast with the generic ones in my examples; between them they eliminate the conclusion that my hard examples are difficult because the referents are generic.

The evidence is admittedly flimsy, but the conclusion is suggestive: processing problems arise only when N_3, the subject of the innermost relative clause, is a common noun. (For simplicity I am ignoring determiners and the rest of the structure of the NP built round the common noun; controversy rages here, but is irrelevant for present purposes.)

(45)The processor cannot handle a sentence in which a finite clause modifies a noun and has its own subject modified by another finite clause whose subject is a common noun.

This is not the first mention in print of the difference between pronouns and common nouns (see Kac 1981, Niv 1993 and Lewis 1996); but nobody has been able to suggest an explanation for the difference because of the basically syntactic orientation that has dominated discussions. Why should it matter so much whether the innermost

subject is a common noun? My guess is that it is because of the semantics of common nouns. A common noun has both a sense and a referent; for example, if I say I read a book, you identify the referent of a book by first identifying its sense, the concept `book', and then creating a concept for an individual object that belongs to that concept. In contrast, pronouns and proper nouns have a referent but (arguably) no sense. To find the referent of I, you look for the person speaking at the moment, rather than for a member of a general category, and for the referent of Bill you search your memory for someone called Bill. My conclusion, then, is that a sentence is hard to process if N_3 has a sense. We shall see below why this is relevant to processing, and we shall relate it to the fact that the clauses in these unprocessable sentences modify nouns - meaning, we shall suggest, that each clause modifies a noun's sense. Word-senses will turn out to be the crux of our explanation for processing difficulties.

5. Some early data revisited

Most of the empirical work on the processing of self-embedded sentences was done in the 1960's, notably by Schlesinger (1966, 1968) and by Fodor and his colleagues (Fodor and Garrett 1967, Fodor, Garrett and Bever 1968). Each of these works throws extra light on the picture developed so far even though the research they report focusses on rather different questions from the ones we are discussing here.

Schlesinger's experiments involved rather small numbers of student subjects and even smaller numbers of test sentences (in English or Hebrew), but the results are interesting. Having read the sentences, the students were asked to judge their grammaticality, to paraphrase them,

and to recall them. In some of the sentences the meaning helped the processor by encouraging some noun-verb pairings and discouraging others, as in the following:

(46)This is the hole [that the rat [which our cat [which the dog bit] caught] ate].

Other sentences had unhelpful meanings:

(47)This is the boy [that the man [whom the lady [which our friend saw] knows] hit].

Yet other sentences involved verbs and nouns which were in the `wrong' order - i.e. the expected semantic links actually contradicted the syntactic order, as in:

(48)This is the hole [that the rat [which our cat [which the dog bit] made] caught].

In this sentence the syntax requires the rat to catch the hole and the cat to make the rat. We can call such sentences `misleading', in contrast with the `helpful' and `unhelpful' ones.

The results of these experiments suggest that semantics plays an important part in the processing. The subjects found the `helpful' sentences much easier to process than the `unhelpful' ones (e.g. four out of eight subjects understood helpful sentences after just one reading, in contrast with the unhelpful ones which no subjects at all could understand even after four readings (ibid:133); as for the `misleading' sentences, these were misunderstood by seven out of eight subjects - i.e. the remainder assumed the sensible meaning rather than the one consistent with the syntax. These results are very hard to reconcile with an efficient push-down store which stores the nouns one on top of the other, and then removes them in reverse order to pair them up with the

verbs. Of course, one way to reconcile these results with the standard syntactic explanation of processing difficulty is to say that the syntactic system collapses so that the semantic/pragmatic system has to take over; but that begs the main question, which is why the syntactic system collapses. After all, we have seen that humans can in fact cope quite easily with center-embedded structures, so there must be something special about these relative-clause examples which makes them hard. From this perspective, the role of semantics is reversed: rather than being the solution to a problem caused by syntax, it is the cause. The syntax is quite straightforward and parsable, but it requires some support from the meaning. When the meaning is unhelpful or misleading, the whole system fails and the processor either gives up or resorts to purely pragmatic reasoning.

The two Fodor papers are relevant because they confirm the conclusion of the last section: processing difficulty is only loosely connected to syntactic structure. Both papers address questions about the details of how we process self-embedded sentences: Are we affected by the presence or absence of relative pronouns (Fodor and Garrett 1967)? Are we affected by valency ambiguity in V_1 or V_2 (Fodor, Garrett and Bever 1968)? In each case the answer was positive, but these answers are less interesting than the `background' data on differences between the test sentences. The relevant experiments all had the same basic form: a student volunteer had to paraphrase a sentence which was presented by the researcher, and the result was scored according to how accurate the paraphrase was, with one point for each of the subject-object relations in the sentence. In one experiment the sentences were presented in a number of different ways (spoken with neutral or

expressive intonation or written), but in each experiment each sentence was presented five times. The results that are reported are the means of all student-scores for each sentence in each experimental condition.

The following sentences are the ones used by Fodor, Garrett and Bever, together with scores based on the accuracy of each of the five occurrences of the sentence, giving a maximum of 15 per sentence. The figure shown is the mean of the scores for the two alternative verbs; for example, the first sentence was tried with saw and with used as V_2, producing scores of 6.5 and 6.7 respectively. Since we are not interested in this variation, I report the mean, 6.6. The sentences are listed in order of increasing success by the students.

(49)a The actors the writer the agent sent saw/used were talented. (6.6)

b The planes the sailors the enemy attacked evaded/feared were bombers. (8.0)

c The material the tailor the designer used cut/required was green. (9.8)

d The plan the lawyer the client interviewed proposed/devised was impractical. (10.1)

e The tiger the natives the hunter preferred/paid hated was fierce. (10.6)

f The letter the secretary the manager employed mailed/expected was late. (10.8)

g The box the man the child knew/met carried was empty. (11.1)

h The deer the man the boy followed fed/heard were timid. (11.1)

i The insult the waiter the lady summoned intended/provoked was obvious. (11.7)

j The results the scientist the committee appointed ignored/predicted

were surprising. (12.0)

k The events the papers the man bought discussed/reported were unsettling. (12.8)

l The tactics the general the soldiers admired suggested/adopted were stupid. (12.9)

The range of scores for sentences is much greater than the variation introduced by varying the verbs, but the authors do not comment on it.

How can we explain this variation? One possibility is that it has to do with the availability of semantic scripts or schemas linking the nouns. For example, having processed the first few words of sentence (l), we can start looking for a plausible semantic link between the meanings of N_1 and N_2 (tactics and a general) which involves looking for a knowledge-schema which mentions both tactics and generals (with the generals in an active, subject-compatible, role). In other words, it is easy to guess a verb to complete the sentence: `The general ___ the tactics'. The same is true of $N_2$ and $N_3$. In contrast, sentence (a) starts with two nouns which are less easily linked by a schema; no verb is an obvious filler for `The writer _____ the actors'. In the absence of an objective measure of semantic linkage this explanation is impossible to test, but it seems promising.

The other Fodor paper (Fodor and Garret 1968) allows us to continue this line of thought. It uses a list of eight sentences in which the presence or absence of relative pronouns is the variable under investigation. They presented the sentences in five different ways (spoken with neutral or expressive intonation, with pauses to compensate for the missing relative pronouns, in written form and with added adjectives), giving five experiments. (In this paper the scores reported show the number of

subject-object relations successfully identified, but this is divided by the time taken to reply, so the numbers are much smaller than in the earlier paper.) In most of the experiments the same sentences were presented (five times each), with and without relative pronouns, and the scores reported are of course those which contrast the presence of pronouns with their absence. We shall ignore this contrast by taking the mean of the two scores, and we shall ignore the (non-significant) effects of the different methods of presentation by summing the means for all five experiments.

These are the figures quoted in the following list, which again presents the sentences in increasing order of ease. However this time they fall into two clearly distinct groups (which were deliberately distinguished by the researchers for other reasons which we shall discuss below) and we shall start with the easier group.

(50)a The car (which) the man (whom) the dog bit drove crashed. (3.68)

b The shot (which) the soldier (that) the mosquito bit fired missed. (4.05)

c The window (which) the ball (that) the boy threw hit broke. (4.18)

In all these sentences N_1 and N_2 are very easily linked by a familiar schema: a man driving a car, a soldier firing a shot and a ball breaking a window. In fact there are few plausible alternatives for filling the verb slot in sentences like `The soldier _____ the shot' or `The ball _____ the window.' Having heard N_1 and N_2, therefore, the hearer can already start to build a semantic structure that links them. The same is true for the schemas linking N_2 and N_3: a dog biting a man, a mosquito biting a person and a boy throwing a ball.

Now consider the harder sentences:

(51)a The tiger (which) the lion (that) the gorilla chased killed was ferocious. (1.37)

b The man (whom) the girl (that) my friend married knew died. (1.47)

c The cigarette (which) the match (that) the flame ignited lit smoked. (1.91)

d The pen (which) the author (whom) the editor liked used was new. (1.92)

e The boats (which) the rocks (that) the waves covered sank were large. (1.99)

Here semantic links are much less helpful, and especially so in the two hardest sentences. Neither N_1 and N_2, nor N_2 and N_3, have obvious links. `The tiger ____ the lion', `The gorilla ___ the lion', `The girl ____ the man' and `My friend ____ the girl' have no one obvious filler. A multiplicity of possible fillers is no help to the processor, for whom the ideal is a single clear candidate.

The other sentences have somewhat higher scores, and seem to constitute a separate intermediate group. It is tempting to explain this in terms of linkages as well. For all of them there is a fairly obvious link for N_1 and N_2, but N_2 and N_3 are more of a problem. It is not at all obvious what a flame does to a match, since in the conventional schema it is the match that produces the flame; and it is equally hard to guess what the editor might have done to the author, or the waves to the rocks.

It seems reasonable to draw the following conclusion from the two Fodor papers: that when we hear N_1 and N_2 we guess that N_2 is the subject of a relative clause (especially if it is preceded by a relative pronoun), but we then have to guess at the meaning of the verb which will eventually link them together (V_2); and similarly for N_2 and N_3,

which generate a hypothetical V_3. For example, take 0a. After hearing N_2, the state of the parser is as follows, where the vertical bar separates the words processed so far from those which are simply anticipated:

(52)The car (which) the man | V_2 V_1

V_2 is the verb which will link man and car, and which the processor suspects means `drive'; and $V_1$ is the anticipated main verb (which will be finite, will have a subject, and so on). After $N_3$, the parser's state is as follows, where $V_3$ is expected to mean `bite'.

(53)The car (which) the man (whom) the dog | V_3 V_2 V_1

The next section will explore the reasons why it is so important to be able to guess the meaning of V_3 and V_2.

This discussion has followed the early literature in focussing exclusively on relative clauses, but we have seen that content clauses can cause problems too. It seems unlikely that these remarks about guessing verb-meanings apply equally when a content clause modifies N_1, because in that case there is no semantic link at all between N_1 and N_2. For example, take a simplified version of one of Gibson's examples:

(54)The information [that the frightened child had survived the crash]
comforted the woman.

The semantic link between the modifying clause and N_1, information, involves the total meaning of the clause (the proposition that the frightened child had survived the crash) and not the kind of direct link found in relative clauses which would allow us to look for a plausible proposition `the frightened child _____ the information'. This makes it impossible to predict the meaning of the anticipated content clause's verb, V_1. On the other hand, the next verb, V_2, may be relatively easy to predict on the basis of a noun such as information, given the small

number of verbs that allow such a noun as subject.

6. Doubly empty concepts?

The research evidence that we have surveyed has led us to the following conclusion about so-called multiple self-embeddings:

(55)As hearers or readers we cannot handle a sentence in which:

i. a finite clause [..N_2..V_2..] modifies a noun N_1,

ii. N_2 is modified by another finite clause [..N_3..V_3..],

iii. N_3 is a common noun,

iv. it is hard to guess the meanings of V_2 and V_3 on hearing N_1, N_2 and N_3.

This definition applies to all the clear examples of sentences which are declared `impossible to process', but allows some variation among these examples according to how true clause iv. is; so any sentence which satisfies i. to iii. is guaranteed to be hard to process, but some are harder than others.

The most general conclusion of this discussion is that it is because of their meaning that these sentences are hard to process, but this definition is expressed mainly in terms of syntactic structures. This cannot be a helpful basis for understanding the problem, so we need to rephrase the definition in semantic terms. The following attempt rests on the uncontroversial assumption that the sense of a phrase like big boy is built compositionally by using the meaning of the dependent, `big', to modify that of the head, `boy'. Somewhat more precisely, if `boy' is the name of the concept which is the sense of boy, then we can use `big boy' as the name of the sense of big boy. Similarly, the sense of boy with long hair is `boy with long hair', and that of boy who lives in our

street is `boy who lives in our street'. (These prose concept-names are meant to be theory-neutral, though they are also the kinds of name which would be used in semantic analyses of the kind I favour, as explained in Hudson 1984, 1990, 1995.) The essential point, however, is not what we should call these concepts, but rather that each has a content which has to be built out of the content of the head word plus its modifiers. It is only the content of the modifier that distinguishes the concept `boy' from `big boy' and from `boy who lives in our street' and so long as the modifier's meaning is unknown, the concept attached to `boy who ..' is hard to distinguish from `boy'. I shall suggest below that this is the main source of difficulty in processing self-embedded sentences.

In terms of processing, the hearer's task is to construct a set of semantic concepts on the basis of words and their modifiers which include the following:

- lexical senses, i.e. the senses that words are given in the lexicon (e.g. `boy' for BOY),
- phrase-senses, built out of the head word's lexical senses and the meanings (typically referents) of its modifiers, and
- referents, the particular (or `generic') individuals who are picked out with or without the help of senses.

For example, if I refer to John Brown as `a boy who lives in our street', the referent of this phrase is John Brown, its phrase-sense (to which the determiner makes no contribution) is `boy who lives in our street', and the lexical sense of boy is `boy'. For each noun in our problem sentences, therefore, the hearer must construct a multiple semantic analysis containing three different concepts: its lexical sense, the phrase-sense of

the phrase that it heads, and its referent. These concepts must be given different names (or `addresses'), but concepts that have similar contents are in danger of collapsing into a single concept - which, I shall suggest, is what happens when we try to process some self-embedded sentences.

When seen from the point of view of an incremental parser, such structures are very frustrating. Suppose I am trying to process the original problem sentence The rat the cat the dog chased ate died. I pass through the following stages:

1. I hear N_1, the rat, for which I find the lexical sense `rat'. I also postulate some particular rat as its referent. So far so good.
2. I hear N_2, the cat, and recognise the start of a relative clause. (This may be only one of several syntactic structures that I consider, but such syntactic ambiguity as there is can be handled quite easily if this is the only relative clause.) This relative clause modifies the sense of N_1, so its sense is not just `rat', but `rat ...' - a particular kind of rat. But what kind? We know that it has something to do with a cat, and that the cat is the syntactic subject of the relative clause, but that's all; we know neither the identity of this cat, nor the situational link between it and the rat. So the best we can do as a definition of the rat-type is to call it:

`rat (cat R_1 it)'

(i.e. `a rat such that some cat has some relation $R_1$ to it', or more colloquially, `a rat that has something to do with some cat') - a definition whose abstractness and lack of real-world links or imagery makes it hard to keep separate from the simpler concept `rat'. For simplicity we can call it an `empty' concepts, although later words will of course `fill it up' by supplying a referent for the cat and a relationship between the cat and the rat.

3. I hear N_3, the dog, introducing another relative clause which doubles the problems just mentioned. I now have to cope with a `doubly-empty' concept - `a rat that has something to do with some cat which has something to do with some dog'. We still have no idea what relation R_1 is, but we now have another unknown relation, R_2. What we know about the cat is that it is:

`cat (dog R_2 it)',

but this uncertainty is built into the meagre description of the rat:

`rat (cat (dog R_2 it) R_1 it)'.

If I can hang onto this structure while processing the next few words, I can replace the variables by concepts with some content; but maybe this much uncertainty in the definition of one concept is more than my mind can cope with, so the structure falls to pieces before I can firm it up. This is the idea that I shall develop in the next section.

7. Towards a theoretical explanation.

My suggestion is that the real source of processing problems is this semantic structure, and not the syntactic structure. But why does it cause problems? Unfortunately this is a question for a psychologist rather than for a linguist. However a linguist may be allowed to speculate. My guess is that the explanation has something to do with the difficulty of keeping the various concepts distinct in memory. Psychological experiments have shown that lists of items are easier to hold in working memory if they are dissimilar:

(56)The span [of immediate memory for a list of words] is reduced if the items of the sequence sound similar. For example, the sequence B V T G P is more difficult than the sequence S K L R N. (Brown

1987b)

(57) An event is difficult to remember if it is one of a number of similar events .. The difficulty probably arises partly because insufficient information has been stored to enable differentiation of one event from another .. (Brown 1987a)

In unprocessable sentences there are three groups of overly similar concepts.

a. First, there are the phrase-senses of N_1, N_2 and N_3, in relation to their respective lexical senses. The idea of a rat which has something or other to do with a cat is too similar to that of a rat; and similarly for the other nouns. We can clearly cope easily with a single `open' phrase-sense which for a few words remains dangerously similar to the relevant lexical sense. This is the case whenever we handle an object-relative clause, however simple it may be: after hearing just the book my students, I know there is a phrase-sense which links `book' to the students, but I have no more idea of what it might be. A sentence with two open phrase-senses is hard, but possible for most of us; an example is our `easy' sentence 0: A book that some Italian I've never heard of wrote will be published soon by MIT Press. In processing this sentence, there is a point just after I where we are holding a phrase-sense for `book' and also one for `Italian'. However this state of uncertainty pushes our capacity to the limit, and we can cope only if the third noun, in this case I, by-passes the sense level altogether. This, then, is the reason why sentences are unprocessable only if N_3 has a sense.

It may be possible to offer a similar explanation for the interesting asymmetry between relative clauses and content clauses which Gibson and colleagues have discovered. As noted earlier in connection with

examples 0a and 0, a relative clause inside a content clause is easier to process than might be expected, in contrast with a content clause inside a relative clause:

0aThe old woman [who the information [that the frightened child survived the crash] had comforted] looked for the rescue worker.

0The information [that the frightened child [who the rescue worker looked for] had survived the crash] comforted the woman.

One possible reason why 0 is so easy is that the semantic relation between N_1 and the modifying clause involves co-reference, i.e. apposition, rather than a refined phrase-sense. To defend this suggestion would take us beyond the scope of the present paper (since it would involve, among other things, defending the view that clauses have a referent, as I have argued in the works cited earlier), but it would relate to the fact that only one determiner, the, is possible before information that ...., in contrast with the relative clause examples. The lack of problems in N_1 leaves plenty of spare capacity for coping with N_2. In contrast, the content clause inside the relative clause of 0a causes problems by its very length and complexity, at a time when the processing system is already working at full capacity because of the uncertainties at N_1.

b. The second set of concepts that are too similar to each other for comfortable processing are the senses of the postulated verbs V_3, V_2 and V_1. The more meaning these have, the easier they are to keep distinct. We have reviewed evidence (produced by Schlesinger and by Fodor and colleagues) which showed that it is easier to handle a sentence if we can guess the meanings of these two verbs. In some respects this

interpretation of their data is similar to the interpretation that they themselves offer. Influenced by Schlesinger (1966), Fodor, Garrett and Bever recognise an important role for meaning in the processing of self-embedded sentences, and indeed it was because of this that they included the three sentences listed in 0 which turned out to be so much easier than the rest. However we have very different ideas about the nature of the role played by meaning. For Fodor et al, its role is to eliminate uncertainty about how to pair up the verbs and nouns retrospectively; for example, since men don't bite cars, we don't try to use bit to link the man to the car in example (a). In contrast, the harder sentences are hard because all pairings of verbs and nouns make sense. In this view, meaning plays no part until the end of the sentence, when all the nouns and verbs have been read in and the processor tries to sort out their relationships. In the view that I am offering, meaning plays an important part from the earliest stage of an incremental parse, by allowing the processor both to anticipate future linking words and to keep the anticipated links active while waiting for the expected words.

Is there any evidence for the second part of my claim, which is that without distinct meanings the verbs tend to merge in memory? According to Gibson (1991:169), `The intuition shared by most people is that the point of breakdown in doubly center-embedded structures occurs somewhere after the first verb has been input, either in the second or third verb of that structure.' As reported by Frazier (1985:178), people sometimes behave as though the anticipated verbs had merged in their memories: they accept as well-formed self-embedded sentences in which one of the last two verbs is omitted, such as the next example:
(58)*The patient the nurse the clinic had hired met Jack.

People who accept such sentences must have lost a verb from their list of anticipated words. The easiest assumption is that the lost verb is V_2, whose meaning cannot be guessed from the sequence The patient the nurse and has therefore merged with V_1.

c. The third dangerous set of concepts, where too much similarity can cause confusion, contains the senses of N_1, N_2 and N_3 themselves. It is noticeable, for example, that the two hardest sentences that Fodor and Garret (1968) quoted are the only two in which all three nouns have similar referents:

0a The tiger (which) the lion (that) the gorilla chased killed was ferocious.

b The man (whom) the girl (that) my friend married knew died.

In the first example they all refer to animals, and in the second to people. Even more clearly, when the modifying clause is a content clause, we noted that reactions varied considerably according to the similarity between N_1 and N_2. In discussing examples 0, 0 and 0 we noticed that The question whether the fact that he sniffs .. was quite easy, in contrast with both The fact that the fact that he snores .. and The idea that the fact he lied ..

In short, the rather meagre factual evidence which is available to us is at least compatible with the idea that processing difficulties are caused by the need to divert resources into preventing very similar concepts from merging with each other in semantic structure, when these resources are also needed for dealing with relatively (but not inordinately) complex syntactic structures. In other words, the difficulties are due to `interference' at the semantic level, comparable with the syntactic interference suggested by Lewis 1996.

If the speculations in this section have any substance then it may be possible to interpret them in terms of some general theory of working memory. However, even if they are completely wrong, at least we may have a clearer idea of what the phenomenon is that we are trying to explain. We have abandoned the vague terms `center-embedding' and `self-embedding', since (however we interpret them) they don't fit the psycholinguistic data reported in the literature; and we have replaced them by a definition in which semantic structure plays an important role.

References

Babyonyshev, Maria. and Gibson, Edward: 1995. Processing overload in Japanese. MIT Working Papers in Linguistics.

Brown, John: 1987a. Forgetting. In Gregory, R. (ed.) The Oxford Companion to the Mind. Oxford University Press, 264-5.

Brown, John: 1987a. Short-term memory. In Gregory, R. (ed.) The Oxford Companion to the Mind. Oxford University Press, 713-4.

Carpenter, Patricia; Miyake, Akira and Just, Marcel: 1994. Working memory constraints in comprehension. Evidence from individual differences, aphasia and aging. In Gernsbacher, M. (ed.) Handbook of Psycholinguistics. New York: Academic Press, 1075-122.

Chomsky, Noam and Miller, George: 1963. Introduction to the formal

analysis of natural languages. In Luce, R.; Bush, R. and Galanter, E. (eds.) Handbook of Mathematical Psychology, Vol 2. New York: Wiley, 269-323.

Crystal, David: 1980/1991. A Dictionary of Linguistics and Phonetics. 3rd Edition. Oxford: Blackwell.

de Roeck, Anne; Johnson, Roderick; King, Margaret; Rosner, Michael; Sampson, Geoffrey and Varile, Nino: 1982. A myth about center-embedding. Lingua 58, 327-40.

Fodor, Jerold and Garrett, Marrill: 1967. Some syntactic determinants of sentential complexity. Perception and Psychophysics 2, 289-96.

Fodor, Jerrold; Garrett, Merrill and Bever, Thomas: 1968. Some syntactic determinants of sentential complexity II: verb structure. Perception and Psychophysics 3, 453-61.

Frank, Robert 1992. Syntactic Locality and Tree Adjoining Grammar: Grammatical, Acquisition and Processing Perspectives. University of Pennsylvania PhD dissertation.

Frazier, Lyn 1985. Syntactic complexity. In D. Dowty; L. Karttunen and Zwicky, A. (eds.) Natural Language Processing: Psychological, Computational and Theoretical Perspectives. Cambridge: Cambridge University Press.

Gibson, Edward 1991. A Computational Model of Human Linguistic Processing: Memory Limitations and Processing Breakdown. Carnegie Mellon University PhD dissertation.

Gibson, Edward; Thomas, James.; Babyonyshev, Maria: 1995. Processing center-embedded and self-embedded structures in English and Japanese. NELS conference handout.

Hagège, Claude: 1976. Relative clause center-embedding and comprehensibility. Linguistic Inquiry 7, 198-201.

Hudson, Richard: 1984. Word Grammar. Oxford: Blackwell.

Hudson, Richard: 1990. English Word Grammar. Oxford: Blackwell.

Hudson, Richard: 1995. Word Meaning. London: Routledge.

Kac, Michael: 1981. Center-embedding revisited. In Proceedings of the Third Annual Conference of the Cognitive Science Society, Hillsdale: Lawrence Erlbaum, 123-4.

Kimball, John: 1973. Seven principles of surface structure parsing in natural language. Cognition 2, 15-47.

Larkin, Phillip: 1947/1975. A Girl in Winter. London: Faber and Faber.

Lewis, Richard: 1996. A theory of grammatical but unacceptable

embeddings. mimeo.

Miller, George and Chomsky, Noam: 1963. Finitary models of language users. In Luce, R.; Bush, R. and Galanter, E. (eds.) Handbook of Mathematical Psychology, Vol 2. New York: Wiley, 419-93.

Miller, George and Isard, Stephen: 1964. Free recall of self-embedded English sentences. Information and Control 7, 292-303.

Newmeyer, Frederick: 1988. Extensions and implications of linguistic theory: an overview. In F. Newmeyer, (ed.) Linguistics: The Cambridge Survey 2. Linguistic Theory: Extensions and Implications. Cambridge: Cambridge University Press, 1-14.

Niv, Michael: 1993. Resolution of syntactic ambiguity: the case of new subjects. Proceedings of the 15th Annual Meeting of the Cognitive Science Society. Boulder, CO.

Pinker, Stephen: 1994. The Language Instinct. Harmondsworth: Penguin.

Sampson, Geoffrey: 1996. From central embedding to corpus linguistics. In Jenifer Thomas and Michael Short (eds.) Using Corpora for Language Research: Studies in the Honourr of Geoffrey Leech. Longman, London, 14-26.

Schlesinger, Izchak: 1966. Sentence Structure and the Reading Process.

Hebrew University of Jerusalem PhD dissertation.

Schlesinger, Izchak: 1968. Sentence Structure and the Reading Process. The Hague: Mouton.

Smith, Neil: 1989. The Twitter Machine. Reflections on language. Oxford: Blackwell.

Stromswold, Karin; Caplan, David; Alpert, Nathaniel and Rausch, Scott: Forthcoming. Localization of syntactic comprehension by Positron Emission Tomography. Brain and Language.

Trask, Larry: 1993. A Dictionary of Grammatical Terms in Linguistics. London: Routledge.

Appendix

Some of the data quoted in this paper are taken from a series of small-scale experiments which I conducted in the autumn of 1995, using students as subjects. In each experiment, I read a list of sentences to a group of between 15 and 30 students, whose task was to write each sentence down as soon as I had read it out. They heard it only once, so the experiment tested their ability to hold it in short-term memory. For some experiments I supplied the first word, so as to be able to use slightly longer (and more complicated) sentences; but for each experiment the length of the sentences (measured in words) was

constant, as was the number of words which I supplied. The dependent variable was the students' success in repeating the sentences, which I assume to be related to their success in understanding it. (Since conducting the experiment I have read Miller and Isard 1963, which describes a somewhat similar experiment in which 24 subjects tried to memorise six tape-recorded 22-word sentences with various degrees of self-embedding. They found that performance correlated very precisely with the degree of self-embedding.)

Aware of the possibility of rote recall without understanding (which worried Fodor and his colleagues), I designed these experiments to test the difference between students' ability to process grammatically structured sentences and random strings. Students who failed on random strings but succeeded on ordinary sentences of the same length must, presumably, have succeeded by understanding them. It turned out that ten words was the limit for pure rote recall (only one student managed to recall a ten-word random string, and even she was beaten by eleven words; Miller and Isard also included a random string, of which their subjects could only recall on average about eleven words even after five hearings), so I reasoned that anyone who could recall a ten-word sentence must have understood it (at least at the level of syntax). Conversely, everyone could recall easy ten-word sentences, so failure on ten words must be evidence for lack of understanding. By supplying words on the students' answer-sheet I could reduce the number of words that they had to remember without reducing the length of the sentences (an important consideration when testing complex sentences).

Like Miller and Isard, I used an expressive intonation pattern for each

sentence which gave the maximum of help with potentially difficult sentences. Clearly any `expressive' intonation is arbitrary when applied to a random list of words, but it may have helped people to `chunk' them, which may explain why they could recall more than the predicted number of words (`seven plus or minus two').

The scoring system was simple: for each sentence, each student either succeeded or failed. A reply was a success if:

- it was a verbatim copy of the original, or
- differences could be attributed to mishearing, or
- differences were meaning-preserving.

(Meaning-preserving differences are firm evidence of successful processing, so it is important to distinguish them from other differences which might suggest difficulty in processing.) Any other response was counted as a fail, but the method could be made more sophisticated by recognising different degrees of failure. The score for each sentence is the number of fails as a percentage of the number of students who took part in the experiment. (I excluded a small number of students whose poor English prevented them from understanding the isolated sentences; most of the remaining students were native speakers.)

The following are the sentences that were used in the experiments, listed in order of increasing failure and difficulty. The number at the beginning of each sentence indicates its position among the other sentences that I read out. A final star (*) marks sentences that contain one clause centre-embedded in another embedded one. Sentences where both of the embedded clauses are relative clauses, as in the classic `unprocessable' examples, have a double star `**'. Random strings of words have `!!'.

Experiment 1 (N = 27; first word supplied, nine to be recalled).

Sentence	% Fails
1. She saw a grey squirrel that was digging for nuts.	0
2. Athletes use the technique of creative imagery to improve their performance.	0
7. The fact that the film you chose was boring didn't matter. *	4
6. The shells that the children we watched collected were piled up. **	7
4. Nobody wants a car that people they dislike gave away. *	11
9. If although you try hard you still don't understand tell me. *	11
8. Sentences that students who I teach memorise can be quite short. **	15
5. The idea that the fact he lied was obvious worried me. **	30
10. People that politicians who journalists interview criticise can't defend themselves well. **	81
3. She windy under a real of with love neither much banana. !!	96

Experiment 2 (N = 36 except sentence 10; first word supplied, ten to be recalled).

Sentence	% Fails
1. It surprises students that linguistics dates back to the Ancient Greeks.	0
3. People that things I say annoy tend to tell me off. **	3
9. People that have good ideas about running the economy successfully compete. *	11
6. That footballers earn fabulous salaries for playing games surpises some people. *	14
4. People were standing around the enormous fire outside watching the fireworks.	17
5. People that footballers hired by managers attack get plenty of excitement. **	17
8. Although when people drink vodka they get headaches some people forget. *	31
2. The fact that the fact that he snores surprised us is interesting. **	36
10. People that criminals who the police recognise attack claim legal compensation. **	68 (N = 25)
10. People that muggers who the police recognise attack claim legal compensation. **	27 (N = 11)
7. People that politicians who journalists interview	75

criticise have no legal protection. **

Experiment 3 (N = 14; no words supplied, eleven to be recalled).

Sentences	% Fails
1. Students complain when they have to write essays about syntactic theory.	0
6. The question whether the fact that he sniffs annoys us is interesting. **	7
7. The apples that the girls Bill watched collected were heaped up. **	14
2. Everything the people who came wanted was waiting in their rooms. **	21
4. People that things which I do upset always complain about the effects. **	21
10. Everything that what Mary earned bought is in that brown suitcase. **	21
5. The farmer that the cow that gave bad milk kicked died. **	29
8. Articles have been written which criticise recent theories about social change.	50
9. Beer students policemen follow buy comes from very many different places. ** [Note: 7/12 fails just omitted very.]	86
3. Much try side more the whom three months knows felt	100

was. !!