Reading relative clauses in English

EDWARD GIBSON, TIMOTHY DESMET, DANIEL GRODNER,
DUANE WATSON and KARA KO*

Abstract

Two self-paced reading experiments investigated several factors that influence the comprehension complexity of singly-embedded relative clauses (RCs) in English. Three factors were manipulated in Experiment 1, resulting in three main effects. First, object-extracted RCs were read more slowly than subject-extracted RCs, replicating previous work. Second, RCs that were embedded within the sentential complement of a noun were read more slowly than comparable RCs that were not embedded in this way. Third, and most interestingly, object-modifying RCs were read more slowly than subject-modifying relative clauses. This result contradicts one of the central tenets of complexity research: that nested sentences are harder to understand than their right-branching equivalents (e.g., Miller and Chomsky 1963). It is hypothesized that this result followed from a combination of two information-flow factors: (1) background information is usually presented early in a sentence; and (2) restrictive RCs—the form of the RCs in Experiment 1—usually convey background information. Experiment 2 tested this hypothesis by comparing restrictive and non-restrictive RCs—which generally provide new information—in both subject- and object-modifying positions. The results of the experiment were as predicted by the information-flow account: Only restrictive RCs were read more slowly when modifying objects. It is concluded that both resource and information-flow factors need to be considered in explaining RC complexity effects.

Keywords:

1. Introduction

During the last four decades, the processing of relative clauses (RCs) has played a prominent role, both in linguistic and psycholinguistic research. One reason for this interest is that RCs represent a type of symbolic

Cognitive Linguistics 16–2 (2005), 313–353 0936–5907/05/0016–0313

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

recursion, one of the most distinctive properties of natural language as a cognitive system. Recursion is the ability to embed one instance of a category inside another instance of that category, and permits the generation of an infinite number of structures. In an RC, a sentence is embedded within another sentence, as in (1).

(1) The scientist collaborated with the professor who advised the student.

Here the sentence *The professor advised the student* is embedded within the sentence *The scientist collaborated with the professor*.

Within the processing literature, *center-embedded* or *nested* structures, a specific case of recursive structures, have received considerable attention. Center-embedding is a formal property of language that necessitates the existence of a memory structure (e.g., a stack) in addition to a finite state automaton (Chomsky 1959; Chomsky and Miller 1963). As a result, center-embedded structures are more difficult to understand than their right-branching counterparts (Chomsky 1957, 1965; Chomsky and Miller 1963; Miller and Isard 1964; Yngve 1960). A syntactic structure A is said to be center-embedded or nested within a structure B if B contains A, such that there is at least one constituent of B to the left and to the right of A. For example in (2a), the RC *who the scientist collaborated with* is nested within the RC *who the professor ... advised*, which is itself nested within the top-level sentence *the student ... copied the article*:

(2) a. The student who the professor who the scientist collaborated with advised copied the article.
    b. The scientist collaborated with the professor who advised the student who copied the article.

The resulting doubly-nested structure in (2a) is much harder to understand than (2b), its right-branching counterpart, containing the same words in the same thematic relations. Although nested sentences are grammatical, increasing the number of nestings soon makes a sentence hard or even impossible to process. This finding has been replicated using a number of different paradigms (e.g., Blaubergs and Braine 1974; Blumenthal 1966; Foss and Lynch 1969; Hakes and Cairns 1970; Miller and Isard 1964; Stolz 1967).

Because nested sentences and their right-branching variants are made up of the same words and have the same meaning, lexical or contextual information cannot explain the complexity differences between them. Researchers have thus proposed that the complexity difference between the two kinds of structures is caused by a difference in the amount of computational resources needed to process them. Miller (see e.g., Miller and

Chomsky 1963; Miller and Isard 1964) noted that in a nested sentence, each RC interrupts the clause in which it is embedded (see [2a]). The same is not true in a right-branching sentence (see [2b]). Miller argued that language comprehension involves a limited capacity short-term processing buffer that holds parts of clauses partially analyzed until their completions are available. As a result, more than a few nestings result in comprehension difficulty or failure.

Miller's so-called 'interruption hypothesis' has been very influential. Numerous theoretical accounts hypothesize that one factor contributing to sentence complexity is the number of partially-processed phrase structure rules or, more generally, the number of incomplete syntactic or thematic dependencies that the parser has to store in memory at a particular parse state, with the goal of forming a grammatical sentence (Kimball 1973; Hakuta 1981; MacWhinney 1987; Gibson 1991, 1998; Pickering and Barry 1991; Lewis 1996; Stabler 1994; Yngve 1960; Chomsky and Miller 1963; Miller and Chomsky 1963; Miller and Isard 1964; Abney and Johnson 1991). We will refer to such accounts as *storage* accounts of nesting complexity. One particular storage account is phrased in terms of the minimal number of predicted syntactic heads that are required to form a grammatical sentence at each parser state (Gibson 1998, 2000). The contrast between a nested structure like (2a) and its right-branching control (2b) is accounted for by this storage account as follows. The point where all theories suggest that the maximal storage load occurs in (2a) is at the point of processing the noun phrase *the scientist*. At this point, there are five predicted syntactic heads, consisting of three predicted verbs for each of the subject NPs (e.g., *copied*, *collaborated* and *advised* in [2a]), and two empty NP positions to be associated with the two RC-pronouns. In contrast, the maximal storage cost at any point in processing the right-branching sentence in (2b) is only one predicted syntactic head. For instance, at the first relative pronoun *who*, only a verb is needed to form a grammatical sentence if the RC pronoun is taken to be the subject of the RC.

In addition to storage costs, other factors have been proposed to affect the processing complexity of embedded structures. These factors include the following (see Gibson 1998, for a recent summary of some relevant factors):

i. Integration distances between dependents that need to be connected together, as proposed in Gibson's (1998, 2000) dependency locality theory (DLT) (cf. Hawkins 1994). For example, although syntactic storage costs differ in (2a) vs. (2b), integration distances also differ in this comparison. In particular, the verbal dependents are linearly farther apart in the nested version (2a) than those in the non-nested version (2b). For example, the verb *advised* is between its arguments *who* and *the student* in (2b), but the same verb is very far from each of these arguments in the nested version (2a). Gibson (1998, 2000) and Grodner and Gibson (in press) provide evidence that longer distance dependencies lead to longer reading times at the right-hand end of dependencies. Some of this evidence comes from the comparison between object-extracted and subject-extracted RCs. In an object-extracted RC like (3a) below, the wh-pronoun is associated with the object position of the verb in the RC, whereas in a subject-extracted RC like (3b) below, the wh-pronoun is associated with the subject of the verb in the RC:

(3) a. The reporter who the senator attacked admitted the error.
b. The reporter who attacked the senator admitted the error.

Much research using a range of methods and tasks has demonstrated that object-extracted RCs are more complex than subject-extracted RCs (Ford 1983; Hakes, Evans and Brannon 1976; Holmes and O'Regan 1981; Just, Carpenter and Keller 1996; King and Just 1991; Stromswold, Caplan, Alpert, and Rauch 1996; Wanner and Maratsos 1978; Waters, Caplan and Hildebrandt 1987). A possible cause for the observed complexity difference is the longer distance connections in object-extracted RCs compared to subject-extracted RCs. In (2a), the RCs are object-extracted, and hence have longer distance dependencies, than the RCs in (2b) which are subject-extracted. Thus part of the reason that the nested version (2a) is more complex than the non-nested version (2b) is probably because of the longer integrations in (2a).

ii. Perspective shift (MacWhinney 1977, 1982; MacWhinney and Pleh 1988; cf. Bever 1970). Under this theory, processing resources are required to shift the perspective of a clause, where the perspective of a clause is taken from the subject of the clause. This theory does not explain the difference between the nested and non-nested versions of (2), but it does offer a potential account of the complexity difference between subject- and object-extractions in (3). Processing the object-extracted RC structure in (3a) requires two perspective shifts: (a) from the perspective of the matrix subject to the subject of the RC and (b) from the perspective of the subject of the RC back to the matrix subject, after the RC is processed. Processing the subject-extracted RC in (3b) requires no perspective shifts, because the matrix subject is also the subject of the RC, so that both clauses come from the same perspective. Thus the object-extraction is more complex than the subject extraction. Recent evidence from the processing of Chinese relative clauses suggest that this theory does not apply in processing Chinese (Hsiao and Gibson 2003), but it may still apply in English.

iii. Differences in canonical vs. non-canonical word order (e.g., MacDonald and Christiansen 2002; cf. Bever 1970; Tabor, Juliano and Tanenhaus 1997; Mitchell, Cuetos, Corley, and Brysbaert 1995). The word order in English is Subject-Verb-Object (SVO). This word order is present in the right-branching subject-extracted RCs in (2b) (e.g., *who advised the student*, *who copied the article*), but not in the nested object-extracted RCs in (2a). Similarly, SVO word order is present in the subject-extracted RC in (3b) (*who attacked the senator*), but the word order in the object-extracted RC in (3a) is non-canonical: OSV (*who the senator attacked*).

Recent research performed by Gibson and colleagues has demonstrated effects of on-line storage independent of the other proposed complexity factors from the literature. For example, Chen, Gibson and Wolf (2003) showed that having more predicted verbs slows reading. In particular, Chen and colleagues showed that the underlined region in (4) is read increasingly slowly across (4a), (4b) and (4c):

(4) a. The employee realized that the boss implied that *the company planned a layoff* and so he sought alternative employment.
b. The employee realized that the implication that *the company planned a layoff* was not just a rumor.
c. The realization that the implication that *the company planned a layoff* was not just a rumor caused a panic.

The critical region *the company planned a layoff* is identical in all conditions, with the consequence that integration costs are the same across the three. In addition, the word order is canonical in all three sentences during the critical region, and there are the same number of perspective shifts in each sentence at the point of processing the critical region. In sentence (4a), the critical region is embedded as the sentential complement of the verb *implied* which is itself part of a clause embedded as the sentential complement of the matrix verb *realized*. Because both verbs *implied* and *realized* are encountered immediately after their respective subject nouns, no additional predicted verbs need to be stored across the critical embedded clause. In sentence (4b), the verb *implied* is nominalized to *implication* with the result that the critical clause is a sentential complement of the noun *implication*. This change to the embedded subject noun phrase *the implication* results in the requirement for an additional verb during the processing of the critical region. Finally, in sentence (4c), both the verbs *realized* and *implied* are nominalized with the result that predictions for two additional verbs must be maintained across the critical region. Thus the storage hypothesis predicts that RTs during the bold region should be slowest in (4c), faster in (4b), and fastest in (4a). These predictions were ratified by two self-paced reading experiments on similar items. Furthermore, Chen and colleagues provide evidence from two other English constructions that demonstrate the existence of storage costs independent of other factors. In addition, Gibson (1998, 2000), Grodner, Gibson and Tunstall (2002) and Gibson and Tunstall (1999) provide evidence from the resolution of ambiguity that syntactic storage costs are utilized independent of integration costs and other factors in the resolution of ambiguity.

Although there is an increasing quantity of evidence for the use of syntactic storage costs in on-line sentence comprehension, one strong prediction of the existence of such costs has failed to be ratified in past experimental investigations: a predicted difference between subject-modifying RCs and object-modifying RCs, as in (5):

(5) a. The reporter *that the senator attacked* ignored the president.
b. The president ignored the reporter *that the senator attacked*.

While processing a subject-modifying RC as in (5a), a verb is still needed to complete the matrix subject-verb dependency. In contrast, there is no such verbal expectation while processing an object-modifying RC as in (5b), because the matrix predicate has already been encountered at that point. The subject-modifying RC therefore requires more storage during its processing.

A number of early studies (e.g., Marks 1968; Blaubergs and Braine 1974) purported to find evidence that subject-modifying RCs are more complex than object-modifying RCs, but these studies confounded modifier position (subject, object) with the type of extraction in the RC: object- or subject-extracted. In these studies, the subject-modifying RCs in the materials were also object-extracted, and the object-modifying RCs were subject-extracted. Hence, the difficulty attributed to subject-modifier position may well have been due to the fact that the RCs in this position were object-extracted.

We know of four studies that directly compared subject- and object-modifying RCs while controlling for extraction type. First, in a sentence recall task, Holmes (1973) found that experimental participants were able to recall a greater number of words from subject-modifying RCs than object-modifying RCs. This result runs directly counter to the prediction of the syntactic storage hypothesis, but because the experiment used an off-line task, it is not clear which components of the sentences caused the purported difference in complexity. Furthermore, the materials in this early experiment did not controlled for a number of factors that we now know affect on-line sentence interpretation, such as plausibility (e.g., Trueswell, Tanenhaus and Garnsey 1994; see Gibson and Pearlmutter 1998 for a review). Second, Baird and Koslick (1974) found no differences between subject-modifying and object-modifying RCs using a fill-in-the-blank questionnaire following auditory presentation of the sentences. At the same time, they found a reliable effect of RC extraction-type, such that object-extracted RCs were more complex than subject-extracted RCs. Third, in the first on-line investigation of this comparison, Hakes et al. (1976) investigated the processing of RCs using a phoneme-monitoring task, and reported results similar to Baird and Koslick's. In particular, Hakes et al. found that object-extracted RCs were more complex than subject-extracted RCs, but they found no significant difference between subject- and object-modifiers. Finally, Gibson and Thomas (1996) studied complex versions of subject- and object-modifying RC sentences using a questionnaire in which sentences were rated according to their intuitive complexity. Like two of the three previous studies, Gibson and Thomas found no difference between subject- and object-modifying RCs, although they found evidence of numerous other complexity effects in comparisons among other conditions.

In summary, the evidence from previous work investigating comparisons between subject- and object-modifying RC structures is equivocal. An early study by Holmes demonstrated an advantage for the subject-modifying structure (contrary to the storage cost hypothesis), but this experiment used an off-line task in less than perfectly controlled materials. Furthermore, the result was not replicated in later studies, using either on-line or off-line methods. The goal of Experiment 1 was to test the storage cost hypothesis in subject- and object-modifying materials using an on-line task, in more rigorously controlled items than had been used in previous on-line studies.

2. Experiment 1

Three factors were crossed in the materials for Experiment 1, resulting in a $2 \times 2 \times 2$ design:

RC modifier position (subject-modifier, object-modifier), RC extraction-type (subject-extraction, object-extraction), and embedding (not embedded, embedded). An example of the eight versions of an item is given in (6).

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

(6) a. Subject modifier, object-extracted (SO), not embedded
The reporter *who the senator attacked* on Tuesday ignored the president.

b. Object modifier, object-extracted (OO), not embedded
The president ignored the reporter *who the senator attacked* on Tuesday.

c. Subject modifier, subject-extracted (SS), not embedded
The reporter *who attacked the senator* on Tuesday ignored the president.

d. Object modifier, subject-extracted (OS), not embedded
The president ignored the reporter *who attacked the senator* on Tuesday.

e. Subject modifier, object-extracted (SO), embedded
The fact that the reporter *who the senator attacked* on Tuesday ignored the president bothered the editor.

f. Object modifier, object-extracted (OO), embedded
The fact that the president ignored the reporter *who the senator attacked* on Tuesday bothered the editor.

g. Subject modifier, subject-extracted (SS), embedded
The fact that the reporter *who attacked the senator* on Tuesday ignored the president bothered the editor.

h. Object modifier, subject-extracted (OS), embedded
The fact that the president ignored the reporter *who attacked the senator* on Tuesday bothered the editor.

The critical manipulation involved the RC modifier position. We concentrate our predictions on the processing of the RC itself, in bold in (6). The storage hypothesis predicts that object-modifiers should be easier to process than subject-modifiers. The second factor, RC extraction-type, was included to ensure that the task was sensitive enough to detect complexity differences that are well documented in the literature. Thus we expected to observe a benefit for subject-extracted items compared to object-extracted items, possibly due to the difference in integration cost between subject- and object-extractions. The third factor—embedding—was included as a control to test the storage hypothesis. Chen and colleagues (2003) found that additional predicted verbs slow processing of embedded clauses. Thus we expected to find that the embedded versions of the RCs should be processed more slowly than the non-embedded versions.

Let us now consider the predictions of the other processing factors that were discussed above. First, consider perspective-shift theory with respect to the non-embedded conditions (6a–d). This theory predicts the least difficulty in processing the RC in (6c) the subject-modifying subject-extracted RC, in which the perspective is unchanged from that of the matrix subject *the reporter*. Perspective-shift theory predicts greater difficulty with the subject-modifying object-extracted RC (6a) during the RC because the perspective is shifted from that of the matrix subject *the reporter* to that of the embedded subject *the senator*. For each of the object-modifying RCs (6b) and (6d), there is one shift in perspective from the matrix subject *the president* to *the senator* for the object-extraction (6b), and to *the reporter* for the subject-extraction (6d). Thus, perspective-shift theory predicts an interaction between modifier position and extraction type, such that subject-extracted RCs should be easier than object-extracted RCs when the RC modifies the subject NP, but there should be no difference when the RC modifies the object. Finally, perspective-shift theory makes no clear predictions for the embedding factor during the processing of the RC. Over the course of processing the sentences, perspective-shift theory predicts that the additionally embedded conditions should be more complex because of an extra perspective shift (which is initiated as the NP prior to the RC—*the reporter* in (6)—is encountered), but the theory does not predict this additional complexity should manifest itself during the processing of the RC.

Second, consider the canonical word-order hypothesis with respect to the non-embedded conditions (6a)–(6d). Like the integration cost hypothesis, the canonical word order theory predicts that the subject-extracted RCs (6c) and (6d) should be easier than their object-extracted counterparts (6a) and (6b), because the word order is canonical SVO in subject-extracted RCs, but non-canonical OSV in object-extracted RCs. Furthermore, like the storage theories, the canonical word order theory predicts that object-modifying RCs should be easier to process than subject-modifying RCs. This prediction is made because the word order is more canonical overall during the processing of the RC for object-modifying conditions. In subject-extracted RCs, the word order in the object-modifying condition (6d) is canonical SVO SVO, whereas the word order in the subject-modifying condition (6c) is S SVO VO, which contains sequences like SSV and OVO, which are less canonical. In object-extractions, the object-modifying word order in (6b) is SVO OSV, which contains one SVO canonical sequence, whereas the subject-modifying word order (6a) is S OSV VO, which contains no canonical SVO sequences. Thus the canonical word order theory makes the same predictions as the storage cost/integration cost theory for the non-embedded conditions.

It is difficult to apply the canonical word order hypothesis to the embedded conditions, because the hypothesis has not been adequately

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

formalized. With the addition of a matrix subject (*the fact that* ... in [6]) before the non-embedded versions of the conditions, none of the conditions consists of canonical SVO order. There are sequences of canonical SVO word orders as described above, but there are non-canonical sequences in all conditions as well. One prediction of a version of a canonical word order theory is that there may be no differences in the embedded conditions, because all are non-canonical with the inclusion of the preceding subject NP. Another version of a canonical word order theory might predict the same pattern of results as in the non-embedded condition, but slower overall RTs, because of the difficulty associated with the non-canonical initial subject NP. But until some version of such a theory is formalized, it is difficult to discuss any potential predictions in detail.

2.1. *Method*

2.1.1. *Participants.*

Seventy-two participants from MIT and the surrounding community were paid for their participation. All were native speakers of English and were naive as to the purposes of the study.

2.1.2. *Materials and design.*

32 sets of sentences were constructed, each with eight conditions, crossing *modifier type* (subject-modifier, object-modifier), *extraction type* (subject-extracted, object-extracted) and *embedding* (non-embedded, embedded). The RC consisted of the same words in each of the conditions, with the noun phrase preceding the verb in the object-extracted RCs and the verb preceding the noun phrase in the subject-extracted RCs. Also, the noun phrase that was modified by the RC (the subject in subject-modifying RCs, the object in object-modifying RCs) was identical in all conditions. The target region—in bold in the example item in (6)—consisted of the RC in all conditions: the wh-pronoun *who* plus an NP and a verb. Note that the RC occurs at the end of the sentence in the non-embedded object-modifying conditions (6b) and (6d). Because people read sentence-ending words more slowly than other words (wrap-up effects), we included a prepositional phrase (PP) at the end of the RC in all conditions. The PP was then at the end of sentence in the non-embedded object-modifying conditions. It should be noted that the PP is *not* part of the critical region of analysis, because (1) it occurs sentence-finally in the non-embedded object-modifying conditions; and (2) there is a PP-attachment ambiguity in the subject-extracted versions (where the PP can initially be attached to the preceding verb or NP) that is not present in the object-extracted versions (where the PP can be attached only to the preceding NP). As a result of these confounds, we did not analyze the PP region, because of the difficulty of interpreting any results here.

All items and their eight versions are given in section 1 of the appendix. An example item is presented in (6). In addition to the target sentences, 74 filler sentences with various syntactic structures were included, including sentence materials from two other experiments. Each participant saw only one of the eight versions of each sentence, and each version was read by the same number of participants, according to a Latin-square design. The stimuli were pseudo-randomized separately for each participant, so that a target sentence never immediately followed another target sentence.

To ensure that processing differences between the object-extracted (*the reporter who the senator attacked*) and subject-extracted versions (*the reporter who attacked the senator*) of the RCs were not due to any plausibility differences, a plausibility survey was conducted. In order to preserve meaning and lexical content, while removing the specific syntactic structure, both versions were transformed into simple descriptions with a subject-verb-object structure (*the senator attacked the reporter* versus *the reporter attacked the senator*). Twenty-four participants from the same population, but who did not participate in the main experiment, rated sentences from 1 (very natural) to 7 (very unnatural) based on the naturalness of the events they describe in the real world. Two lists of 32 items, consisting of 16 each from the subject-extracted and object-extracted conditions, were constructed. Each list was given to an equal number of participants. The results of this plausibility survey showed that the subject-extracted (rating of 3.58) and object-extracted RCs (rating of 3.49) we used in the present experiment are equally natural (both $F1$ and $F2 < 1$).

2.1.3. *Procedure.* The task was self-paced word-by-word reading with a moving window display (Just, Carpenter and Woolley 1982) using a Macintosh computer running software developed in the lab. The Macintosh display allowed for up to 100 characters to appear on each line. Each trial began with a series of dashes marking the length and position of the words in the sentences. Participants pressed the spacebar to reveal each word of the sentence. As each new word appeared, the preceding word disappeared. The amount of time the participant spent reading each word was recorded as the time between key-presses. To make sure the participants read the sentences for meaning, a comprehension question appeared after the final word of each sentence which asked about information contained in the sentence they just read. Participants pressed one of two keys to respond *yes* or *no* to the comprehension question. After an incorrect answer, the word *INCORRECT* flashed briefly on the screen. No feedback was given for correct responses. Participants were asked to read sentences at a natural rate and to be sure that they understood what

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

they read. They were told to answer the questions as quickly and accurately as they could and to take wrong answers as an indication to read more carefully.

Before the main experiment started, a short list of practice items and questions was presented in order to familiarize the participants with the task. Participants took approximately 20 minutes to complete the experiment. For most participants, this experiment was combined with an unrelated experiment using the same self-paced reading task. Participants were able to take a short break between the two experiments.

2.2. *Results*

2.2.1. *Comprehension question performance.*

The comprehension questions for the experimental items were answered correctly on 80.3% of the trials. The percentages of correct answers per condition are presented in Table 1. A three-factor ANOVA crossing Modifier Type, Extraction Type and Embedding on the these question-answering data revealed that questions about embedded sentences (76.9% correct) were significantly harder to answer than questions about non-embedded sentences (83.6%), both in the analysis over subjects ($F1(1,71) = 18.35$, $p < 0.001$) and in the analysis over items ($F2(1,31) = 19.51$, $p < 0.001$). The only other significant effect was an interaction between Modifier Type and Embedding, which was significant in the analysis over items ($F1(1,31) = 5.80$, $p < 0.05$) but marginal in the analysis over subjects ($F1(1,71) = 3.95$, $p = 0.05$). In particular, the effect of Embedding was smaller in sentences containing object-modifiers (77.4% embedded versus 80.7% not-embedded) compared to sentences with subject-modifiers (76.4% embedded versus 86.5% not-embedded). This interaction was predicted by none of the theories that we considered. It may have been caused by differences in the difficulty of the questions across the conditions.

2.2.2. *Reading times.*

To adjust for differences in word length across conditions as well as overall differences in participants' reading rates, a

Table 1. *Experiment 1 comprehension question performance, as a function of modifier type, extraction type and embedding*

	Subject-modifier		*Object-modifier*	
	Subj-extracted	Obj-extracted	Subj-extracted	Obj-extracted
Non-embedded	88.2	84.7	81.3	80.2
Embedded	78.5	74.3	75.7	79.2

regression equation predicting reading times from word length was derived for each participant, using all filler and target items (Ferreira and Clifton 1986; see Trueswell et al. 1994, for discussion). At each word position, the reading time predicted by the participant's regression equation was subtracted from the actual measured reading time to obtain a residual reading time.

Because the comprehension questions were mainly included to make sure that the participants were reading for comprehension, all items were analyzed, regardless of how the comprehension question was answered. In any case, the statistical analyses that are reported below gave identical results whether or not we excluded trials in which the comprehension questions were answered correctly, or whether we analyzed raw reading times.

Because the predictions concerned the comprehension complexity of the RCs (e.g., *who the senator attacked/who attacked the senator* in (6)), we will focus on these reading times. Figure 1 shows the residual reading times of the RCs by condition. Tables of raw and residual reading times organized by condition are presented in section 2 of the appendix.

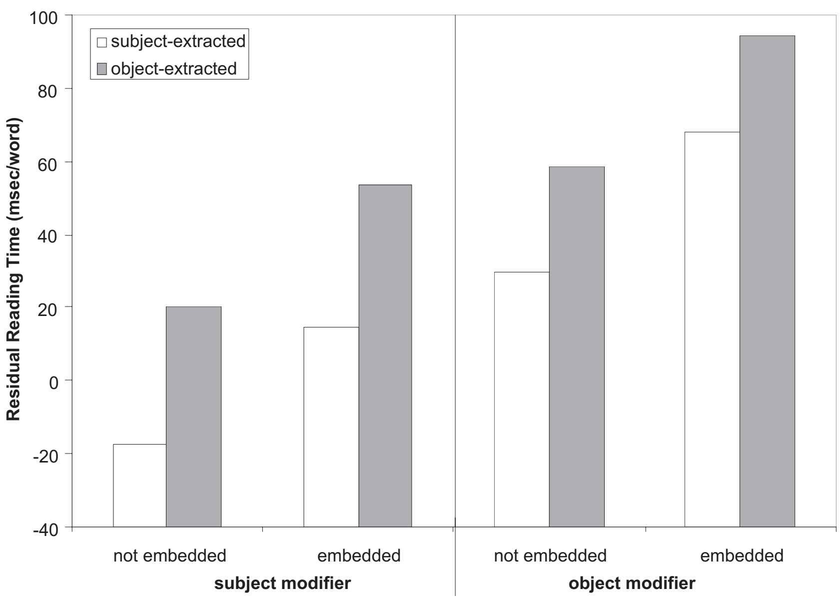

Figure 1. *Experiment 1 residual reading times in the RC, as a function of modifier type, extraction type and embedding.*

A three-factor ANOVA crossing modifier type, extraction type and embedding, revealed that all three factors had significant main effects, with no interactions among the factors. The main effect of embedding ($F1(1,71) = 6.47,\ p < 0.05;\ F2(1,31) = 10.40,\ p < 0.01$) revealed that the residual reading times of RCs were faster for non-embedded sentences (23 msec per word) than for embedded sentences (58 msec per word). This result was as predicted by the syntactic storage hypothesis. The main effect of extraction type ($F1(1,71) = 9.51,\ p < 0.01;\ F2(1,31) = 9.07,\ p < 0.01$) showed that it is harder to read object-extracted RCs (57 msec per word) than subject-extracted RCs (24 msec per word). This result was predicted by the integration hypothesis, and also the canonical word order hypothesis. Finally, there was a main effect of modifier type ($F1(1,71) = 9.56,\ p < 0.01;\ F2(1,31) = 25.37,\ p < 0.001$), but it was in the opposite direction as expected by syntactic storage based accounts, or the canonical word order hypothesis: subject-modifying RCs (18 msec per word) were read more quickly than RCs modifying the object (63 msec per word). Additional analyses revealed that all three effects were additive. Neither the three-way interaction, nor any of the two-way interactions showed any hints of being significant (all Fs < 1). In particular, the interaction that was predicted by perspective-shift theory was not significant (Fs < 1).

One possible source for the modifier position effect observed here is word position in the sentence: Earlier words in a sentence might be read more quickly than later words in a sentence simply because people might slow down in the course of reading sentences in our self-paced reading task. Because subject-modifiers always occurred earlier in the sentences than object-modifiers, such a slow-down with word position could potentially account for the observed result. Before we report the results of an analysis of this hypothesis, we should first point out that the general tendency in reading sentences is the *reverse* of this hypothesis: People generally read more quickly as they get later into sentences, presumably because they have more context to which to connect the incoming words (Just and Carpenter 1980). In order to test the hypothesis that people are slowing down through our materials, we analyzed RTs in the PP at the end of the RC (e.g., *on Tuesday* in (6)). If the later word position was the cause of the slower RTs in the object-modifiers, then we should see the same effect for the PPs following the object-modifying RCs: They should be read more slowly than the PPs in the subject-modifying RCs. There was no such effect (Fs < 1), in spite of the fact that the PP is the sentence-final region in two of the object-modifying conditions, and people tend to read the final regions of sentences more slowly than earlier regions (Just and Carpenter 1980). In fact, the numerical tendency was in

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

the opposite direction: 43 msec/word for the PP in the subject-modifying RC vs. 31 msec/word for the PP in the object-modifying RC. This analysis therefore excludes the possibility that the modifier position effect might have been caused by a general tendency for participants to read later words in sentences more slowly, especially in light of the fact there is no such tendency in any other previous reading study in English that we know of.

2.3. *Discussion*

The three experimental manipulations in this experiment had additive effects on the reading times of the RC. Embedded RCs were read more slowly than non-embedded RCs. Object-extracted RCs were read more slowly than subject-extracted RCs. Finally, and perhaps most surprisingly, object-modifying RCs were read more slowly than subject-modifying RCs. This last finding, which replicates Holmes (1973) using an on-line measure, rebuts the generally accepted idea that nested sentences are universally harder than right-branching sentences. The first result is as predicted by the storage hypothesis, and the second result is predicted by the integration hypothesis as well as some versions of the canonical word order hypothesis. But the third result runs counter to the predictions of all current theories of nesting, including the storage hypothesis and the canonical word order hypothesis. The results are also not consistent with the on-line application of perspective shift theory. Perspective shift theory predicted an interaction between extraction type and modifier position, but there was no such interaction in our results. In particular, the extraction-type effect was just as large for object-modifying RCs (27 msec per word in the RC) as for subject-modifying RCs (38 msec per word). These results suggest that the extraction-type difference observed for subject-modifiers is not due to perspective shift, because the same difference persists when there is no difference in perspective shifts, in the object modifying RCs. These results support either the integration cost interpretation or the canonical word order interpretation of the subject- vs. object-extraction results.

How can we reconcile the current results, which demonstrate a benefit for nested structures over non-nested structures in single embeddings, with earlier results, which show a benefit for right-branching structures over doubly-nested structures?[1] It is possible that some of the results may be explained by the canonical word order hypothesis. But because this hypothesis has not been adequately formalized, it is difficult to see what it predicts, even for the conditions that have been considered here. One version of this hypothesis is inadequate in two ways: (1) it

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

predicts that object-modifying structures should be easier than subject-modifying RCs for the non-embedded conditions; and (2) it predicts that this effect should disappear in the embedded conditions, because all of the embedded conditions include non-canonical word order. In contrast to these predictions, there was an advantage for subject-modifying RCs in both the embedded and non-embedded conditions, of approximately equal effect size. Although it is possible that other versions of the canonical word order hypothesis may better account for the observed data pattern, we will not consider this hypothesis in more depth here, because of the difficulty of pursuing an inadequately formalized theory.

The proposal that we will pursue here is that there are three independent factors at play. The first is storage, in terms of predicted categories or partially processed phrase structure rules, as in the storage cost hypothesis. Storage accounts for the effect of embedding, because the prediction of an additional verb must be stored during the processing of the critical RC in the embedded conditions. The second factor is integration, such that longer distance dependencies are more complex than shorter ones. Integration costs explain why object-extracted RCs are more complex than subject-extracted RCs. These two factors form the basis of the dependency locality theory, first presented in Gibson (1998), and more extensively in Gibson (2000).

A third factor is required to account for the observation that subject-modifying RCs were read more quickly than object-modifying RCs. We hypothesize that differences in the information-flow properties of the RCs in the two positions can account for this effect (Chafe 1976, 1987; Du Bois 1987; Givón 1979, 1983, 1984; Prince 1981). According to information flow, intonational, grammatical, and word choices in sentence production can in part be determined by conventions or interactionally determined choices between speakers. Of interest for our purposes is that English exhibits a general pattern in which the material in the subject position is usually old, sometimes indicated by terms as theme, topic or presupposed background. On the other hand, new information that comprises the core assertion of the utterance tends to come at the end of the sentence, within the predicate (Halliday 1970; Givón 1984; Chafe 1987; Gundel et al. 1988). We propose that people will experience comprehension difficulty in the form of slower processing when there is a conflict between the type of information being conveyed, and its position in the sentence. Thus people will slow down when old, background information is presented late in a sentence, or when new information is presented early in a sentence. We refer to this hypothesis as the information flow hypothesis for English:

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

(7) The information flow hypothesis: Old, background information is comprehended more easily early in a sentence, such as in a position modifying the subject; new, foreground material is processed more easily later in a sentence, such as in a position in the main predicate of the sentence.[2]

The information flow hypothesis is relevant to this study because restrictive RCs—the form of the RCs in Experiment 1—typically contain background information. In particular, one of the primary discourse functions of a restrictive RC is to identify a particular referent from among a group of entities. In order to perform this identification, background information which is common to both the speaker/writer and the hearer/reader is usually used to select the target referent from the group. For example, consider (8):

(8) The boy that studied for the exam aced the test.

A sentence like (8) is typically produced when the information in the restrictive RC 'that studied for the exam' is already available in the context. That is, (8) would typically be uttered in a context in which it is known to both the speaker/writer and the hearer/reader which boys studied for the exam and which boys did not.

Because restrictive RCs typically contain background information, the information flow hypothesis predicts that they are processed more easily earlier in sentences rather than later in sentences. Thus the information flow hypothesis accounts for the fact that subject-modifying RCs are read faster than object-modifying RCs.

As stated in (7), the information flow hypothesis is descriptive. That is, we have not yet proposed why a conflict between sentence position and informational content should cause processing difficulty. Before we address this issue, we first test the hypothesis further in a second experiment. We return to the issue of the potential cognitive underpinnings of (7) in the General Discussion. There, we also return to the issue of how the proposed three-factor account can explain the contrast between doubly-nested structures and right-branching structures.

3. Experiment 2

Experiment 2 directly tested the information flow hypothesis in (7) by comparing restrictive RCs to non-restrictive RCs in both subject- and object-modifying positions. In contrast to a restrictive RC, the discourse function of a non-restrictive RC is to provide extra information about the entity being modified, but which is not part of the core assertion of the utterance. Thus, unlike restrictive RCs, non-restrictive RCs typically contain new information about the entity that they modify. For example, consider the non-restrictive RC in (9):

(9) My father, who ate ham this morning, became extremely ill.

The non-restrictive RC *who ate ham this morning* provides a plausible cause for the assertion in the main clause. This information is often new to the discourse. This situation contrasts with the case of a restrictive RC, such as in (8): The information in a restrictive RC is usually present in the discourse.

If the restrictiveness of the RCs in Experiment 1 is responsible for the fact that subject-modifiers were read more quickly than object-modifiers, then we should replicate this finding for the restrictive versions of the conditions in Experiment 2, but not in the non-restrictive conditions. That is, the information flow hypothesis predicts that the restrictive RCs should be read more quickly in subject-position than in object position. Depending on the content of the RCs, the information flow hypothesis predicts the reverse effect in the non-restrictive conditions. That is, because non-restrictive RCs generally contain new information, the information flow hypothesis predicts that non-restrictive RCs should be processed more quickly in object-position than in subject-position. The information flow hypothesis therefore predicts an interaction between the restrictiveness of the RC (restrictive, non-restrictive) and the RC position (subject-modifying, object-modifying) during the processing of the RC.

A number of extra-sentential and intra-sentential cues were used to make sure that the participants noticed the difference between restrictive and non-restrictive RCs. An example item is given in (10).

(10) a. Subject-modifier, restrictive
A group of film critics praised a director at a banquet and another director at a film premiere. The director *that the critics praised at a banquet* insulted an actor from a big action movie during an interview.

b. Object-modifier, restrictive
A group of film critics praised a director at a banquet and another director at a film premiere. An actor from a big action movie insulted the director *that the critics praised at a banquet* during an interview.

c. Subject-modifier, non-restrictive
A group of film critics praised a director and a producer. The director, *who the critics praised at a banquet*, insulted an actor from a big action movie during an interview.

d. Object-modifier, non-restrictive
A group of film critics praised a director and a producer. An actor from a big action movie insulted the director, *who the critics praised at a banquet*, during an interview.

First, we presented a single-sentence context before the target sentences. The contexts either supported a restrictive or a non-restrictive interpretation. In the restrictive condition the context contained two possible referents for the noun phrase that was modified. The RC in the target sentence was then used to single out one of these two referents by using information that was given in the context. For example, two directors are introduced in (10a) and (10b), one of which is praised at a banquet, while the other is praised at a film premiere. Subsequently, the restrictive RC makes clear which of the two directors is intended in the target sentence. Sentences with non-restrictive RCs followed contexts where only one possible referent was presented. For example, only one director is introduced in (10c) and (10d). The non-restrictive RC then conveys some new information about the modified noun phrase, at the point of processing the prepositional phrase (PP, *at a banquet* in (10)).

In addition to the explicit manipulation of the preceding context, we provided two intra-sentential cues to indicate the difference between restrictive and non-restrictive RCs. Whereas the restrictive RCs were introduced with the complementizer *that*, the non-restrictive RCs began with the wh-pronouns *who* or *which*. For most American English speakers, the overt complementizer *that* cannot be used in a non-restrictive RC and therefore unambiguously signals a restrictive RC. Second, the non-restrictive RCs were separated from the noun phrase they modified by a comma, while no comma was present in the sentences with restrictive RCs. A comma imposes an intonation break between the noun and the modifying clause, which is inconsistent with restrictive modification (Selkirk 1984).

Because the contents of the non-restrictive RCs in our examples always included both old information from the preceding context (everything except the PP at the end of the RC) together with some new information (the PP), the information flow hypothesis does not make a strong prediction as to whether subject- or object-modifying RCs should be faster in the non-restrictive RCs. In particular, if all the information in the RC were new, then the information flow hypothesis would predict that the object-modifications should be faster. But because most of the information in the non-restrictive RCs is necessarily old information (in order to be minimally different from the restrictive conditions), the non-restrictive RCs contain conflicting sources of information: on the one hand, old

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

information from the context; and on the other hand, some new information and the syntax of a non-restrictive modifier, which suggests new information. The presence of old information in the non-restrictive RCs might then lead to faster RTs for the subject-modifiers than otherwise might be expected if only new information were present in the RC. This speed-up might offset an RT preference for object-modifiers over subject-modifiers. This conflict may then result in little or no difference between the subject- and object-modified non-restrictive RCs. In any case, the critical prediction for the experiment is that there should be less of a subject-modifier advantage for the non-restrictives than for the restrictives.

3.1. *Method*

3.1.1. *Participants.* 48 participants from MIT and the surrounding community were paid for their participation. All were native speakers of English and were naive as to the purpose of the study.

3.1.2. *Materials and design.* Sixteen sets of sentences were constructed. Each set contained four versions, crossing *restrictiveness* (restrictive, non-restrictive) with *modifier type* (subject-modifier, object-modifier). Each item consisted of two sentences: a context sentence and the target sentence containing the RC. The context sentence consisted of an indefinite subject NP (e.g., *a group of film critics* in [10]) followed by a verb (*praised* in [10]), and an object NP having one of two forms, depending on the restrictiveness factor. The first type of object NP was used in the restrictive conditions to introduce two entities to be referred to using the same head noun. This NP consisted of two conjoined indefinite NPs with the same head noun, the first introduced by the indefinite determiner *a/an* and the second introduced by the determiner *another*. Each of these indefinite NPs was modified by a prepositional phrase (e.g., *a director at a banquet and another director at a film premiere* in [10]). The second type of object NP was used in the non-restrictive conditions. In these items, the object NP consisted of two indefinite NPs conjoined together, with no prepositional phrase modification (e.g., *a director and a producer* in [10]).

The target sentence had one of two forms depending on the subject-/object-modification factor. In the subject-modification conditions, the target sentence consisted of a definite subject NP which referred to one of the object NPs of the previous sentence (e.g., *the director*), followed by the critical RC (e.g., *that/who the critics praised at a banquet*), then followed by the main verb of the sentence (e.g., *insulted*), and an indefinite object NP which included a PP modifier (e.g., *an actor from a big action* *movie*), and finally a PP which preferentially modified the matrix verb in the sentence (e.g., *during an interview*). In the object-modification conditions, the target sentence consisted of the same elements as in the subject-modification conditions, but with the subject and object switched. That is, the target sentence in the object-modification conditions consisted of an indefinite subject NP which included a PP modifier (e.g., *an actor from a big action movie*) followed by the main verb of the sentence (e.g., *insulted*), then a definite NP which referred to one of the object NPs from the previous sentence (e.g., *the director*) and the critical RC. Finally, a PP which preferentially modified the verb in the RC completed the target sentence (e.g., *during an interview*). As in Experiment 1, the final PP was included so that the target region—the RC—was not in sentence-final position in the object-modification conditions, which could have led to sentence wrap-up effects during this region. There was sometimes some ambiguity of attachment of the sentence-final PP, but this was not a critical region of analysis for the experiment, so this ambiguity did not matter to the hypotheses in question.

The target RC in all four conditions consisted of an object-extracted RC, with a PP modifying the verb (*that/who the critics praised at a banquet*). In the restrictive conditions, the RC was introduced by the relative pronoun *that*, whereas in the non-restrictive conditions the RC was introduced by the relative pronoun *who* and was separated from the subject and the main verb of the sentence by commas.

The critical region for analysis in this experiment consisted of the whole RC not including the first word of the RC (*that/who*), because this differed across the restrictive/non-restrictive conditions. It should be noted that the PP in the RC could logically be interpreted as modifying the main verb in the object-modification conditions, but not in the subject-modification conditions. For example in (10), the PP *at a banquet* can modify either the verb in the RC *praised* or the main verb of the sentence *insulted*. Although this ambiguity is present in the object-modification conditions and not in the subject-modification conditions, this is likely not an important confound in the design of the materials. Most importantly, there have been a number of studies that have demonstrated a strong locality preference in the case of ambiguities involving potential attachments to two preceding VPs (e.g., Altmann, van Nice, Garnham, and Henstra 1998; Pearlmutter and Gibson 2001). Thus although the PP could logically attach to the non-local verb, it is likely that participants rarely noticed this alternative. In any case, to be safe we analyzed the RC with and without the PP included.

A full list of items is given in section 3 of the appendix. In addition to the experimental sentences, 40 filler items with various syntactic

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

structures were included. Each participant saw only one of the four versions of each sentence, and each version was read by the same number of participants, according to a Latin-square design. The stimuli were pseudo-randomized separately for each participant, so that at least one filler item was presented between two target sentences.

3.1.3. *Procedure.* The task was the same self-paced moving-window word-by-word reading task that was used in Experiment 1. Each experimental session averaged 20 minutes. Most participants also took part in a second unrelated self-paced reading experiment. Participants were given short breaks between the two experiments.

3.2. *Results*

3.2.1. *Comprehension question performance.* The comprehension questions were answered correctly 76.6% of the time, broken down as follows. When the sentences contained a restrictive RC, the accuracy was 75.3% in the subject-modifying condition and 78.4% in the object-modifying condition. When the sentences contained a non-restrictive RC, the percentages were 74.6% in the subject-modifying condition and 77.9% in the object-modifying condition. A two factor ANOVA revealed no main effects nor interaction (Fs < 1.64, ps > 0.20).

3.2.2. *Reading times.* The analysis was similar to that for Experiment 1. First, we localized our analysis to the RC. We excluded the complementizer from analysis, because this differed between the restrictive (*that*) and non-restrictive conditions (*who*, *which*). In the first analysis we report, we included the prepositional phrase (e.g., *at a banquet* in [10]). In a second analysis, we examined the RCs without the PP.

As in Experiment 1, residual reading times were calculated, and all trials were analyzed, whether the associated comprehension question was answered correctly or not. The pattern of results was the same when only correct trials were analyzed. Mean residual reading times for the RC are presented in Figure 2. Analyses of raw times revealed the same patterns as for residual times, although not all effects reached significance in the raw time analyses. Tables of raw and residual reading times organized by condition are presented in section four of the appendix.

A two-factor ANOVA over the RC revealed three significant effects. First, there was a main effect of restrictiveness, such that restrictive RCs were read more quickly than non-restrictive RCs ($F1(1,47) = 3.98$, $p = 0.05$; $F2(1,15) = 6.31$, $p < 0.05$). Second, there was a main effect of modifier position, such that subject-modifying RCs were read

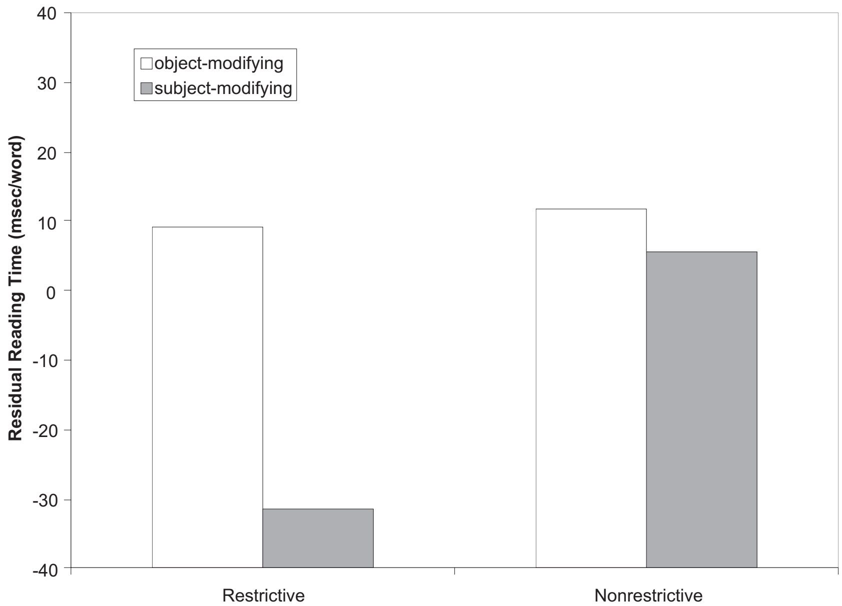

Figure 2. *Experiment 2 residual reading times for the subject NP, verb and prepositional phrase in the RC (e.g.,* the critics praised at a banquet *in [10]), as a function of modifier type and restrictiveness.*

more quickly than object-modifying RCs ($F1(1,47) = 5.78$, $p < 0.05$; $F2(1,15) = 11.13$, $p < 0.005$). Third, and, most importantly, there was an interaction between restrictiveness and modifier position ($F1(1,47) = 3.67$, $p = 0.06$; $F2(1,15) = 4.73$, $p < 0.05$), although this effect did not quite reach significance in the participants analysis. We also performed planned comparisons between subject- and object-modifying RCs separately for the restrictive and nonrestrictive contexts. In the restrictive conditions, subject-modifying RCs were read more quickly than object-modifying RCs ($F1(1,47) = 11.70$, $p < 0.001$; $F2(1,15) = 21.42$, $p < 0.001$), replicating the results from Experiment 1. In contrast, in the nonrestrictive conditions there was no difference between the reading times for the subject- and object-modifying RCs (*F*s < 1).

An analysis of the RC excluding the PP region was also performed. The only significant effect in this analysis was an effect of modifier position, such that subject-modifying RCs were read more quickly than object-modifying RCs ($F1(1,47) = 4.36$, $p < 0.05$; $F2(1,15) = 5.58$, $p < 0.05$). There was a tendency toward an interaction between restrictiveness and modifier position in this region in the items analysis ($F2(1,15) = 3.76$,

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

$p = 0.07$), but this effect was not significant in the participants analysis ($F1(1, 47) = 1.24$, $p = 0.24$). As in the full RC region, planned comparisons between subject- and object-modifying RCs revealed that restrictive subject-modifying RCs (-39.2 msec/word) were read more quickly than restrictive object-modifying RCs (-4.7 msec/word; $F1(1, 47) = 17.21$, $p < 0.001$; $F2(1, 15) = 21.42$, $p < 0.001$). In contrast, in the non-restrictive conditions there was no difference between the reading times for the subject- and object-modifying RCs (-13.0 versus -4.6 msec/word, Fs < 1). These results are therefore very similar to those from RTs across the full RCs. Because the effects are present in the early part of the RC as well as in the full RC including the PP, the observed effects are probably not due to ambiguity of attachment of the PP in the object-modification conditions.

3.3. *Discussion*

The results of this experiment were generally as predicted by the information flow hypothesis in (7). In particular, subject-modifying restrictive RCs were read more quickly than object-modifying restrictive RCs, replicating the results from Experiment 1. In addition, the advantage for subject-modifying RCs disappeared for the non-restrictive RCs: There was no difference in reading times between subject-modifying non-restrictive RCs and object-modifying non-restrictive RCs. As discussed above, the lack of a difference in RTs for the non-restrictive conditions may have been due to the fact that there was a lot of old information in the content of the RCs, from the preceding context sentence, leading to conflicting cues in the non-restrictive RCs: (a) some old information, leading to relatively faster RTs for the subject-modifiers; and (b) some new information and the syntax of a non-restrictive, leading to relatively faster RTs for the object-modifiers. This conflict may have then led to roughly equal RTs in the two conditions.

In principle, one way to investigate possible sources of the similar RTs in the non-restrictive conditions is to examine RTs at different points in the RCs. Because the last word of the RC (*banquet* in [10]) is new information in the non-restrictive versions, this word may be processed more quickly in the object-modifying condition if the information-flow hypothesis is correct. Analyses revealed no such difference, but this may be because (a) this is only a single word region, leading to a lack of statistical power (and no additional words can be included in the region, because they differ across the subject- and object-modifying conditions); and (b) this word also happened to have been presented along with the RC-final comma, which would lead to clause wrap-up effects in both conditions, potentially masking differences. Thus, although the results of Experiment 2 suggestively support the information flow hypothesis, further work is still needed to evaluate the hypothesis further.

4. General discussion

It has long been thought that non-nested structures are universally less complex than nested structures, as predicted by Miller's interruption hypothesis (Miller and Chomsky 1963; Miller and Isard 1964). One of the most surprising results of the studies presented here is the demonstration that singly-embedded right-branching restrictive RCs are read more slowly than corresponding nested RCs. This result was obtained in Experiment 1 and replicated in Experiment 2 (cf. Holmes 1973). In order to account for this result, we proposed the information flow hypothesis: Background information (like that in restrictive RCs) is processed more quickly earlier in a sentence rather than later in a sentence. This hypothesis accounts for the observation that subject-modifying restrictive RCs are read more quickly than corresponding object-modifying restrictive RCs because 1) restrictive RCs usually include background information and 2) subject-modifying RCs occur earlier in a sentence than object-modifying RCs. The information flow hypothesis also generally predicted the pattern of results of Experiment 2: that non-restrictive RCs would not show the same advantage for subject-modifiers over object-modifiers, because non-restrictive RCs are not associated with any particular grammatical position, and so are not expected early in a sentence.

Although the information flow hypothesis in (7) can account for the modifier position effects observed here, we have yet to provide specific cognitive motivations for why a conflict between sentence position and informational content should cause processing difficulty. One possible explanation for this observation is that it may derive from differences in people's syntactic expectations in the two environments. It is well established that people have difficulty when they encounter a word that is not a possible continuation of the input string that they have processed thus far, thus resulting in (temporary) ungrammaticality. Following Gibson (1991) and Elman (1991), we hypothesize that people have syntactic expectations in the form of predictions about what the next potential words and syntactic categories will be at every parse state, based on the current syntactic structure(s) for the input thus far. Furthermore, following Jurafsky (1996), Tabor and colleagues (1997), Hale (2001) and Rohde (2002), we hypothesize that there is a continuum between predicted and unpredicted input words, such that there is more difficulty in integrating less expected input words, as determined by experience with the language.

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

Thus a word/syntactic category that is highly expected given the current structure and the current state of the grammar (as determined by the learner's experience with the language) will be processed quickly. At the other end of the continuum, when a word/syntactic category is very unexpected, it will be processed slowly, reflecting the processor's difficulty in finding a matching prediction.

Applying this general idea to the current complexity difference, we propose that people have difficulty with restrictive RCs that modify objects because these are unusual in their linguistic experience, whereas restrictive RCs that modify subjects are much more frequent, and therefore expected. A similar explanation applies to the non-restrictive RCs. Of course, any explanation of processing difficulty that relies on linguistic experience (e.g., Mitchell et al. 1995; Jurafsky 1996; Tabor et al. 1997; MacDonald 1999; Hale 2001; Rohde 2002) begs the question of why the differences should be there in the corpus in the first place. In this case, the relevant question is why it is that syntactic expressions marking old information tend to come earlier in a sentence, whereas as syntactic expressions marking new information tend to come later. We assume that this difference arises from cognitive mechanisms in production, such that it is cognitively easier for people prefer to start with information that they already know about. Thus, following MacDonald (1999), we hypothesize that differences in the production process give rise to differences in comprehensibility.

There has been some corpus work that is consistent with the experience-based syntactic-expectation hypothesis for these types of structures. Fox and Thompson (1990) examined a corpus of spoken speech and found that RCs that modified object NPs were more likely to provide new information about the NP, whereas RCs that modified subject NPs were more likely to link the head to entities in the discourse. In spoken speech, the cues that distinguish a restrictive RC from a non-restrictive RC are partly intonational (i.e., the placement of intonational boundaries around non-restrictive information) and partly discourse based (Watson and Gibson in press). Although Fox and Thompson did not code their corpus for intonational information, it is plausible that the RCs that linked their head nouns to entities in the discourse were restrictive RCs, and that those that provided new information were generally non-restrictive RCs. Thus Fox and Thompson's corpus data are consistent with the experience-based syntactic-expectation hypothesis for restrictive and non-restrictive RCs.

In addition to providing support for the hypothesis that information flow differences constrain sentence comprehension, Experiment 1 also provided evidence for integration and storage resource constraints on

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

sentence comprehension, the core components of the dependency locality theory (Gibson 1998, 2000). First, subject-extracted RCs were processed more quickly than object-extracted RCs, as predicted by a distance-based integration cost function. Second, RCs that were embedded within the sentential complement of a noun were read more slowly than comparable RCs that were not embedded in this way. This result is predicted by a storage theory such as the dependency locality theory that keeps track of predicted categories or partially processed phrase structure rules.

Given the proposed syntactic-expectation explanation of the information flow effects, it is worth considering whether the same explanation could be used to account for effects that are usually attributed to resource constraints, such as those exhibited in Experiment 1. An examination of the kinds of effects that resource theories account for suggests that a single experience-based syntactic-expectation constraint will not suffice to account for either integration or storage effects. First, consider English integration effects. A number of studies have shown that there is difficulty at the embedded verb in English object-extracted RCs (Gibson 1998; Grodner and Gibson in press; King and Just 1991), in spite of the fact that the verbal position is highly predictable given the previous context. That is, given a relative pronoun and a subject NP, a verb is highly expected, and yet RTs are relatively slow when such a verb is encountered. Thus syntactic expectations seem like an unlikely explanation for integration effects. Similarly, syntactic storage effects are unlikely to be explained in terms of syntactic expectations as conceived here. Once there is an open dependency (e.g., from a subject NP that takes a sentence complement, like *the fact that* ...), people process the following material slowly until the open dependency is resolved. But the words in the embedded clause are no more or less predictable from the preceding syntactic context whether or not there is an open dependency. For example, a verb is just as predictable following an embedded subject as following a main clause subject: in both cases a verb is 100% expected. Thus it seems unlikely that syntactic expectation constraints can account for resource effects.

We therefore propose a multiple constraint framework for sentence comprehension in which three of the constraints are (a) syntactic expectations, giving rise to information flow effects; (b) integration resources; and (c) storage resources. In this framework we hypothesize that each constraint is independent, contributing a cost to the processing difficulty at the point of processing a word in an input sentence. For example, consider the syntactic expectations constraint. Under the current proposal, the difficulty at a word depends on the expectedness of the word in that syntactic context: people will read more slowly and have more difficulty with more unexpected syntactic continuations (Jurafsky 1996; Hale 2001). Similarly, integration costs and storage costs are proposed to be additive to the total difficulty at a word, depending on the difficulty of the integrations and number of open syntactic predictions, respectively (Gibson 1998, 2000). Other constraints are proposed to be additive as well, including lexical frequency constraints (less frequent lexical items lead to more difficulty) and plausibility constraints (less plausible local continuations lead to more difficulty). The proposed framework is therefore generally consistent with earlier proposals in which multiple constraints interact in the word-by-word construction of sentence representations (see Gibson and Pearlmutter 1998, and Tanenhaus and Trueswell 1995, for summaries of relevant evidence). Furthermore, the proposed framework can account for the complexity of unambiguous materials, as well as preferences in resolving (temporary) ambiguities, such that people prefer ambiguity resolutions associated with less overall difficulty/cost.

Let us now work through how the proposed constraints may interact to provide the results from the current experiments. Consider the materials from Experiment 1 once again. First, the integration cost factor explains the uniform slowdown of object-extractions relative to subject-extractions across all the conditions. Second, the storage cost component of the theory explains the uniform slowdown when sentences are embedded in the sentential complement of a noun. The most interesting case is showing how the three factors account for the observation that a subject-modifying RC as in (11a) is processed more quickly than an object-modifying RC as in (11b):

(11) a. The reporter *that the senator attacked* ignored the president.
b. The president ignored the reporter *that the senator attacked.*

The RCs in (11a) and (11b) are both object-extracted, so the integration factor does not make differing predictions during their processing. The syntactic storage constraint contributes the cost associated with one additional predicted syntactic head to processing the RC in (11a) relative to (11b), because an additional category (the top-level verb) is needed when processing the RC to form a grammatical sentence in (11a). The syntactic expectations constraint—which is proposed to derive the information flow differences between the two—favors the subject-modification in (11a) over the object-modification in (11b), simply because a restrictive RC is more likely to modify a subject than an object. In order to account for the observation that subject-modifiers are processed more quickly than object-modifiers, we hypothesize that the syntactic-expectations constraint is strongly biased against the presence of a restrictive RC modifying an object NP, with the consequence that this cost is greater than the storage cost associated with processing the RC in subject position. The resultant sum of costs therefore favors the subject-modifying RC.

Let us now return to the contrast between doubly-nested structures and their right-branching controls (12a) and (12b), which was the original evidence in support of the interruption hypothesis:

(12) a. The student who the professor who the scientist collaborated with advised copied the article.
b. The scientist collaborated with the professor who advised the student who copied the article.

The nested structure in (12a) is much harder to understand than its right-branching counterpart in (12b). But the information flow factor predicts the opposite pattern: As in the singly-nested sentences in (11), the information flow factor favors the subject-modifying RC in (12a) over the object-modifying RC in (12b). The greater complexity of the nested version in (12a) can be accounted for by the other constraints within the multiple constraint approach to sentence comprehension assumed here. In particular, the integration and storage factors are heavily biased in favor of the non-nested structure in (12b) over the nested structure in (12a). First, consider integration. All the integrations are local in (12b), whereas the integrations in the nested (12a) are far longer. This contributes a heavy processing cost to the nested structure in (12a). Second, there is a larger storage cost difference between the doubly-nested (12a) and its right-branching counterpart (maximally five predicted syntactic heads in (12a) vs. only one in (12b)) than between the singly-nested versions in (11). Thus, although information flow favors the nested structure in (12a) over the non-nested structure in (12b), integration and storage factors greatly outweigh this tendency, with the result that (12b) is much easier to understand than (12a).

The results of Experiment 1 are also relevant to the question of how syntactic and resource constraints interact in sentence comprehension. In particular, the fact that the two resource constraints and the information flow factor had additive non-interactive effects indicates that the three factors may be independent. It is especially interesting that the two resource constraints do not appear to interact. This observation is counter to the claim made by Gibson (1998) who, following Just and Carpenter (1992), hypothesized that integration and storage would interact because they probably tapped the same resource pool. The results here suggest that Gibson's (1998) hypothesis was incorrect. Rather, it seems that storage and integration may tap into separate pools of resources. It is possible that the resource pool was not pushed close to its limit when participants were processing the items in Experiment 1, so that a potential interaction

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

was not visible. This seems unlikely, however, because the most complex items in Experiment 1 were quite complex, resulting in degraded question-answering performance.

In conclusion, this paper has provided evidence against the simplest form of the interruption hypothesis, which predicted that singly nested RCs should be harder to process than their right-branching counterparts. The evidence supports the view that constraints in information flow, possibly implemented in terms of differences in syntactic expectations, also contribute to sentence complexity alongside resource constraints in a multiple constraint sentence comprehension mechanism.

Received 23 October 2003 *Revision received 15 December 2003*

Massachusetts Institute of Technology
Ghent University, Belgium

Appendix

1. *Items used in Experiment 1.*

There were eight conditions in Experiment 1, crossing three factors: embedding (embedded, not embedded), extraction type (subject-extracted, object-extracted) and modifier type (subject-modifying, object-modifying). The four embedded versions of item 1 are presented below. The non-embedded versions are obtained by omitting the parenthesized material. For the remainder of the items, only the embedded, subject-extracted, subject-modifying version is given. The object-extracted versions may be obtained by switching the position of the noun phrase and the verb within the relative clause (e.g., by swapping *the senator* and *attacked* in item 1 below). The object-modifying versions may be obtained by switching the position of the matrix sentence object noun phrase with the matrix sentence subject noun phrase, which includes the modifying relative clause (e.g., by swapping *the president* and *reporter who the senator attacked on Tuesday* in item 1 below). Embedded versions are obtained by including the material in parentheses. Non-embedded versions are obtained by omitting this material.

1. a. (The chance that) the reporter who the senator attacked on Tuesday ignored the president (bothered the editor).
   b. (The chance that) the reporter who attacked the senator on Tuesday ignored the president (bothered the editor).
   c. (The chance that) the president ignored the reporter who the senator attacked on Tuesday (bothered the editor).
   d. (The chance that) the president ignored the reporter who attacked the senator on Tuesday (bothered the editor).

2. (The knowledge that) the babysitter who the parents liked very much played with the child (pleased the grandparents).
3. (The perception that) the banker who the chairman praised during lunch distrusted the broker (annoyed the clients).
4. (The information that) the violinist who the sponsors flattered at the rehearsal insulted the singer (disappointed the conductor).
5. (The realization that) the burglar who the police negotiated with on Monday had frightened the dog (distressed the neighbors).
6. (The speculation that) the carpenter who the plumber punched in the nose yelled at the painter (worried the contractor).
7. (The implication that) the accountant who the engineer advised during the meeting spoke to the secretary (irritated the boss).
8. (The observation that) the model who the artist teased after the debut winked at the journalist (excited the onlookers).
9. (The reminder that) the student who the professor trusted for a long time met with the provost (tormented the teaching assistant).
10. (The rumor that) the mobster who the media criticized on Monday kidnapped the spy (intimidated the attorney).
11. (The news that) the player who the coach screamed at after practice wrestled with the trainer (surprised the team).
12. (The thought that) the actor who the starlet annoyed a great deal forgot about the leading lady (amused the comedian).
13. (The fact that) the criminal who the lawyer sued for millions of dollars stared at the judge (unnerved the jury).
14. (The idea that) the suitors who the king entertained during the evening wanted to see the princess (overjoyed the queen).
15. (The discovery that) the bachelor who the socialite pursued with passion resented the millionaire (fascinated the tabloids).
16. (The discovery that) the councilman who the radio host provoked last week married the secretary (shocked the entire city).
17. (The observation that) the contestant who the judges joked with about the host turned toward the cameraman (pleased the audience).
18. (The revelation that) the child who the psychologist talked to during the therapy session had hurt the woman (worried the young couple).
19. (The news that) the diplomat who the prime minister insulted on Friday angered the dictator (discredited the government).
20. (The fact that) the tourists who the guide walked with during the visit waved at the nuns (embarrassed the priest).
21. (The report that) the politician who the voters spoke to during the campaign smiled at the preacher (softened the critics a bit).

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

22. (The suggestion that) the farmer who the aliens had communicated with at dusk phoned the newspaper (had strengthened the credibility of the article).
23. (The information that) the official who the governor argued with very loudly avoided the mayor (disillusioned the apprentice).
24. (The impression that) the clerk who the manager disliked very much smiled at the customer (intrigued the security guard).
25. (The suggestion that) the guitarist who the band played with at concerts despised the agent (upset the drummer).
26. (The speculation that) the salesman who the cashier resented for stealing merchandise ridiculed the shoppers (hurt business).
27. (The fact that) the waiter who the cook despised for being lazy ignored the busboy (bothered the owner).
28. (The disclosure that) the medic who the doctor worked with on the weekends scolded the patient (startled the board of directors).
29. (The evidence that) the passenger who the navigator had met at the party talked to the pilot (proved the identity of the criminal).
30. (The evidence that) the dog which the bear chased up a tree scratched the cubs (infuriated the owner).
31. (The report that) the cheerleader who the quarterback dated for a month denounced the track star (amused the team).
32. (The claim that) the raccoon which the fox bit on the leg ran from the deer (interested the nature show host).

2. *Residual reading times (raw times in parentheses) for Experiment 1 (msec/word), as a function of modifier type (subject- vs. object-modifying), extraction type (subject- vs. object-extraction) and embedding (embedded vs. non-embedded).*

	Subject-modifying sentences						
	Em-bedding *The fact that*	Subject NP *the reporter*	RC *who the senator . . .*	PP *on Tuesday*	Main verb *ignored*	Object NP *the president*	End *both-ered . . .*
Non-embedded Subj-ext	N/A	−24.27 (377)	−17.43 (377)	23.58 (411)	13.96 (419)	−26.93 (369)	1.86 (400)
Non-embedded Obj-ext	N/A	−33.60 (366)	20.00 (415)	72.10 (462)	34.47 (437)	−5.32 (389)	7.73 (404)
Embedded Subj-ext	−25.72 (372)	−60.39 (342)	14.78 (409)	14.85 (406)	59.41 (464)	26.44 (421)	108.91 (523)
Embedded Obj-ext	−40.01 (358)	−53.98 (346)	53.56 (449)	65.34 (454)	110.12 (515)	36.45 (432)	75.28 (492)

	Object-modifying sentences						
	Em-bedding *The fact that*	Subject NP *the reporter*	Main verb *ignored*	Object NP *the president*	RC *who the senator*	PP *on Tuesday*	End *both-ered ...*
Non-embedded Subj-ext	N/A	−23.23 (371)	−27.48 (379)	−12.29 (389)	29.88 (425)	30.55 (421)	14.83 (412)
Non-embedded Obj-ext	N/A	−12.67 (381)	−10.25 (392)	−17.95 (383)	58.69 (454)	36.82 (427)	13.86 (410)
Embedded Subj-ext	−21.47 (377)	−60.66 (334)	−8.86 (398)	15.47 (418)	68.41 (464)	27.64 (418)	58.78 (474)
Embedded Obj-ext	−13.17 (385)	−48.06 (347)	34.15 (436)	16.87 (416)	94.34 (489)	30.31 (418)	29.46 (444)

3. *Items used in Experiment 3*

There were four conditions in Experiment 2, crossing two factors: restrictiveness of the relative clause (restrictive, non-restrictive) and modifier position (subject-modifying, object-modifying). All four versions of item (1) are presented below. For the remainder of the items, only the non-restrictive subject-modifying target sentence of each item is presented. The restrictive context sentence is presented first, followed by the non-restrictive context sentence. The non-restrictive subject-modifying target sentence follows. The restrictive version of the target sentence can be formed by deleting the commas around the relative clause, and by replacing the relative pronoun *who/which* with *that*. The object-modifying versions are formed by swapping the subject NP (*the director* in (1)) with the object NP (*an actor from a big action movie* in (1)). The object NP consisted of the indefinite determiner *a/an* followed by a noun, followed by a prepositional phrase. There was a line break presented after the context sentence, so that the target sentence always started on a new line.

1. a. restrictive, subject-modifying: A group of film critics praised a director at a banquet and another director at a film premiere. The director that the critics praised at a banquet insulted an actor from a big action movie during an interview.
   b. restrictive, object-modifying: A group of film critics praised a director at a banquet and another director at a film premiere.

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

An actor from a big action movie insulted the director that the critics praised at a banquet during an interview.

c. non-restrictive, subject-modifying: A group of film critics praised a director and a producer. The director, who the critics praised at a banquet, insulted an actor from a big action movie during an interview.

d. non-restrictive, object-modifying: A group of film critics praised a director and a producer. An actor from a big action movie insulted the director, who the critics praised at a banquet, during an interview.

2. (A vicious guard dog bit a postman on the leg and another postman on the arm./ A vicious guard dog bit a postman and a milkman.) The postman, who the dog bit on the leg, saw a doctor from a nearby hospital but the bite got infected anyway.
3. (An art professor read a paper in the library and another paper in a pub./ An art professor read a paper and a book review.) The paper, which the professor read in the library, criticized an archaeologist at a Dutch university although some of the criticisms were unfounded.
4. (An 18th century British admiral captured a pirate off the coast and another pirate near an island./ An 18th century British admiral captured a pirate and a smuggler near England.) The pirate, who the admiral captured off the coast, taunted an officer of the British navy before the pirate was imprisoned.
5. (A talk show host interviewed a celebrity at a wedding and another celebrity at a fund-raiser./ A talk show host interviewed a celebrity and a politician.) The celebrity, who the host interviewed at a wedding, punched a cameraman with a red goatee after insults had been exchanged.
6. (A clerk helped a customer at the register and another customer at the tie rack./ A clerk helped a customer and a cashier.) The customer, who the clerk helped at the register, flirted with the owner of the clothing store while looking for a stack of sweaters.
7. (An evil villain imprisoned a superhero in a fortress and another superhero in his hideout./ An evil villain imprisoned a superhero and a police chief.) The superhero, who the villain imprisoned in a fortress, kissed a woman with long blond hair after the hero escaped to safety.
8. (A dean misquoted a philosopher at a party and another philosopher at a meeting./ A dean misquoted a philosopher and a famous novelist.) The philosopher, who the dean misquoted at a party,

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

wrote to a colleague in a different department because the dean's error upset him.

9. (A young woman carried a child in her arms and another child on her back as she walked through the airport./ A young woman carried a child and a backpack full of toys through the airport.) The child, who the woman carried in her arms, waved to a ticket agent at the gate before boarding the plane.
10. (The owner of a mansion hired a sculptor for a fountain and another sculptor for a statue./ The owner of a mansion hired a sculptor and a landscaper.) The sculptor, who the patron hired for a fountain, talked to the gardener of the enormous estate because remodeling was needed.
11. (A bully hit a student with a rock and another student with a binder./ A bully hit a student and a teacher after eating too much sugar.) The student, who the bully hit with a rock, visited the nurse at the high school so that the injury could receive treatment.
12. (A movie studio sued a producer over a contract and another producer over a budget dispute./ A movie studio sued a producer and a script writer.) The producer, who the studio sued over a contract, confronted a lawyer from the legal department despite warnings from his friends.
13. (A soccer coach scolded a player for being late and another player for poor defensive play./ A soccer coach scolded a player and a parent.) The player, who the coach scolded for being late, pushed an opponent from the other team because the two disliked each other.
14. (A senator attacked a reporter for bad journalism and another reporter for bribing a cop./ A senator attacked a reporter and a congressional leader.) The reporter, who the senator attacked for bad journalism, ignored the editor of the political news instead of addressing the claims.
15. (An FBI agent pursued a kidnapper for two years and another kidnapper for six months./ An FBI agent pursued a kidnapper and a counterfeiter across the country.) The kidnapper, who the agent pursued for two years, tackled a deputy with a black mustache after the police found him.
16. (A soldier hated a diplomat for political reasons and another diplomat for personal reasons./ A soldier hated a diplomat and a pentagon official.) The diplomat, who the soldier hated for political reasons, supported a general in the Army due to his military expertise.

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

4. *Residual reading times (raw times in parentheses) for Experiment 2 (msec/word), as a function of modifier type and restrictiveness.*

	Subject-modifying sentences							
	Context sentence *A group . . . premiere*	Subject NP *The director*	RC *that/ who the critics praised*	PP_1 *at a banquet*	Main verb *insulted*	Object NP *an actor*	PP_2 *from a big action movie*	End *during an inter-view*
Restrictive	4.11 (363)	51.60 (419)	−33.97 (328)	−23.96 (328)	−23.55 (371)	14.71 (373)	10.60 (367)	75.78 (444)
Non-restrictive	19.83 (379)	68.09 (443)	−1.97 (357)	24.00 (380)	78.21 (475)	35.62 (392)	−3.22 (354)	53.56 (422)

	Object-modifying sentences							
	Context sentence *A group . . . premiere*	Subject NP *An actor*	RC *from a big action movie*	PP_1 *insulted*	Main verb *the director*	Object NP *that/ who the critics praised*	PP_2 *at a ban-quet*	End *during an inter-view*
Restrictive	5.36 (365)	88.82 (444)	−8.16 (348)	15.84 (409)	4.84 (373)	3.80 (366)	23.26 (374)	70.28 (438)
Non-restrictive	23.78 (383)	100.6 (458)	−11.55 (344)	−30.4 (363)	29.31 (402)	11.17 (371)	28.84 (383)	68.93 (438)

5. *Residual reading times (raw times in parentheses) for the RCs in Experiment 2 (msec/word), as a function of modifier type and restrictiveness. Differences between these numbers and the RC numbers in section 4 of the appendix reflect the fact that the RC region in the current table includes RTs from the RC and following PP region (which is always part of the RC), whereas the RC region in section 4 of the appendix does not include the following PP region. In addition, the RC region in section 4 of the appendix includes RTs from the wh-pronoun in the RC, whereas the RC region in the current table does not include this word.*

	Object-modifying	Subject-modifying
Restrictive	9.3 (368)	−31.6 (326)
Non-restrictive	11.6 (372)	5.5 (366)

Notes

* Author's address: NE20-459, Department of Brain and Cognitive Sciences, MIT, Cambridge, MA 02139. Email: ⟨egibson@mit.edu⟩. We would like to thank the following people for helpful discussions on presentations of this work: Neal Pearlmutter, Doug Rohde and Tessa Warren. We would also like to thank Susanne Tunstall, Matt Cain and Lauren Tsai for running subjects in the experiments reported here. Funding for this work was provided by a grant provided by the National Science Foundation to the first author, award number SBR-9729037. Additional funding for the project was provided by the JST/MIT joint international research project *Mind/Articulation*. Timothy Desmet is a Postdoctoral Fellow of the Fund for Scientific Research—Flanders (Belgium) (F.W.O.—Vlaanderen).

1. One possibility, suggested by Holmes (1973), is that singly-embedded structures may be processed in a fundamentally different way from doubly-nested structures. Although this is a logical possibility, such a hypothesis should only be a last resort. Here, we pursue uniform theories of processing the two kinds of structures.
2. Information structure generalizations are usually stated in terms of subject and predicate positions, rather than early and late. These two ways of conceiving information structure are conflated in the examples under consideration here, so we cannot distinguish the two positions here.
3. There was one item in which the head noun for the RC was inanimate. The relative pronoun *which* initiated the non-restrictive conditions for this item.

References

Abney, Steven P. and Mark Johnson
1991 Memory requirements and local ambiguities of parsing strategies. *Journal of Psycholinguistic Research* 20, 233–250.

Altmann, Gerry T. M., Kathy Y. van Nice, Alan Garnham, and Judith-Ann Henstra
1998 Late closure in context. *Journal of Memory and Language* 38, 459–484.

Baird, R. and J. D. Koslick
1974 Recall of grammatical relations within clause-containing sentences. *Journal of Psycholinguistic Research* 3, 165–171.

Bever, Thomas G.
1970 The cognitive basis for linguistic structures. In John R. Hayes (ed.), *Cognition and the Development of Language*. New York, NY: Wiley, 279–362.

Blaubergs, Maija S. and Martin D. S. Braine
1974 Short-term memory limitations on decoding self-embedded sentences. *Journal of Experimental Psychology* 102, 745–748.

Blumenthal, Arthur L.
1966 Observations with self-embedded sentences. *Psychonomic Science* 6, 453–454.

Chafe, Wallace
1976 Givenness, contrastiveness, definiteness, subjects, topics, and point of view. In Li, Charles N. (ed.), *Subject and Topic.* New York, NY: Academic Press, 25–55.
1987 Cognitive constraints on information flow. In Tomlin, Russell (ed.), *Coherence and Grounding in Discourse*. Amsterdam/Philadelphia: Benjamins, 21–51.

Chen, Evan, Edward Gibson, and Florian Wolf
2003 Online Syntactic Storage Costs in Sentence Comprehension. Manuscript under submission.
Chomsky, Noam
1957 *Syntactic Structures*. The Hague: Mouton.
1959 On certain formal properties of grammar. *Information and Control* 2, 137–167.
1965 *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
Chomsky, Noam and George A. Miller
1963 Introduction to the formal analysis of natural languages. In Luce, R. Duncan, Robert R. Bush, and Eugene Galanter (eds.), *Handbook of Mathematical Psychology*, Volume 2. New York, NY: Wiley, 269–321.
Du Bois, John W.
1987 The discourse basis of ergativity. *Language* 63, 805–855.
Elman, Jeffrey L.
1991 Distributed representations, simple recurrent networks and grammatical structure. *Machine Learning* 7, 195–225.
Ferreira, Fernanda and Charles Clifton, Jr.
1986 The independence of syntactic processing. *Journal of Memory and Language* 25, 348–368.
Ford, Marilyn
1983 A method for obtaining measures of local parsing complexity throughout sentences. *Journal of Verbal Learning and Verbal Behavior* 22, 203–218.
Foss, Donald J. and R. H. Lynch
1969 Decision processes during sentence comprehension: Effects of surface structure on decision times. *Perception and Psychophysics* 5, 145–148.
Fox, Barbara A. and Sandra A. Thompson
1990 Relative clauses in English conversation. *Language* 66, 297–316.
Gibson, Edward
1991 A computational theory of human linguistic processing: Memory limitations and processing breakdown. Unpublished Ph.D. thesis, Carnegie Mellon University.
1998 Linguistic complexity: Locality of syntactic dependencies. *Cognition* 68, 1–76.
2000 The dependency locality theory: A distance-based theory of linguistic complexity. In Miyashita, Yasushi, Alec Marantz, and Wayne O'Neil (eds.), *Image, Language, Brain*. Cambridge, MA: MIT Press, 95–126.
Gibson, Edward and Neal Pearlmutter
1998 Constraints on sentence comprehension. *Trends in Cognitive Science* 2, 262–268.
Gibson, Edward and James Thomas
1996 The processing complexity of English center-embedded and self-embedded structures. *Proceedings of the NELS 26 Workshop on Language Processing, MIT Working Papers in Linguistics*. Cambridge, MA: MIT Press.
Gibson, Edward and Suzanne Tunstall
1999 The parser need not keep track of contingent frequencies: A lexical account of the English determiner/complementizer ambiguity. Poster presented at the 12th CUNY sentence processing conference.
Givón, Talmy
1979 *On Understanding Grammar.* New York, NY: Academic Press.
1983 *Topic Continuity in Discourse*. Amsterdam/Philadelphia: Benjamins.
1984 *Syntax*, Volume I. Amsterdam/Philadelphia: Benjamins.

Grodner, Daniel and Edward Gibson
In press Consequences of the serial nature of linguistic input. *Cognitive Science*.
Grodner, Daniel, Edward Gibson, and Suzanne Tunstall
2002 Syntactic complexity in ambiguity resolution. *Journal of Memory and Language* 46, 267–295.
Gundel, Jeanette K.
1988 Universals of topic-comment structure. In Michael Hammond, Edith Moravczik, and Jessica R. Wirth (eds.), *Studies in syntactic typology*. Amsterdam/Philadelphia: Benjamins, 209–239.
Hakes, David T. and Helen S. Cairns
1970 Sentence comprehension and relative pronouns. *Perception and Psychophysics* 8, 5–8.
Hakes, David T., Judith S. Evans, and Linda L. Brannon
1976 Understanding sentences with relative clauses. *Memory and Cognition* 4, 283–290.
Hakuta, Kenji
1981 Grammatical description versus configurational arrangement in language acquisition: The case of relative clauses in Japanese. *Cognition* 9, 197–236.
Hale, John
2001 A probabilistic early parser as a psycholinguistic model. *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics*, 159–166.
Halliday, Michael A. K.
1970 Language structure and language function. In Lyons, John (ed.), *New Horizons in Linguistics*. Aylesbury: Penguin, 140–165.
Hawkins, John A.
1994 *A Performance Theory of Order and Constituency*. Cambridge, UK: Cambridge University Press.
Holmes, V. M.
1973 Order of main and subordinate clauses in sentence perception. *Journal of Verbal Learning and Verbal Behavior* 12, 285–293.
Holmes, V. M. and J. K. O'Regan
1981 Eye fixation patterns during the reading of relative clause sentences. *Journal of Verbal Learning and Verbal Behavior* 20, 417–430.
Hsiao, Franny and Edward Gibson
2003 Processing relative clauses in Chinese. *Cognition* 90, 3–27.
Jurafsky, Daniel
1996 A probabilistic model of lexical and syntactic access and disambiguation. *Cognitive Science* 20, 137–194.
Just, Marcel and Patricia A. Carpenter
1980 A theory of reading: From eye-fixations to comprehension. *Psychological Review* 87, 329–354.
1992 A capacity theory of comprehension: Individual differences in working memory. *Psychological Review* 99, 122–149.
Just, Marcel A., Patricia A. Carpenter, and Timothy A. Keller
1996 The capacity theory of comprehension: New frontiers of evidence and arguments. *Psychological Review* 103, 773–780.
Just, Marcel A., Patricia A. Carpenter, and J. D. Woolley
1982 Paradigms and processing in reading comprehension. *Journal of Experimental Psychology: General* 111, 228–238.

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

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

King, Jonathan and Marcel A. Just
1991 Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language* 30, 580–602.

Lewis, Richard
1996 A theory of grammatical but unacceptable embeddings. *Journal of Psycholinguistic Research* 25, 93–116.

MacDonald, Maryellen C.
1999 Distributional information in language comprehension, production, and acquisition: Three puzzles and a moral. In MacWhinney, Brian (ed.), *The Emergence of Language*. Mahwah, NJ: Erlbaum, 177–196.

MacDonald, Maryellen C. and Morten Christiansen
2002 Reassessing working memory: Comment on Just and Carpenter (1992) and Waters and Caplan (1999). *Psychological Review* 109, 35–54.

MacWhinney, Brian
1977 Starting points. *Language* 53, 152–168.
1982 Basic syntactic processes. In Kuczaj, Stan (ed.), *Language Acquisition*: Volume I, *Syntax and Semantics.* Hillsdale, NJ: Erlbaum, 73–136.
1987 The competition model. In MacWhinney, Brian (ed.), *Mechanisms of Language Acquisition*. Hillsdale, NJ: Erlbaum, 249–308.

MacWhinney, Brian and Csaba Pleh
1988 The processing of restrictive relative clauses in Hungarian. *Cognition* 29, 95–141.

Marks, Lawrence E.
1968 Scaling of grammaticalness of self-embedded English sentences. *Journal of Verbal Learning and Verbal Behavior* 7, 965–967.

Miller, George A. and Noam Chomsky
1963 Finitary models of language users. In Luce, R. Duncan, Robert R. Bush, and Eugene Galanter (eds.), *Handbook of Mathematical Psychology*, Volume 2. New York, NY: Wiley, 269–321.

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

Mitchell, Donald C., Fernando Cuetos, Martin M. B. Corley, and Marc Brysbaert
1995 Exposure-based models of human parsing: Evidence for the use of coarse-grained (non-lexical) statistical records. *Journal of Psycholinguistic Research* 24, 469–488.

Pearlmutter Neal J. and Edward Gibson
2001 Recency in verb phrase attachment. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27, 574–590.

Pickering, Martin and Guy Barry
1991 Sentence processing without empty categories. *Language and Cognitive Processes* 6, 229–259.

Prince, Ellen
1981 Toward a typology of given-new information. In Cole, Peter (ed.), *Radical Pragmatics*. New York, NY: Academic Press, 223–255.

Rohde, Douglas L. T.
2002 A connectionist model of sentence comprehension and production. Unpublished Ph.D. thesis, Carnegie Mellon University.

Selkirk, Elizabeth O.
1984 *The Relation between Sound and Structure*. Cambridge, MA: MIT Press.

Stabler, Edward P.
1994 The finite connectivity of linguistic structures. In Clifton, Charles Jr., Lyn Frazier, and Keith Rayner (eds.), *Perspectives on Sentence Processing*. Hillsdale, NJ: Erlbaum, 303–336.

Stolz, Walter S.
1967 A study of the ability to decode grammatically novel sentences. *Journal of Verbal Learning and Verbal Behavior* 6, 867–873.

Stromswold, Karin, David Caplan, Nathaniel Alpert, and Scott Rauch
1996 Localization of syntactic comprehension by Positron Emission Tomography. *Brain and Language* 52, 452–473.

Tabor, Whitney, Cornell Juliano, and Michael K. Tanenhaus
1997 Parsing in a dynamical system: An attractor-based account of the interaction of lexical and structural constraints in sentence processing. *Language and Cognitive Processes* 12, 211–272.

Tanenhaus, Michael K. and John C. Trueswell
1995 Sentence comprehension. In Miller, Joanne L. and Peter D. Eimas (eds.), *Speech, Language and Communication*. San Diego, CA: Academic Press, 217–262.

Trueswell, John C., Michael K. Tanenhaus, and Susan M. Garnsey
1994 Semantic influences on parsing: use of thematic role information in syntactic disambiguation. *Journal of Memory and Language* 33, 285–318.

Wanner, Eric and Michael Maratsos
1978 An ATN approach in comprehension. In Halle, Morris, Joan Bresnan, and George Miller (Eds.), *Linguistic Theory and Psychological Reality*. Cambridge, MA: MIT Press, 119–161.

Waters, Gloria S., David Caplan, and Nancy Hildebrandt
1987 Working memory and written sentence comprehension. In Coltheart, Max (ed.), *Attention and Performance*: Volume 12, *The Psychology of Reading*. Hillsdale, NJ: Erlbaum, 531–555.

Watson, Duane and Edward Gibson
In press The relationship between intonational phrasing and syntactic structure in language production. *Language and Cognitive Processes.*

Yngve, Victor H.
1960 A model and a hypothesis for language structure. *Proceedings of the American Philosophical Society* 104, 444–466.

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.

(V7 12/5/05 12:05) WDG/G J-1300 Cognitive Linguistics, 16:2 PMU:(IDP(CKN)W)12/5/2005 Times_M (0).3.04.05 pp. 313–354 cogl_16-2_313-354 (p.