Anticipating Explanations in Relative Clause Processing

Hannah Rohde, Roger Levy, Andrew Kehler
Department of Linguistics
University of California, San Diego

September 16, 2008

Abstract

In support of a model of language comprehension in which pragmatic biases are integrated with syntactic processing, we show that expectations about upcoming discourse continuations influence the resolution of local structural ambiguity. An off-line sentence-completion study and an on-line self-paced reading study examined readers' expectations for high/low relative clause attachments following implicit-causality and non-implicit-causality verbs (*John detests/babysits the children of the musician who...*). In the off-line study, the widely reported low-attachment preference for English is observed in the non-implicit causality condition, but this preference gives way to more high attachments in the implicit causality condition in cases in which (i) the verb's causally implicated referent occupies the high-attachment position and (ii) the relative clause provides an explanation for the event described by the matrix clause (e.g., *...who are arrogant and rude*). In the on-line study, a similar preference for high attachment emerges in the implicit causality context, whereas the low-attachment preference is consistent elsewhere. These results suggest that comprehenders construct discourse contexts dynamically during sentence processing, using available pragmatic cues mid-sentence to generate expectations about the structural analysis of the remainder of the sentence.

Keywords:
Discourse Processing; Relative Clause Attachment Ambiguity; Implicit Causality

1 Introduction

A foundational question in psycholinguistics asks how it can be that comprehenders, working merely from linear sequences of words, are able to recover the array of latent relationships necessary to recover the speaker's intended underlying message. At the sentence level, for instance, a considerable body of research has focused on the comprehender's need to infer intrasentential *structural* relationships between underlying constituents in building a syntactic representation of an utterance. When processing (1) below, for example, a comprehender infers that *his coworkers* stands in the structural relation of objecthood with *detests*, that *obnoxious* is coordinated with *rude*, and so on. Likewise, in a largely independent line of research, considerable work has also examined how relationships are recovered above the level of the sentence, that is, how comprehenders establish *discourse* dependencies. When processing (1), for instance, the comprehender is likely to infer that the phrase *they* corefers with the coworkers mentioned in the first sentence, and further that the second sentence stands in a particular discourse relationship with the first sentence, specifically that of providing an *explanation* for why John detests these coworkers.

(1) John detests his coworkers. They are obnoxious and rude.

Considerable progress has been made in our understanding of how these latent relationships are recovered at both levels. This progress, however, has produced models of sentence and discourse processing that bear strikingly little resemblance to one another. When considering processing mechanisms at both levels as an ensemble, therefore, several key questions are naturally brought to the fore:

1. Do syntactic (i.e. intrasentential) and discourse (i.e. intersentential) comprehension processes happen serially, or are they interleaved?
2. Is there information flow from processing at the higher level (discourse) that influences decisions at the lower (intrasentential) level, and if so, how is it best characterized?
3. Are syntactic and discourse processing inherently different, or will common processing mechanisms emerge when viewed from an appropriate perspective?

In this paper, we attempt to address these questions through two experiments focusing on the juxtaposition of a quintessentially structure-

oriented, intrasentential comprehension process—the resolution of syntactic attachment of relative clauses (RCs)—with what is generally taken to be a quintessentially meaning-oriented, discourse-level comprehension process—the establishment of the informational relations that hold between clauses. Corresponding to the three questions just posed, we argue that (1) syntactic and discourse comprehension processes are interleaved, (2) that discourse processing can indeed influence the inference of intrasentential structural relations, and (3) when cast within a suitable evidential and/or probabilistic model, these processes and their interdependencies can be naturally captured by appealing to common interpretation mechanisms.

1.1 Structural Disambiguation: Relative-Clause Attachment

In this paper we focus on a particular aspect of syntactic processing that has attracted considerable attention in the field: relative clause (RC) attachment ambiguity, which arises in sentences such as (2).

(2) Someone shot <u>the servant</u> of <u>the actress</u> who was on the balcony.
NP1 (high) NP2 (low)

Example (2) contains an RC (*who was on the balcony*) whose attachment site is ambiguous: The RC can be interpreted to modify one or the other of the two noun phrases (NPs) in the preceding complex NP (*the servant of the actress*). The RC is said to attach HIGH if it is interpreted to modify NP1 (*the servant*), which occupies the higher position in the syntactic structure. It is likewise said to attach LOW if it modifies NP2 (*the actress*), the possessor NP within the complex NP. Following an early account that predicted a low-attachment preference through the principle of Late Closure (Frazier, 1978), a low-attachment preference in contexts like (2) in English is now widely accepted as the default, having been confirmed in off-line studies with questionnaires and completion tasks and in most online studies (Frazier & Clifton, 1996; Carreiras & Clifton, 1999; Fernandez, 2003; but see also Traxler, Pickering, & Clifton, 1998).

This preference is not universal across languages, however, and starting with Cuetos and Mitchell (1988), a considerable body of research has emerged on RC attachment preferences within and across languages (see reviews in Cuetos, Mitchell, & Corley, 1996; Mitchell & Brysbaert, 1998; Desmet, Brysbaert, & De Baecke, 2002; and Papadopoulou & Clahsen,

2006). Studies involving constructions such as (2) in languages including Spanish, French, German, and Dutch (among others) have found evidence of a bias towards a high-attachment interpretation (Cuetos & Mitchell, 1988; Zagar, Pynte, & Rativeau, 1997; Hemforth, Konieczny, & Scheepers, 2000; Brysbaert & D.C., 1996; a.o.). The lack of a universal attachment preference has thus been problematic for theories of sentence processing which posit crosslinguistic syntactic constraints and strategies. Moreover, it has also been shown that within a single language, attachment preferences vary with lexical properties, such as animacy (Desmet, Baecke, Drieghe, Brysbaert, & Vonk, 2006), referentiality (Gilboy, Sopena, Clifton, & Frazier, 1995), and even specific lexical head (Desmet & Gibson, 2003). Evidential accounts such as the competition-integration model (MacDonald, 1994; Spivey & Tanenhaus, 1998; McRae, Spivey-Knowlton, & Tanenhaus, 1998) and probabilistic models (Jurafsky, 1996; Narayanan & Jurafsky, 1998, 2002; Crocker & Brants, 2000; Hale, 2001; Levy, 2008) are compatible with the idea that a host of factors involving multiple information sources conspire to determine attachment preferences in any particular sentence. In this view, so-called "default" attachment preferences in a given language are simply the consequence of the distribution of relevant information-source particulars in the language.

There has also been a limited amount of work on the role of discourse processing in on-line sentence comprehension. Perhaps the earliest account relating discourse-level processing and syntactic disambiguation was the referential theory (Crain & Steedman, 1985; Altmann & Steedman, 1988), which focused on the ability of NP postmodifiers to restrict the domain of possible reference of the modified NP. According to this theory, an NP with a restrictive postmodifier such as *the horse raced past the barn* can, in a typical discourse context, be taken not only to presuppose the existence of a horse that was raced past a barn, but also to conversationally implicate (Grice, 1975) the existence of a horse that was not. This implicature results, according to Gricean reasoning, from that fact that if there were only one horse in the context, the speaker would be expected to have chosen the less informative and less prolix NP *the horse*. As a result, when there is ambiguity as to whether material after a given NP constitutes a postmodifier of that NP, the postmodifier analysis should be favored when the preceding context implies that the NP would otherwise be referentially ambiguous. It has been argued by some researchers that it is implausible to expect that the inference of a conversational implicature could affect on-line syntactic comprehension, e.g.:

> To make a conversational implicature, a listener must have already parsed the sentence, assigned it its literal interpretation, realised that additional inferences must be added to make it conform to the Gricean maxim, and determined what these inferences are. Such activity could not reasonably affect the initial steps of parsing. (Clifton & Ferreira, 1989)

Nevertheless, it has since been shown by Ni, Crain, and Shankweiler (1996) and Sedivy (2002) that invoking implicit referential contrast sets can affect the main-verb/reduced-relative ambiguity in classic garden-path sentences. Likewise, one might expect that manipulating the number of referents of a given type in preceding discourse context can affect RC attachment preferences. In (2), for example, if it is clear that there is more than one actress in the discourse context, then the use of the definite NP *the actress* is not felicitous. As a result, one might predict a low-attachment bias in a context with multiple actresses as in (3) below, but a high-attachment bias in a context with multiple servants as in (4):

(3) There was a servant working for two actresses. Someone shot the servant of the actress who was on the balcony.

(4) There were two servants working for a famous actress. Someone shot the servant of the actress who was on the balcony.

There has been evidence supporting this prediction from research using off-line methodologies on a variety of languages (French: Zagar et al., 1997; Dutch: Desmet, de Baecke, & Brysbaert, 2002; Greek: Papadopoulou & Clahsen, 2006). The evidence for on-line effects has been more mixed. Papadopoulou and Clahsen (2006) reported significant effects using self-paced reading and van Berkum, Brown, and Hagoort (1999) study found on-line effects with ERPs, but the Zagar et al. (1997) and Desmet et al. (2002) studies found no significant effects using eye-tracking.[1]

[1]The van Berkum et al. study used a manipulation that involved a temporary ambiguity between an RC and a complement clause (*David told the girl that...*), rather than an attachment ambiguity for an RC. The conclusion is the same, however, in that contexts that required that the reference of the critical NP be disambiguated caused comprehenders to favor the RC analysis.

1.2 The Construction of Coherent Discourse

Of course, language understanding involves more than sentence processing and reference resolution: Comprehenders must also link together larger linguistic units to construct a coherent discourse. Theories of discourse coherence (Hobbs, 1979; Kehler, 2002; a.o.) posit the existence of COHERENCE RELATIONS that comprehenders must infer to fully interpret a discourse. For instance, on its most natural interpretation, the clauses in example (5) below are related by an Explanation relation: The state described in the second sentence is interpreted as a cause of the event expressed in the first sentence. Upon hearing (5), a comprehender might therefore be surprised to discover that John has no intention of visiting his family, even though the passage never states that he does.

(5) John took a train to San Francisco. He has family there.

Passage (6) below, on the other hand, does not involve causality on its most natural interpretation. Instead, it is characterized by an Elaboration relation, in which the second clause is understood to elaborate the first.

(6) John took a train to San Francisco. He boarded at the local Caltrain station and enjoyed the scenery on the way to the city's 4th and King Street station.

As a result, the comprehender infers that the two sentences jointly describe a single event rather than two. Explanation and Elaboration are but two relations from a larger inventory which have been proposed within theories of discourse coherence by Hobbs, Kehler, and others.

Recent studies by Rohde, Kehler, and Elman (2006; 2007) and Kehler, Kertz, Rohde, and Elman (2008) provide evidence that comprehenders generate expectations concerning what coherence relations are likely to ensue based on the current context, and argue that any successful account of pronoun interpretation necessarily must incorporate those expectations. If such coherence-driven expectations could similarly be shown to influence *local syntactic processing* decisions – particularly in a scenario in which the felicity of the utterance was not at stake, as it was in the aforementioned research on RCs as referential restrictors – it would constitute a fairly radical demonstration of the range of information sources that are brought to bear in on-line syntactic comprehension. The goal of the present paper is to do exactly this, by examining comprehenders' behavior when processing sentences involving RC attachment ambiguity such as (2), but in cases in

which the matrix and relative clauses share both a syntactic and a discourse relation, and hence in which syntactic and discourse processing might be expected to interact. Our design brings together three independent observations involving pragmatic interpretation:

i. Matrix clauses that contain so-called *implicit causality* (IC) verbs create a strong expectation that an ensuing clause will contain an explanation of the eventuality denoted by the matrix clause.

ii. An RC can be used to provide an explanation of the matrix clause.

iii. Certain IC verbs, specifically *object-biased* IC verbs, create an expectation that the locus of an ensuing explanation is likely to be the verb's direct object.

On an expectation-based theory in which discourse processing occurs simultaneously with syntactic comprehension and can influence on-line attachment decisions, these three factors should conspire to generate a bias toward high attachments for RCs in a particular class of ambiguous sentences. We explain why in the following subsections.

1.3 Observation 1: Implicit Causality Creates an Expectation for an Explanation

Our first observation is that a clause containing a certain type of verb, of which *detest* in (7) below is an example, typically creates a strong expectation that it will be followed by an explanation of the eventuality that the clause denotes.[2] Garvey and Caramazza (1974) coined the term IMPLICIT CAUSALITY (IC) to describe the verbs of interest here, for which one of the role participants "is implicated as the assumed locus of the underlying cause of the action or attitude" (see Section 1.5 below). Kehler et al. (2008) conducted a story continuation experiment using IC and non-IC verbs as in (7-8).

(7) John detests Mary. _____. [IC VERB]

[2]Unless, of course, the cause of the eventuality has already been provided previously in the discourse (cf. Simner & Pickering, 2005).

(8) John babysits Mary. _____. [NON-IC VERB]

Kehler et al. found that context sentences with IC verbs (7) yielded significantly more Explanation coherence continuations (60%) than context sentences with non-IC verbs (8) did (24%). At an intuitive level, the lexical semantics of verbs like *detest* appear to lead the comprehender to ask *Why?* in a way that verbs like *babysit* do not.

1.4 Observation 2: RCs can Provide an Explanation

Restriction of reference is only one of the possible functions of an RC. Our second observation brings another role that RCs play to the fore, one that has not, to our knowledge, previously been utilized in psycholinguistic work: An RC can serve the additional pragmatic function of providing an explanation of the event described in the matrix clause. For instance, on a natural interpretation, the RC in (9) not only restricts the reference of *the kindergartner*, but also carries an implicature that the student's tardiness is the reason for the reproach.

(9) The teacher reproached the kindergartner who always shows up late.

This implicature depends on the world knowledge that persistent lateness is a plausible reason to be reproached,[3] and thus, crucially, this implicature is not triggered by a property internal to the RC itself. That is, the same RC in a different context can merely serve to modify the preceding NP if the world knowledge necessary to support a causal interpretation does not exist, as in (10).

(10) The teenager babysits the kindergartner who always shows up late.

Likewise, it is possible to use a different RC with the matrix clause of (9) such that a causal relationship is not inferred:

(11) The teacher reproached the kindergartner who reminds me of Harry Potter.

[3]The fact that this is an implicature and not an entailment is demonstrated by the fact that it is cancelable, as in (i).

(i) The teacher reproached the kindergartner who always shows up late.
The kindergartner had forgotten his lunch money for the third day in a row.

In (11), for instance, the comprehender is not normally led to the inference that reminding the speaker of Harry Potter is the reason for the reproach.

This kind of intrasentential inference process mirrors the process that comprehenders use to establish discourse coherence across sentences, as already described with respect to (5) and (6). That is, the causal reasoning that links the clauses in (9) is similar to that used in establishing an Explanation relationship between the two sentences in (12). Similarly, the lack of inferred causation in (10) mirrors the non-causal Elaboration relation that holds between the sentences in (13).

(12) The teacher reproached the kindergartner. The kindergartner always shows up late.

(13) The teenager babysits the kindergartner. The kindergartner always shows up late.

Paralleling our observations regarding (9) and (10), the different relationship between the sentences in (12) and (13) stems from the fact that lateness provides a plausible reason to reproach someone but not a plausible reason to babysit someone.

1.5 Observation 3: Implicit Causality Influences Next Mention Expectations

Our third observation is already well established in the literature: That IC verbs impute causality primarily to one of the participants of the eventuality they denote, and thus create a strong bias toward mentioning that participant in any ensuing explanation of that eventuality (Garvey & Caramazza, 1974; Brown & Fish, 1983; Au, 1986; McKoon, Greene, & Ratcliff, 1993). Some verbs, like *detest* in (7) and (14), are OBJECT-BIASED, meaning that it is the direct object that comprehenders expect to hear mentioned again in connection with an explanation: If John detests Mary, then the cause is likely to originate from a property of Mary. On the other hand, verbs like *annoy* in (15) are SUBJECT-BIASED: If John annoys Mary, then the cause presumably originates from a property of John. Non-IC verbs, such as *babysit* in (16), are reported to have weaker and less consistent biases.

(14) John detests Mary because... [OBJECT-BIASED IC VERB]
...she is rude and arrogant.

(15) John annoys Mary because... [SUBJECT-BIASED IC VERB]
...he is rude and arrogant.

(16) John babysits Mary because... [NON-IC VERB]
...he needs the money. / ...she is too young to be left alone.

These next-mention biases have commonly been measured using story completion tasks with *because*-prompts like those in (14-16) (typical continuations are listed below each prompt). Kehler et al. (2008) also found that a similar pattern of biases emerges for prompts like (14-16) even without the *because* connective when, crucially, only the subset of continuations in which an Explanation relation is operative are analyzed. That is, the bias is tied specifically to causes per Explanation relations, and is not dependent on the conjunction *because*.

1.6 Implicit Causality and RC Attachment

We are now ready to introduce the manipulation of interest, illustrated in the matrix clauses of (17-18), with sample RCs shown below in (a-b):

(17) John babysits the children of the musician who ...
a. ...lives in La Jolla. [low]
b. ...are students at a private school. [high]

(18) John detests the children of the musician who...
a. ...lives in La Jolla. [low]
b. ...are arrogant and rude. [high]

The matrix clauses in these examples differ only in the verb: *detests* is an object-biased IC verb, whereas *babysits* is non-IC. The default low-attachment preference attested in English predicts uniform biases across (17-18); for instance, in a passage completion experiment, we would expect to see more low-attaching completions (like 17a and 18a) than high-attaching ones (like 17b and 18b). We would likewise expect the RC verb *is* in (17a) and (18a) to be easier to process on-line than *are* in (17b) and (18b).

However, if comprehenders utilize the coherence-based pragmatic information described in Sections 1.3-1.5, we might expect to see a difference in these measures: The bias toward high attachments should be greater in (18) than in (17). The reasoning goes as follows. We saw in Section 1.3

that clauses containing IC verbs, such as *detest* in (18), create a strong expectation for an ensuing explanation. We then saw in Section 1.4 that this explanation can be delivered via an immediately-following RC. Finally, we saw in Section 1.5 that if an explanation were to follow, that an object-biased verb like *detest* creates a strong expectation that the explanation will be about the verb's direct object, which, in cases like (18), is *the high attachment site for the relative clause*. That is, we conjecture that upon hearing the first clause of (18), the three types of pragmatic knowledge just discussed will conspire to contribute a bias toward expecting the RC to provide an explanation that expresses a predication about, and therefore attaches (high) to, *the children*. Whereas we would expect this same system of coherence-based knowledge to exert parallel influences in examples like (17), the results are expected to be considerably weaker, since non-IC verbs do not create a strong expectation toward an ensuing explanation (Section 1.3), nor is the direct object favored to be the locus of such explanations (Section 1.5).

Crucially, this reasoning only goes through by making a significant assumption: That comprehenders, having processed the initial part of a matrix clause, implicitly integrate all three of these types of pragmatic knowledge and biases and use them when making a syntactic attachment decision in mid-sentence. We submit, therefore, that a difference in the comprehender's behavior in processing passages like (17) and (18) in the predicted direction would constitute a novel and significant piece of evidence in support of an expectation-based model of syntactic processing that incorporates a broad set of cross-modular information sources.

The remainder of this paper presents two experiments that test this hypothesis. If comprehenders are indeed using coherence-driven biases mid-sentence, then one would expect to see effects with respect to the types of RC completions they generate following IC matrix-clause verbs (Experiment 1, a sentence completion study). If these biases contribute to on-line processing, then one would expect to see processing difficulty associated with those RC attachments that violate the biases introduced by the matrix-clause verb (Experiment 2, a self-paced reading study).

2 Experiment 1: Sentence Completion Study

The first experiment uses an off-line sentence-completion task to investigate our main hypothesis—namely, whether object-biased IC verbs influ-

ence sentence completions for ambiguously-attached RCs in ways consistent with the three observations above. The prompt pair in (19) involves a single manipulation (henceforth referred to as 'verbtype'): whether the main clause has an IC verb (19b) or a non-IC verb (19a).

(19) a. John *babysits* the children of the musician who ... [NON-IC]
b. John *detests* the children of the musician who ... [IC]

According to our hypothesis, comprehenders will utilize discourse-level expectations about upcoming coherence relations mid-sentence. In that case, the presence of an IC verb in (19b) should elicit a greater number of completions that explain the matrix clause event than elicited by the non-IC variant (19a), and, in light of the fact that *detests* is object-biased, these explanations should tend to attach high. If, on the other hand, such expectations are not utilized during syntactic processing, RCs following both verbtypes should attach low.

2.1 Methodology

Participants

Fifty-two monolingual English-speaking undergraduates at UC San Diego participated in the study for course credit in Linguistics courses.

Materials

Stimuli consisted of twenty-one pairs of sentences differing only in the matrix verb, as in (19); the complete stimuli can be found in the appendix. The subject of the sentences was a proper name, and the direct object consisted of a complex NP containing two NPs connected by the genitive marker *of*. Both NPs denoted human referents so that the relative pronoun *who* could plausibly be used to modify either NP. In order to make disambiguation easier for the judges, the complex NP consisted of a singular NP and a plural NP so that number agreement on the embedded verb could be used to assess the intended attachment site of the RC. The order of singular and plural was balanced across stimuli (10 singular-plural, 11 plural-singular).

The verb in the matrix clause was either an object-biased IC verb or a non-IC verb. IC verbs were selected from two lexical semantic categories that Levin (1993) labels 'psych' and 'judgment' verbs. Non-IC verbs were adapted from those used by McKoon et al. (1993) in their study of IC

and pronoun interpretation. For our stimuli, psych verbs appeared in the present tense since they describe non-eventive states (e.g., *detest, adore*), whereas judgment verbs appeared in the simple past (e.g., *scolded, praised*). Each pair of IC and non-IC verbs was matched for tense as in (19).

In addition to the experimental items, the experiment included twenty-one fillers and twenty-one additional stimuli for an unrelated experiment, pseudorandomized to create eight lists.[4] The additional fillers consisted of sentences with non-transfer verbs and a variety of prompts as well as sentences with complex NPs and unambiguous RC prompts.[5] Half of the unambiguous RC fillers enforced a low attachment and half enforced a high attachment.

Procedure

Sentence completions were collected via a web-based interface that participants could access from their own computer. Each item was presented on a page by itself with a text box in which participants were instructed to write their completion. The entire experiment took roughly thirty minutes, but participants were encouraged to have an hour available so that the experiment could be completed in one session. (Participants could leave and return at a later time by identifying themselves with an ID number.) They were instructed to imagine a natural sentence completion for each prompt, writing the first completion that came to mind and avoiding humor.

Evaluation and Analysis

Two trained judges—the first author of this paper and an undergraduate Linguistics student—annotated all responses for the type of RC ('restriction/modification' or 'explanation') and the intended attachment site (low or high). An RC was labeled 'restriction/modification' if it provided additional information about one of the nouns without providing additional information about the event in the matrix clause *per se*. Explanation RCs, on the other hand, required that a causal link be inferred between the information conveyed by the matrix clause and the information in the RC. Disagreements were resolved through discussion. The intended attachment site was assessed in light of the matrix clause context and the elicited RC.

[4]The stimuli for the interleaved experiment contained sentences with transfer-of-possession verbs followed either by a full stop and a completion prompt or a full stop and an ambiguous pronoun prompt: *Matt passed a sandwich to David. (He) ...*

[5]The potential ambiguity in these cases was resolved by the relative pronoun *who* in contexts in which only one of the two nouns in the complex NP was human.

An RC was excluded from the analysis if at least one coder assessed its attachment height to be ambiguous, as well as in the few cases in which the coders disagreed.

The sample completions in (20) and (21) show examples of modification and explanation RCs that attach either low or high. (Examples of actual completions corresponding to each of (20a-d), (21a-d) can be found in the appendix.) Note that the sample modification RCs listed do not differ between the non-IC and IC contexts (20a-b, 21a-b) because such RCs need not reflect information about the event described in the matrix clause. Explanation RCs, on the other hand, provide an explanation of the matrix clause event and therefore are shown varying with the matrix clause (20c-d, 21c-d).

(20) NON-IC VERB: John babysits the children of the musician who ...
  a. ...lives in La Jolla. [MOD - LOW]
  b. ...are in elementary school. [MOD - HIGH]
  c. ...works a late shift every night. [EXP - LOW]
  d. ...are left home on Friday nights. [EXP - HIGH]

(21) IC VERB: John detests the children of the musician who ...
  a. ...lives in La Jolla. [MOD - LOW]
  b. ...are in elementary school. [MOD - HIGH]
  c. ...encourages their 3am drum solos. [EXP - LOW]
  d. ...are arrogant and rude. [EXP - HIGH]

As (20) and (21) show, both verb types can be followed by an explanation RC or a modification RC, and neither RC type enforces a particular attachment. The hypothesis is that the combination of coherence biases and next-mention biases will render high-attaching explanation RCs more likely following IC verbs than non-IC verbs: Completions like (21d) will be more common than (20d). The low-attaching modification RCs, on the other hand, are predicted to be more expected following non-IC verbs than IC verbs: Completions like (20a) will be more common than (21a).

We conducted analyses of variance on the assessed RC completion types and on the assessed attachment sites to test for a main effect of verbtype. Because these measures involve examining proportions of categorical outcomes, we first applied an arcsine transformation (Sokal & Rohlf, 1995) to the percentages of each outcome. For clarity of presentation, we present means as raw proportions. The observed RC types and RC attachment sites were also modeled using mixed-effects multinomial logistic regressions with

random subject-specific and item-specific intercepts (Jaeger, *in press*). We report the coefficient estimate and p-value (based on the Wald Z statistic; Agresti, 2002) for the factor verbtype in models fitted to the observed RC completion types and to the observed RC attachments.

2.2 Results

As predicted, IC verbs yielded significantly more explanation-providing RCs than non-IC verbs (main effect of verbtype on RC type: $F_1(1,51)=292.22$, $p<0.001$; $F_2(1,20)=87.665$, $p<0.001$) and significantly more high-attaching RCs (main effect of verbtype on attachment height: $F_1(1,51)=27.158$, $p<0.001$; $F_2(1,20)=6.8475$, $p<0.05$).[6] In the logistic regressions, verbtype was a significant factor for modeling the binary outcome of RC type (whether an RC provided an explanation: $\beta=4.530$, $p<0.001$) and the binary outcome of attachment height (whether the RC attached high: $\beta=0.803$, $p<0.005$).

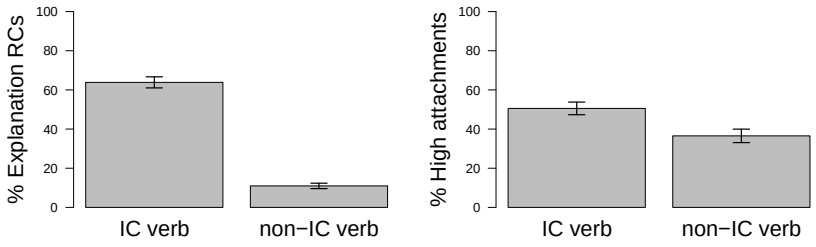

a. RC type b. RC attachment

Figure 1: Percentages of explanation RCs and high attachments

As can be seen in Figure 1a, verbtype affected the types of RCs participants produced: More than half (63.9%) of the RCs following IC verbs provided an explanation of the event in the matrix clause, while only a

[6]This analysis represents a conservative analysis in which an RC was excluded if at least one coder assessed it as ambiguous (22.5% of the total number of RC completions). The results remain significant if RCs are included if at least one coder assigned a non-ambiguous interpretation (Attachment: $F_1(1,51)=53.52$, $p<0.001$; $F_2(1,20)=8.1197$, $p<0.01$; RC type: $F_1(1,51)=356.07$, $p<0.001$; $F_2(1,20)=96.407$, $p<0.001$). We restricted the analysis to subject-extracted RCs since object-extracted RCs made up fewer than 1% of the total completions, and their inclusion does not affect the overall results.

small proportion (11.0%) of RCs following non-IC verbs provided an explanation. Figure 1b shows that the pattern of RC attachment differs by verbtype as well. In the non-IC context, only 36.5% of the unambiguous elicited completions contained a high attachment, which matches the reported low-attachment preference for English. In the IC context, the low-attachment preference disappears with 50.6% of unambiguous completions containing a high attachment. All of these proportions represent subject means and are shown in Figure 1 with error bars for standard error of the mean.

Figure 2 shows the results broken down by verbtype and RC type. Regardless of verb type, explanation-providing RCs had a higher incidence of high attachment (66.3% for IC verbs, 47.0% for non-IC verbs) than RCs that did not provide explanations (26.0% for IC verbs, 35.9% for non-IC verbs). Pairwise comparisons of explanation vs. modification RCs were significant in the IC condition ($F_1(1,49)=35.351$, $p<0.001$; $F_2(1,20)=36.419$, $p<0.001$; logistic regression: $\beta=2.9391$, $p<0.001$) but insignificant in the non-IC condition ($F_1(1,32)=0.4819$, $p=0.49$; $F_2(1,8)= 0.6325$, $p=0.45$; logistic regression: $\beta=0.6246$, $p=0.15$). The greater incidence of high attachment for explanation RCs in the IC condition follows intuitively from the fact that the IC verbs used in this experiment impute causality to their direct object (the high NP). Even for non-IC verbs, it is unsurprising that explanation RCs might tend to attach to the high NP more than modification RCs, since the high NP is a direct participant in the matrix-clause event being explained, whereas the low NP is not (though see Section 2.3). Unlike the case for IC verbs, however, the difference in the proportion of high attachments was not significant.

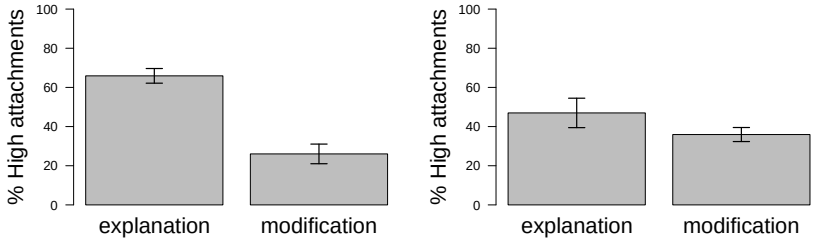

a. RCs following IC contexts b. RCs following non-IC contexts

Figure 2: Percentage of high attachments by verbtype and by RC type

2.3 Discussion

The sentence completion study was designed to test a model of sentence processing that integrates discourse-level biases with those governing the resolution of local syntactic ambiguity. As predicted, IC verbs yielded significantly more explanation-providing RCs and significantly more high-attaching RCs compared with non-IC verbs. These results have implications for models of both sentence processing and discourse processing.

First, with respect to models of RC processing, the pattern of high attachments following IC verbs provides evidence that any model relying primarily on a default low-attachment preference will be insufficient. Further, the fact that RC attachment height biases are dependent on the discourse relation between the RC and the matrix clause suggests that models need to incorporate factors beyond merely those derived from properties of the NPs themselves; in particular, they must incorporate the type of coherence-driven biases revealed here as well.

Second, the results are likewise relevant to models of discourse processing that aim to uncover the factors that guide comprehenders' interpretation of discourse and the nature of the linguistic elements over which such factors operate. The results confirm Kehler et al.'s (2008) conclusion that comprehenders are sensitive to two types of biases invoked by IC verbs: a clause-level coherence bias toward upcoming explanations and an entity-level next-mention bias conditioned on the presence of an explanation relation. These results go beyond previous work in demonstrating that intersentential coherence relations like Explanation can be inferred to hold intrasententially and that RCs can embody such relations. If an RC can provide an explanation of an event in the matrix clause, then the process of linking together elements of a discourse into a larger coherent structure cannot be cast as a process that occurs only after a sentence has been digested in its entirety.

Although the experiment confirmed our predictions, a closer analysis of the elicited sentence completions revealed that several factors likely conspired to reduce the effect from what might otherwise have been found. Recall that in our model, the way in which the preceding context influences attachment preferences is complex, depending simultaneously on coherence biases and the next-mention biases given the operative coherence relation and the presence of an object-biased IC verb. In this regard, our analysis identified two subpatterns of behavior that are not apparent in the

aggregate effects: one concerning the bias towards explanations and one concerning the attachment biases.

Regarding the coherence bias, we found that some verbs that have been classified in the literature as non-IC actually yielded a larger number of explanation-providing RCs than some IC verbs. For example, the verb *watch*, which McKoon et al. (1993) included in a non-IC condition, yielded 46.2% explanations — more than some IC verbs such as *like* (26.7% explanations) and *value* (22.7% explanations). The appendix contains the full list of the percentages of explanation RCs that each verb elicited. Caramazza et al. (1977) previously commented that the next-mention biases of IC verbs lie along a continuum; here we find that the same is true for their biases towards ensuing coherence relations as well. We therefore would have expected a stronger effect if the IC verbs used had uniformly stronger biases towards explanations than their non-IC counterparts.

Regarding the attachment biases, we found that the presence of both an IC verb and an explanation-providing RC are not in themselves sufficient to yield a high-attachment preference; the relationship between the NPs in the complex NP also has an influence. The two IC items in (22) show how specific complex NPs can shift the bias to yield more high-attaching or more low-attaching RCs.

(22) a. Alan punished the accountant of the businessmen who ...
b. Bill congratulated the teacher of the second graders who ...

Example (22a) yielded a large proportion of explanation-providing RCs (85.7%), and those RCs consistently attached to the higher NP (100%). Example (22b) also yielded many explanation-providing RCs (81%), but in this case, the RCs tended to attach low (only 29% high attachment). Example (22b) differs from (22a) in that the lower NP in (22) refers to a set of individuals (NP2: *the second graders*) who are under the care or responsibility of the individual referenced by the direct object (NP1: *the teacher*). Because of the possibility of attributing responsibility to the NP1 referent for the NP2 referent's behavior, the explanation-providing RCs could plausibly mention either the guardian or the guardian's wards.[7] As such, avoiding such relationships between NP1 and NP2 would presumably also have yielded a stronger effect.

[7]Similarly behaving items included *scold the landlady of the actors who...*, *detest the father of the students who...*; and *pity the bodyguards of the celebrity who...*

Although we found a significant effect in Experiment 1, the results thus far are restricted to an off-line completion task. If, as we hypothesize, the coherence-driven biases that emerge are indeed deployed mid-sentence, then one would expect to see effects in comprehenders' incremental processing in a self-paced reading time experiment. The goal of Experiment 2 is to test this hypothesis.

3 Experiment 2: Self-Paced Reading Study

Consistent with our hypothesis, Experiment 1 demonstrated that the discourse-coherence biases introduced by IC verbs affect expectations about the content and attachment level of upcoming relative clauses. The goal of Experiment 2 was to test a further aspect of our hypothesis: That these expectations are deployed rapidly in on-line comprehension, and crucially, are active before comprehenders have been exposed to complete clauses. That is, we predict that inferences about intraclausal coherence relations are *fully incremental*, and can therefore affect local syntactic disambiguation.

To test this aspect of our hypothesis, we conducted a moving-window self-paced reading study to examine the timecourse of the biases found in the sentence completion study. We adapted the stimuli from Experiment 1 to create a 2×2 design, varying the verbtype and the RC attachment height, as in (23) and (24). Underscores connect words presented together as a single region in the study.

(23) NON-IC MATRIX: John babysits the_children of the_musician ...
  a. [LOW ATTACHMENT] ... who *is* generally arrogant and rude.
  b. [HIGH ATTACHMENT] ... who *are* generally arrogant and rude.

(24) IC MATRIX: John detests the_children of the_musician ...
  a. [LOW ATTACHMENT] ... who *is* generally arrogant and rude.
  b. [HIGH ATTACHMENT] ... who *are* generally arrogant and rude.

In (23) and (24), the attachment height of the RC is signaled by the number agreement information of the finite verb (e.g., the verb *is* in (23a) agrees in number with NP2, the NP at the low-attachment site). On our hypothesis, just before encountering the finite verb, the comprehender should already have formed expectations about the attachment height of the RC initiated by *who*, based on the likelihood (conditioned by the preceding context) that the RC may express an explanation for the matrix clause event, and

on the likelihoods of high attachment if the RC expresses an explanation or only a modification, as detailed in Sections 1.3 through 1.5. The degree to which these expectations match the RC attachment height signaled by the finite verb should then determine the difficulty of processing the finite verb. This finite verb therefore constitutes the CRITICAL region of the experiment, with a low-attachment bias expected to emerge for non-IC verbs ((23a) easier than (23b)), but predicted to be reduced, neutralized, or even reversed following IC verbs ((24a) not as much easier than, equally difficult as, or harder than (24b)). On the other hand, if our hypothesis is incorrect and integration into the larger discourse structure occurs only after the entire sentence has been processed, then the default low-attachment bias should hold across the board ((23a,24a) easier than (23b,24b)). Because differences in processing difficulty in self-paced reading often show up a region or two downstream of the critical region, especially when the critical region is short as it is here (e.g., Mitchell, 1984), the immediately postcritical word (*generally* in (23) and (24)) was always an adverb chosen to be relatively non-indicative of attachment height. This word and the immediately following word (*arrogant* in (23) and (24)) are the SPILLOVER regions for this experiment. On the null hypothesis of discourse-insensitive attachment preference, we expect a main effect of attachment height on reading times at the critical and/or spillover regions; on the alternative hypothesis that discourse coherence relations affect the online formation of attachment preferences, we expect an interaction between attachment height and verbtype on reading times at the critical and/or spillover regions.

Participants

58 monolingual English speakers participated in the experiment for credit in Linguistics and Psychology courses.

Materials

Each experimental item consisted of a matrix clause with a proper name, a verb, and a complex NP direct object, followed by a temporarily ambiguous RC, as in (23) and (24). The complex NP contained a singular NP and a plural NP so that number agreement on the critical embedded verb (*is/are* in (23) and (24)) served to disambiguate the attachment site of the RC. The order of singular and plural in the complex NP was balanced across stimuli so that high attachment was signaled with plural agreement for half the items and with singular agreement for the other half. The embedded verb (from here on, the 'RC verb') was always a *be* or *have* verb form that was inflected for number agreement; depending on the item, it served either

as an auxiliary or as the main verb of the relative clause. A semantically neutral adverb always appeared immediately after the critical region.

We observed in Section 2.3 that for some of the stimuli used in Experiment 1, the relationship between the high and low NPs was such that an explanation relation remained strongly compatible with reference to the low NP. Because our analysis and the default low-attachment analysis make the same predictions for such cases, and thus these cases cannot be used to discriminate between the analyses, we sought to avoid such stimuli for this experiment. Similarly, we also observed in Section 2.3 that the biases towards ensuing explanations for the non-IC and IC verbs were not always as consistent as we intended, which also affected the extent to which the two analyses could be discriminated. We therefore selected IC verbs from Experiment 1 that had strong biases towards explanation RCs, and supplemented them with a few additional ones taken from Kehler et al.'s (2008) study (which were in turn taken from McKoon et al.'s (1993) study, with some minor substitutions). The 20 non-IC verbs similarly consisted of a mix of verbs from McKoon et al. (1993), Levin (1993), and from Experiment 1. The full set of verbs used can be found in the appendix, as can the complete set of experimental items.

The experiment consisted of 10 practice items, followed by 20 experimental items mixed with 30 fillers, pseudorandomized for each subject. The filler items were similar to the stimuli in that some included proper names and RCs or other subordinate clauses.

Procedure

Items were presented in a moving-window self-paced reading paradigm, using DMDX experiment software (Forster & Forster, 2003). Sentences appeared in white letters on a dark background, left-justified on a 19" CRT screen, and no sentence was longer than one line of text. Sentences initially appeared as a series of dashes ($--\!-\!-$) obscuring the words, and participants pushed a button on a Logitech USB gamepad to reveal each region. The presentation was non-cumulative such that previous regions were replaced with dashes when the next region appeared. The critical region and the spillover regions were revealed one word at a time, but multi-word regions were used elsewhere to present short phrases such as a verb and a preposition (*stared at*, *stood near*) or a determiner and a noun (*the children*). Multi-word regions are indicated in the stimuli set in the appendix. Participants pushed either a YES or NO button on the gamepad to answer a comprehension question after every sentence, and they

received automatic feedback whenever they answered incorrectly. They were instructed to read as quickly and carefully as possible, making sure they understood the complete sentence and slowing down if they answered multiple questions incorrectly. We recorded reading times for each region as well as the participant's response to the comprehension question.

3.1 Results

After excluding three participants whose comprehension-question accuracy was not significantly better than chance, the percentage of correct responses was 93.03% for fillers and 85.07% for experimental items (percentages over subject means), indicating that participants paid attention to the task. Table 1 shows the raw reading times by condition for the critical region and the spillover regions, as well as the mean accuracy on comprehension questions for each condition.

	RC Verb	**Spillover1**	**Spillover2**	**Accuracy**
IC.high	395.70 ±16.83	430.43 ±18.90	442.81 ±18.84	.873 ±.02
IC.low	398.83 ±16.71	474.16 ±23.26	477.19 ±26.34	.780 ±.02
nonIC.high	402.03 ±16.55	501.48 ±24.26	473.59 ±20.22	.862 ±.02
nonIC.low	403.96 ±13.83	462.63 ±20.03	437.50 ±15.91	.887 ±.02

Table 1: Raw RTs and question accuracy (subject means ± standard error)

Figure 3 shows the residual reading times for each of the four conditions starting at the matrix verb. Comprehension-question accuracy and reading times were analyzed with 2×2 ANOVAs, by subject and by item. The results were also analyzed using linear mixed-effects models with random subject-specific and item-specific intercepts. For these models, we report the coefficient estimates and MCMC-derived p-values (Baayen, Davidson, & Bates, *in press*).

Comprehension question accuracy

All question-accuracy ANOVAs were conducted on arcsine transformed proportions of correct answers. There was a marginal main effect of attachment height favoring high-attaching RCs (marginal by subject: $F_1(1,54)=3.889$, $p=0.054$; $F_2(1,19)=2.778$, $p=0.112$). There was also a main effect of verbtype favoring non-IC verbs (significant only by subject: $F_1(1,54)=4.59$, $p<0.05$; $F_2(1,19)=2.206$, $p=0.154$). This main effect

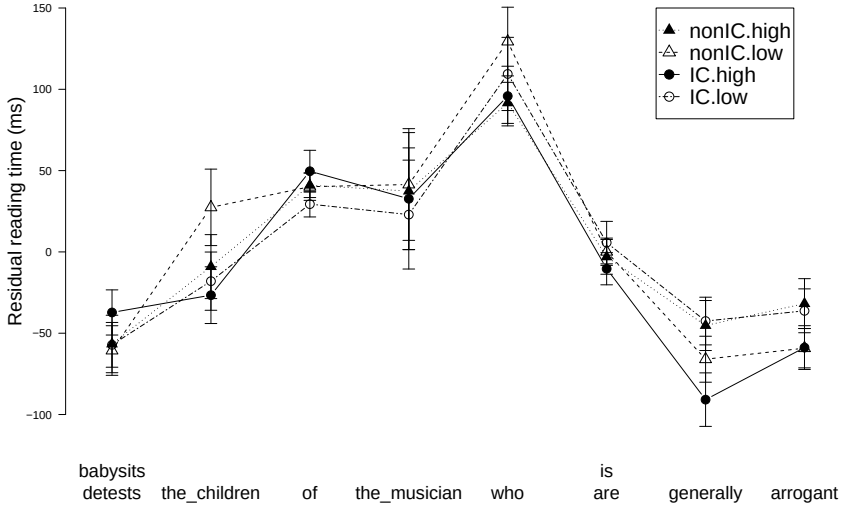

Figure 3: Residual reading times

is driven by an interaction in which low-attaching RCs in the IC condition yielded lower accuracies than any other condition. The crossover interaction that emerges is consistent with the predicted interaction for processing difficulty in cases in which the RC violates the expectations of the preceding context (verbtype×attachment-height interaction significant by subject and marginal by item: $F_1(1,54)=7.346$, $p<0.01$; $F_2(1,19)=3.89$, $p=0.063$). In a mixed-effects logistic regression, attachment height was not a significant factor for modeling question accuracy (attachment: β=-0.273, $p=0.13$), whereas verbtype and the verbtype×attachment interaction were significant (verbtype: $\beta=0.445$ $p<0.05$; verbtype×attachment interaction: $\beta=0.972$, $p<0.01$).

Reading time results

We analyzed residual reading times at the critical region and two spillover regions. Residual reading times adjust for overall differences in participants' reading rates as well as differences in readers' sensitivity to word length. Residual RTs were calculated as the difference between the actual reading time on a word and the reading time predicted by a regression

equation (computed separately for each participant, using all experimental and filler items) relating word length to reading time (F. Ferreira & Clifton, 1986). We removed residual RTs that were more than four standard deviations away from the mean, per region and per condition (0.36% of the data). The following results represent an analysis of all non-outlier items, regardless of comprehension-question accuracy.

At the disambiguating RC verb (*is*/*are*), there were no significant effects for verbtype (Fs<1), attachment height ($F_1(1,54)=1.071$, p=0.31; $F_2(1,19)=1.124$, p=0.30), or the verbtype×attachment interaction (Fs<1). In a mixed-effects linear regression, the factors for verbtype, attachment height, and the verbtype×attachment interaction were not significant factors for modeling residual reading time (verbtype: $\beta=-0.1556$, p=0.99; attachment: $\beta=9.84$, p=0.31; verbtype×attachment interaction: $\beta=-13.258$, p=0.51).

At the first spillover region (*generally*), there were again no main effects of verbtype (Fs<1) or attachment height ($F_1(1,54)=1.295$, p=0.26, $F_2<1$). A significant crossover interaction was observed in the predicted direction: High attachments were read more slowly than low attachments in non-IC conditions but faster in IC conditions ($F_1(1,54)=5.522$, $p<0.05$; $F_2(1,19)=6.167$, $p<0.05$). In a mixed-effects linear regression, the factors for verbtype and attachment were not significant (verbtype: $\beta=8.672$, p=0.47; attachment: $\beta=12.027$, p=0.31), whereas the verbtype×attachment interaction was (verbtype×attachment interaction: $\beta=-63.60$, $p<0.01$).

At the second spillover region (*arrogant*), there were again no main effects (Fs<1), but the same crossover interaction was significant ($F_1(1,54)=6.588$, $p<0.05$; $F_2(1,19)=4.967$, $p<0.05$). In a mixed-effects linear regression, the main factors of verbtype and attachment height were not significant, but the interaction was (verbtype: $\beta=3.106$, p=0.78; attachment: $\beta=-3.279$, p=0.77; verbtype×attachment interaction: $\beta=-47.10$, $p<0.05$).

Analyses of the raw reading times were qualitatively the same, as were

analyses of residual RTs with incorrectly answered items excluded.[8]

Outside the critical region, there was a marginal effect of verbtype at NP1 (*the_children*), with non-IC verbs yielding slower reading times (verbtype: $F_1(1,54)=2.977$, $p=0.09$; $F_2(1,19)=3.945$, $p=0.062$; attachment: $F_1(1,54)=1.616$, $p=0.21$; $F_2(1,19)=1.854$, $p=0.19$; verbtype×attachment: $Fs<1$; regression for NP1 with factors for verbtype: $\beta=33.30$, $p=0.06$; attachment: $\beta=23.00$, $p=0.19$; verbtype×attachment: $\beta=29.613$, $p=0.40$). No other pre-critical regions yielded significant effects.

3.2 Discussion

Experiment 2 was designed to test whether expectations about discourse continuations have an impact on the immediate processing of RCs. Not only was the bias toward low-attaching RCs reduced in the IC condition as compared to the non-IC condition (as predicted), but the bias showed a full reversal, such that high-attaching RCs were actually read more quickly than low-attaching RCs in the IC condition. As predicted, the reverse was true in the non-IC condition. The effect was observed in the region immediately

[8]Considering the raw reading times, there were no effects at the disambiguating verb ($Fs<1$; regression with factors for verbtype: $\beta=4.112$, $p=0.69$; attachment: $\beta=5.577$, $p=0.61$; verbtype×attachment: $\beta=-7.953$, $p=0.72$). At the first spillover region, there was an effect of verbtype by subject and a significant crossover interaction (verbtype: $F_1(1,54)=7.075$, $p<0.05$; $F_2(1,19)=3.548$, $p=0.075$; attachment: $Fs<1$; verbtype×attachment: $F_1(1,54)=6.853$, $p<0.05$; $F_2(1,19)=5.434$, $p<0.05$; regression for spillover1 with factors for verbtype: $\beta=30.466$, $p<0.05$; attachment: $\beta=1.787$, $p=0.88$; verbtype×attachment: $\beta=-81.31$, $p<0.005$). At the second spillover region, there were no main effects ($Fs<1$) and the same interaction was significant ($F_1(1,54)=6.705$, $p<0.05$; $F_2(1,19)=6.078$, $p<0.05$; regression for spillover2 with factors for verbtype: $\beta=-2.968$, $p=0.83$; attachment: $\beta=-2.738$, $p=0.81$; verbtype×attachment: $\beta=-70.41$, $p<0.005$).

Considering the residual reading times with incorrectly answered items excluded, there were no effects at the disambiguating verb (verbtype: $F_1(1,54)=1.442$, $p=0.24$; $F_2(1,19)=2.116$, $p=0.16$; attachment: $Fs<1$; verbtype×attachment: $Fs<1$; regression with factors for verbtype: $\beta=7.338$, $p=0.43$; attachment: $\beta=6.813$, $p=0.51$; verbtype×attachment: $\beta=-12.77$, $p=0.53$). At the first and second spillover regions there were no main effects but significant crossover interactions (Spillover1: verbtype: $F_1(1,54)=2.646$, $p=0.11$; $F_2(1,19)=1.701$, $p=0.21$; attachment: $Fs<1$; verbtype×attachment: $F_1(1,54)=6.117$, $p<0.05$; $F_2(1,19)=5.216$, $p<0.05$; regression for spillover1 with factors for verbtype: $\beta=22.064$, $p=0.073$; attachment: $\beta=7.779$, $p=0.56$; verbtype×attachment: $\beta=-74.05$, $p<0.005$; Spillover2: verbtype: $Fs<1$; attachment: $Fs<1$; verbtype×attachment: $F_1(1,54)=7.598$, $p<0.01$; $F_2(1,19)=5.465$, $p<0.05$; regression for spillover2 with factors for verbtype: $\beta=0.1607$, $p=0.98$; attachment: $\beta=-2.4821$, $p=0.84$; verbtype×attachment: $\beta=-64.89$, $p<0.01$).

following the disambiguation point.

These results show that the coherence biases obtained in the earlier completion study also have an impact on reaction times in a self-paced reading experiment. Crucially, comprehenders appear to use these discourse cues in their incremental processing such that predictions at the matrix verb can impact the attachment of an RC structure in the same sentence. In order for participants' reading times to be affected at the first spillover region of the RC, several key assumptions must all hold: (i) comprehenders' biases regarding upcoming discourse continuations must be conditioned on the IC context, (ii) comprehenders must be implicitly aware that RCs can provide explanations of preceding material, and (iii) comprehenders' biases regarding next mention must be conditioned on the likelihood of an explanation. Despite the fact that several different discourse cues contribute and that their integration requires a fairly complicated chain of reasoning, the timecourse of the effect suggests that participants are invoking these biases mid-sentence and using them to generate expectations about upcoming syntactic structures.

4 General Discussion

The field of psycholinguistics has placed considerable emphasis on RC processing as a testing ground for a variety of phenomena, including ambiguity resolution (Cuetos & Mitchell, 1988), extraction (Kluender, 1992), optionality (V. S. Ferreira & Dell, 2000), expectation-based parsing (Levy, 2008), and thematic fit (McRae et al., 1998), among others. In this paper, we have shown how RC processing can provide answers to different questions: the timecourse over which people construct a discourse context and the ways in which discourse-level information impacts syntactic disambiguation. Of particular interest are the questions of whether or not comprehenders generate expectations about upcoming discourse coherence relations (Kehler, 2002) and whether those expectations influence syntactic processing.

Our sentence-completion study showed first that IC verbs yielded RCs that are more likely to explain the event in the matrix clause than RCs following non-IC verbs, and second, that these explanation-providing RCs are more likely to attach high following IC verbs. These results are in keeping with previous work on intersentential coherence (Kehler et al., 2008) suggesting that the presence of an IC verb raises the expectation for an upcoming explanation and creates a concomitant bias for a subsequent men-

tion of the causally-implicated referent when the explanation is reached. Our self-paced reading study yielded a reversal of the RC low-attachment preference in cases in which the causally-implicated referent was located at the high attachment site. These on-line results suggest that comprehenders track such expectations mid-sentence, and that these expectations influence syntactic processing. In the remainder of this section, we situate our results relative to previous work on RC attachment ambiguity and discourse-sensitive models of sentence processing. We also briefly discuss how an expectation-based processing model that incorporates information about discourse continuations can open up new avenues for sentence processing research.

4.1 The Role of Discourse in RC Processing

Our experiments complement several previous studies showing that properties of the discourse context can influence RC processing (van Berkum et al., 1999; Zagar et al., 1997; Desmet et al., 2002; Papadopoulou & Clahsen, 2006). These previous studies manipulate referential context (uniqueness of the NPs) in order to influence RC interpretation: Comprehenders are biased toward attachments that prevent failure of unambiguous definite reference. Our study goes beyond this previous work in two key respects. First, it is the first study to look at the effect of discourse processing beyond referential ambiguity on syntactic comprehension. Second, the modulation of attachment preferences observed in our study cannot be reduced to a simple bias against infelicity in definite descriptions (or in any other aspect of the sentence). For example, (18a), repeated below as (25), is in no way infelicitous despite the fact that the RC does not provide an explanation of the detesting; the need for an explanation can easily be satisfied by a subsequent sentence (26).

(25) John detests the children of the musician who lives in La Jolla.

(26) The children are arrogant and rude.

Taken together, this work suggests that discourse-level factors are as important as, and fully integrated with, lexical and morphosyntactic cues in resolving ambiguity.

We have taken advantage of the fact that an RC can be used to express an explanation for a matrix-clause event in modulating attachment preferences. Notably, as described in Section 1.4, such explanations are

only pragmatically *implicated* when provided by an RC – and as such, our study demonstrates that a comprehender's expectations about ensuing conversationally-implicated information can indeed influence the initial steps of parsing (cf. Clifton & Ferreira, 1989, as discussed in Section 1.1). The experiments also establish a potentially intriguing connection between the next mention biases in Explanation relations that Kehler et al. (2008) found for pronoun interpretation and those illustrated here for RC attachment, which itself can be seen as involving the anaphoric binding of a *wh*-pronoun (see Hemforth et al., 2000 for discussion of the relationship between pronoun interpretation and *wh*-pronoun binding). We now turn to the question of what types of on-line processing models can account for the results reported in this paper as well as previous studies on the effect of discourse context on syntactic comprehension.

4.2 Expectation-Based Models for Discourse-Sensitive Syntactic Comprehension

There are several aspects of our experimental results that any processing model needs to account for. First, there was a close match between off-line attachment preferences in sentence completion in Experiment 1 and on-line reading preferences in Experiment 2: In both cases, the presence of an object-biased IC verb eliminated low-attachment bias. Experiment 1 makes it clear that this shift in attachment preference is driven by a much higher proportion of RC completions providing an explanation for the matrix-clause event in the IC condition than in the non-IC condition. Second, the on-line preferences in Experiment 2 emerged at the first spillover region after the disambiguating RC verb (i.e., *is/are* in 23–24), well before participants encountered the lexical material necessary to establish that the RC provides an explanation for the matrix-clause event. These results are problematic for any theory in which the establishment of interclausal discourse coherence relations (including those between matrix and relative clauses) is not fully incremental.

The question remains of what theories *can* naturally capture these results, and how. We believe that these results can be handled by a range of evidential models of on-line comprehension, so long as (a) syntactic attachment preferences are made fully incrementally (c.f. the delay model of Rayner & Frazier, 1987), and (b) discourse-based preferences are taken into account and can interact fully with any other biases that may be active. These models include at least the competition-integration model

(MacDonald, 1994; Spivey & Tanenhaus, 1998; McRae et al., 1998), probabilistic disambiguation/pruning and attention-shift models (Jurafsky, 1996; Narayanan & Jurafsky, 1998, 2002; Crocker & Brants, 2000), and surprisal (Hale, 2001; Levy, 2008). Although none of these models as presented to date has explicitly included discourse constraints, their probabilistic architectures allow for the incorporation of potentially arbitrary information sources, and there is no reason why discourse factors could not be smoothly incorporated into any of these models. As an illustration, we now briefly describe how the results in this paper can be accounted for by surprisal.

4.2.1 Surprisal-Based Analysis of Experiment 2

Assuming that the difficulty at the first spillover region in Experiment 2 is indeed spillover generated by the disambiguation at the preceding RC verb, we focus on the conditional probability of the RC verb. For ease of illustration, we will cast our formulation specifically in terms of examples (23) and (24), repeated below as (27) and (28):

(27) NON-IC MATRIX: John babysits the children of the musician ...
  a. [LOW ATTACHMENT] ... who *is* generally arrogant and rude.
  b. [HIGH ATTACHMENT] ... who *are* generally arrogant and rude.

(28) IC MATRIX: John detests the children of the musician ...
  a. [LOW ATTACHMENT] ... who *is* generally arrogant and rude.
  b. [HIGH ATTACHMENT] ... who *are* generally arrogant and rude.

All simplifications that result from focusing on (27) and (28) are made without loss of generality.

First we introduce some notation for several events:

- RC {high, low}: the presence of a subject-extracted relative clause (SRC) that attaches high or low into a two-level NP
- M: the context preceding the word *who*, i.e., the matrix clause
- C: the context preceding the RC verb, i.e., the word *who* and M
- NP_i: the event that the discourse referent mentioned as the subject of the clause following M (whether as an SRC or new sentence) will be the i-th NP of the NP complex. (In the case that an SRC follows, NP_i is taken to be the referent of *who*.)

- *Reln*: the coherence relation holding between M and the next clause

Under surprisal, the difficulty of a word w in its context is determined by the log of its inverse conditional probability, $\log \frac{1}{P(w|C)}$. The difference between the difficulties of *is* and *are* in the same context is thus determined by the ratio of their conditional probabilities: $\frac{P(are|C)}{P(is|C)}$. Since *is* is only possible with a low-attaching RC in (27) and (28), and *are* with a high-attaching RC, we can rewrite this ratio as:

$$\frac{P(are|C)}{P(is|C)} = \frac{P(are|\text{RC high}, C)P(\text{RC high}|C)}{P(is|\text{RC low}, C)P(\text{RC low}|C)} \tag{I}$$

Our first assumption is that the probability of using a *be* verb of the appropriate form is approximately the same regardless of whether the RC is attached low or high—that is, $P(are|\text{RC high}, C) \approx P(is|\text{RC low}, C)$. We can then rewrite Equation (I) as

$$\frac{P(are|C)}{P(is|C)} \approx \frac{P(\text{RC high}|C)}{P(\text{RC low}|C)} \tag{II}$$

A high-attaching SRC implies the event NP_1—i.e. that the next-mentioned referent is the high NP—and likewise a low-attaching SRC implies NP_2. We also take advantange of the fact that C is composed of the events M and *who*, giving us

$$\frac{P(are|C)}{P(is|C)} \approx \frac{P(\text{RC high}, NP_1|who, M)}{P(\text{RC low}, NP_2|who, M)} \tag{III}$$

By the definition of conditional probability we can rewrite the numerator $P(\text{RC high}, NP_1|who, M)$ as $P(\text{RC high}, NP_1, who|M)/P(who|M)$, and similarly for the denominator, giving us

$$\frac{P(are|C)}{P(is|C)} \approx \frac{P(\text{RC high}, NP_1, who|M)/P(who|M)}{P(\text{RC low}, NP_2, who|M)/P(who|M)} \quad \text{(IV)}$$

$$= \frac{P(\text{RC high}, NP_1, who|M)}{P(\text{RC low}, NP_2, who|M)} \quad \text{(V)}$$

$$= \frac{P(who|\text{RC high}, NP_1, M)P(\text{RC high}, NP_1|M)}{P(who|\text{RC low}, NP_2, M)P(\text{RC low}, NP_2|M)} \quad \text{(VI)}$$

We introduce our second assumption here: That the probability of using *who* to begin an SRC is approximately independent of the attachment level. This allows us to cancel the first terms of the numerator and denominator, giving us

$$\frac{P(are|C)}{P(is|C)} \approx \frac{P(\text{RC high}, NP_1|M)}{P(\text{RC low}, NP_2|M)} \quad \text{(VII)}$$

$$= \frac{P(\text{RC high}|NP_1, M)P(NP_1|M)}{P(\text{RC low}|NP_2, M)P(NP_2|M)} \quad \text{(VIII)}$$

Our third assumption is that the probability of expressing the next clause as an SRC is also approximately independent of whether the high or low NP is the subject of the next clause. This allows us to cancel the first terms once more, and making the marginalization over discourse coherence relations explicit gives us:

$$\frac{P(are|C)}{P(is|C)} \approx \frac{\sum_{Reln} P(NP_1|Reln, M)P(Reln|M)}{\sum_{Reln} P(NP_2|Reln, M)P(Reln|M)} \quad \text{(IX)}$$

The role of discourse coherence relations now becomes clear. In Experiments 1 and 2, we contrasted conditions with and without object-biased IC verbs in M. It has been independently verified that the presence of such an IC verb causes two things to happen: First, it makes the coherence relation *Reln* with the next clause more likely to be an Explanation (Kehler et al., 2008), and second, it makes the next-mentioned NP more likely to be the direct object of that verb *when* Reln *is an Explanation* (Garvey & Caramazza, 1974; Brown & Fish, 1983; Au, 1986; McKoon et al., 1993; Kehler et al., 2008, a.o.). The

first effect corresponds to a rise in $P(\textit{Reln=Explanation}|M)$, and the second effect to a rise in $P(NP_1|\textit{Reln=Explanation}, M)$ and a corresponding drop in $P(NP_2|\textit{Reln=Explanation}, M)$. Thus, the most important term in the sums of Equation (IX) will be that for which *Reln=Explanation*, and this term will be large in the numerator (because $P(NP_1|\textit{Reln=Explanation}, M)$ is large) and small in the denominator (because $P(NP_2|\textit{Reln=Explanation}, M)$ is small). As a result, we can expect the probability ratio $\frac{P(are|C)}{P(is|C)}$ to favor *are* more strongly in the IC condition than in the non-IC condition. This matches the empirical results of Experiment 2, where we found an interaction between RC attachment and verbtype immediately after the RC verb.

Modeling the surprisal at the RC verb in this way draws a direct connection between RC processing and earlier literature on the effects of IC verbs. By making three assumptions of approximate probabilistic independence to draw this connection – involving the probability of using an RC verb at all, using the word *who* to mark the RC onset, and expressing the next clause as an RC – our formulation makes totally explicit all the conditions that are required to predict the surprisal-based processing differences that occurred in our experiment. Importantly, however, we do not claim that these independence assumptions are imposed absolutely by comprehenders, but rather that they are approximately correct for the materials we used. Furthermore, if we constructed materials that broke these independence assumptions, we would expect different patterns of processing difficulty. As a clear example, if one of the NPs were inanimate, then the second assumption (independence of *who*-marking) would be false, and would lead to the prediction that attachment preferences should consistently favor the animate NP. Likewise, if we were able to construct matrix clauses M such that explanations involving one NP were more likely to be expressed as a new sentence than explanations involving the other NP, we would again see different predictions.

4.2.2 Discourse Continuations as a Unit of Prediction

Although these results can be incorporated into incremental models of probabilistic syntactic comprehension, they further constrain these models in terms of the information sources that they must include. That is, models of sentence processing can no longer be built separately from models of discourse processing. If the goal is to capture the behavior of the human sentence processor as it encounters each new word (within a sentence

and at sentence boundaries), we must reconsider the units over which such predictions are made, and the range of features over which those predictions are conditioned. In the analysis of the last section, for example, intraclausal discourse coherence relations were both a unit of prediction—$P(Reln|M)$—and a feature implicated indirectly in a syntactic prediction—$P(NP_1|Reln, M)$ and $P(\text{RC high}|NP_1, Reln, M)$.

By considering discourse continuations as a unit over which comprehenders may make predictions, we are beginning to address the long-standing issue of incorporating discourse information (the nebulous pragmatic wastebasket) into formal models of sentence processing. Most researchers would not deny that complex inferencing affects language usage and comprehension, and there is a rich history of its use in the realms of formal semantics and artificial intelligence (Hobbs, 1979), but finding a quantifiable unit over which to estimate predictions for processing effects has been hard to do. There is consensus about the importance of appealing to discourse factors, especially for our current RC models — both in the psycholinguistics community (Gilboy et al., 1995; Frazier & Clifton, 1996; Traxler et al., 1998; Hemforth et al., 2000; Desmet et al., 2002) and in the natural language processing literature (Siddharthan, 2002). In both areas, a host of morphosyntactic biases and heuristics have already been identified but the problem remains unsolved. New insights into language processing will require models that appeal to novel discourse-level cues.

Discourse coherence provides one way of structuring elements of a discourse into a unified whole and allowing expectations to be calculated over concrete, quantifiable features of the local or global discourse context. The advantage of the approach taken in this paper is that a concrete factor (verbtype) was manipulated in order to generate repercussions at the level of coherence. As such, 'discourse-level factors' need not be relegated to the status of haphazard or fuzzy cues (see Kadmon, 2001 for a discussion of what constitutes a pragmatic explanation) nor do we need to restrict our analyses to 'neutral' discourse contexts and make claims concerning processing biases that are active *all else being equal*. In fact, all else is never equal, and our hope is that acknowledging this lack of neutrality in the discourse context will lead to more research quantifying the properties and structure of the surrounding discourse.

5 Conclusions

The studies reported here have established that certain contexts increase the likelihood of certain discourse coherence relations and that these biases have an impact on RC attachment ambiguity. Expectations regarding upcoming coherence relations were shown to arise from properties of the matrix clause, specifically the matrix verb. By varying the type of verb in the matrix clause, we found that a subsequent RC completion was more likely to provide an explanation of the event in the matrix clause if the verb belonged to the class of IC verbs (Experiment 1). The RC was also more likely to attach to the direct object of an object-biased IC verb than to the object of a non-IC verb in cases in which the direct object occupied the higher position of a complex NP. We also found that these off-line preferences were mirrored in on-line comprehension: In a self-paced reading study (Experiment 2), high-attaching RCs following IC verbs were read more quickly than low-attaching RCs in non-IC contexts. Crucially, these effects occur before comprehenders have been exposed to complete clauses; that is, *expectations about interclausal discourse coherence relations are updated fully incrementally, and have moment-by-moment influence on syntactic disambiguation.* Existing models that have been proposed to account for the widely reported dispreference for high attachments in English fail to account for these results since, in our experiments, the structure of the complex NP, the lexical properties of the individual nouns, and the referential context were the same across conditions. What changed between conditions was the expectation triggered by the matrix clause verb regarding the likelihood of an upcoming Explanation relation. The fact that comprehenders appear to be sensitive to coherence-level biases mid-sentence attests to the importance of constructing models of sentence processing that incorporate information about discourse coherence relations.

Acknowledgments

This research was supported by NIH grant T32-DC000041 to the Center for Research in Language at the University of California, San Diego and NSF IGERT grant DGE-0333451. The results from this paper have been presented at the 21st CUNY Conference on Human Sentence Processing, Chapel Hill, NC, March 2008. We also thank Klinton Bicknell, Laura Kertz, and research assistant Erica Gold.

References

Agresti, A. (2002). *Categorical data analysis* (second ed.). Wiley.

Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, *30*, 191–238.

Au, T. K. (1986). A verb is worth a thousand words: The causes and consequences of interpersonal events implicit in language. *Journal of Memory and Language*, *25*, 104-122.

Baayen, R. H., Davidson, D. J., & Bates, D. M. (*in press*). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* (Special issue on Emerging Data Analysis).

Brown, R., & Fish, D. (1983). The psychological causality implicit in language. *Cognition*, *14*, 237-273.

Brysbaert, M., & D.C., M. (1996). Modifier attachment in sentence parsing: Evidence from Dutch. *Quarterly Journal of Experimental Psychology*, *49A*, 664-695.

Caramazza, A., Grober, E., Garvey, C., & Yates, J. (1977). Comprehension of anaphoric pronouns. *Journal of Verbal Learning and Verbal Behaviour*, *16*, 601-609.

Carreiras, M., & Clifton, C. (1999). Another word on parsing relative clauses: Eyetracking evidence from Spanish and English. *Memory and Cognition*, *27*, 826-833.

Clifton, C., Jr, & Ferreira, F. (1989). Ambiguity in context. *Language and Cognitive Processes*, *4*(3), 77–103.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological syntax processor. In D. Dowty, L. Karttunen, & A. Zwicky (Eds.), (pp. 443–467). Cambridge.

Crocker, M., & Brants, T. (2000). Wide-coverage probabilistic sentence processing. *Journal of Psycholinguistic Research*, *29*(6), 647–669.

Cuetos, F., & Mitchell, D. (1988). Cross-linguistic differences in parsing: Restrictions on the use of the late closure strategy in Spanish. *Cognition*, *30*, 73-105.

Cuetos, F., Mitchell, D., & Corley, M. (1996). Parsing in different languages. In M. Carreiras, J. García-Albea, & N. Sebastián-Gallés (Eds.), *Language processing in Spanish* (p. 145-187). Mahwah, New Jersey: Lawrence Erlbaum.

Desmet, T., Baecke, C. de, Drieghe, D., Brysbaert, M., & Vonk, W. (2006). Relative clause attachment in Dutch: On-line comprehension corresponds to corpus frequencies when lexical variables are taken into account. *Language and Cognitive Processes*, *21*, 453-485.

Desmet, T., Brysbaert, M., & De Baecke, C. (2002). The correspondence between sentence production and corpus frequencies in modifier attachment. *The Quarterly Journal of Experimental Psychology*, *55*, 879-896.

Desmet, T., de Baecke, C., & Brysbaert, M. (2002). The influence of referential discourse context on modifier attachment in Dutch. *Memory and Cognition*, *30*, 150-157.

Desmet, T., & Gibson, E. (2003). Disambiguation preferences and corpus frequencies in noun phrase conjunction. *Journal of Memory and Language*, *49*, 353–374.

Fernandez, E. (2003). *Bilingual sentence processing: Relative clause attachment in bilinguals and monolinguals*. Amsterdam: John Benjamins.

Ferreira, F., & Clifton, J. (1986). The independence of syntactic processing. *Journal of Memory and Language*, *25*, 348-368.

Ferreira, V. S., & Dell, G. S. (2000). The effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive Psychology*, *40*, 296-340.

Forster, K. I., & Forster, J. C. (2003). DMDX: A Windows display program with millisecond accuracy. *Behavior Research Methods, Instruments and Computers*, *35*, 116124.

Frazier, L. (1978). *On comprehending sentences: Syntactic parsing strategies*. Unpublished doctoral dissertation, University of Connecticut.

Frazier, L., & Clifton, C. (1996). *Construal*. Cambridge, Mass: MIT Press.

Garvey, C., & Caramazza, A. (1974). Implicit causality in verbs. *Linguistic Inquiry*, *5*, 459-464.

Gilboy, E., Sopena, J., Clifton, C., & Frazier, L. (1995). Argument structure and association preferences in Spanish and English complex NPs. *Cognition*, *54*, 131-167.

Grice, H. P. (1975). Logic and conversation. In P. Cole & G. Morgan (Eds.), *Syntax and semantics III* (pp. 41–58). London: Academic Press.

Hale, J. (2001). A probabilistic Earley parser as a psycholinguistic model. In *Proceedings of the Second Meeting of the North American Chapter of the Asssociation for Computational Linguistics.* Pittsburgh, PA, June 2-7, 2001.

Hemforth, B., Konieczny, L., & Scheepers, C. (2000). Syntactic attachment and anaphora resolution: The two sides of relative clause attachment. In M. Crocker, M. Pickering, & C. Clifton (Eds.), *Architectures and mechanisms for language processing* (pp. 259–281). Cambridge: Cambridge University Press.

Hobbs, J. R. (1979). Coherence and coreference. *Cognitive Science*, *3*, 67-90.

Jaeger, T. F. (*in press*). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language* (Special issue on Emerging Data Analysis).

Jurafsky, D. (1996). A probabilistic model of lexical and syntactic access and disambiguation. *Cognitive Science*, *20*(2), 137–194.

Kadmon, N. (2001). *Formal pragmatics: Semantics, pragmatics, presupposition, and focus*. Malden, Mass.: Blackwell.

Kehler, A. (2002). *Coherence, reference, and the theory of grammar*. Stanford, CA: CSLI Publications.

Kehler, A., Kertz, L., Rohde, H., & Elman, J. (2008). Coherence and coreference revisited. *Journal of Semantics*, *25*, 1-44.

Kluender, R. (1992). Deriving island constraints from principles of predication. In H. Goodluck & M. Rochemont (Eds.), *Island constraints: Theory, acquisition and processing* (p. 223-258). Dordrecht: Kluwer.

Levin, B. (1993). *English verb classes and alternations: A preliminary investigation*. Chicago: University of Chicago Press.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, *106*, 1126–1177.

MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes*, *9*(2), 157–201.

McKoon, G., Greene, S., & Ratcliff, R. (1993). Discourse models, pronoun resolution, and the implicit causality of verbs. *Journal of Experimental Psychology*, *19*, 1040-1052.

McRae, K., Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*, *38*(3), 283–312.

Mitchell, D. (1984). An evaluation of subject-paced reading tasks and other methods for investigating immediate processes in reading. In D. Kieras & M. Just (Eds.), *New methods in reading comprehension research* (p. 6989). Hillsdale, NJ: Earlbaum.

Mitchell, D., & Brysbaert, M. (1998). Challenges to recent theories of crosslinguistic variation in parsing: Evidence from Dutch. *Syntax and Semantics*, *31*, 313-335.

Narayanan, S., & Jurafsky, D. (1998). Bayesian models of human sentence processing. In *Proceedings of the Twelfth Annual Meeting of the Cognitive Science Society.*

Narayanan, S., & Jurafsky, D. (2002). A Bayesian model predicts human parse preference and reading time in sentence processing. In *Advances*

in neural information processing systems (Vol. 14, pp. 59–65).

Ni, W., Crain, S., & Shankweiler, D. (1996). Sidestepping garden paths: Assessing the contributions of syntax, semantics and plausibility in resolving ambiguities. *Language and Cognitive Processes*, *11*(3), 283–334.

Papadopoulou, D., & Clahsen, H. (2006). Ambiguity resolution in sentence processing: The role of lexical and contextual information. *Journal of Linguistics*, *42*, 109-138.

Rayner, K., & Frazier, L. (1987). Selection mechanisms in reading lexically ambiguous words. *Journal of experimental psychology. Learning, memory, and cognition*, *15*(5), 779–790.

Rohde, H., Kehler, A., & Elman, J. (2006). Event structure and discourse coherence biases in pronoun interpretation. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society.* Vancouver, July 26-29, 2006.

Rohde, H., Kehler, A., & Elman, J. (2007). Pronoun interpretation as a side effect of discourse coherence. In *Proceedings of the 29th Annual Conference of the Cognitive Science Society.* Nashville, TN, August 1-4, 2007.

Sedivy, J. C. (2002). Invoking discourse-based contrast sets and resolving syntactic ambiguities. *Journal of Memory and Language*, *46*, 341–370.

Siddharthan, A. (2002). Resolving relative clause attachment ambiguities using machine learning techniques and WordNet hierarchies. In *5th National Colloquium for Computational Linguistics in the UK (CLUK 2002)* (p. 45-49). University of Leeds, January 8-9, 2002.

Simner, J., & Pickering, M. J. (2005). Planning causes and consequences in discourse. *Journal of Memory and Language*, *52*, 226-239.

Sokal, R. R., & Rohlf, F. J. (1995). *Biometry.* W.H. Freeman.

Spivey, M. J., & Tanenhaus, M. K. (1998). Syntactic ambiguity resolution in discourse: Modeling the effects of referential content and lexical frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(6), 1521–1543.

Traxler, M., Pickering, M., & Clifton, C. (1998). Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of Memory and Language*, *39*, 558-592.

van Berkum, J., Brown, C. M., & Hagoort, P. (1999). Early referential context effects in sentence processing: Evidence from event-related brain potentials. *Journal of Memory and Language*, *41*, 147-182.

Zagar, D., Pynte, J., & Rativeau, S. (1997). Evidence for early-closure attachment on first-pass reading times in French. *The Quarterly Journal*

of Experimental Psychology, *50*, 421-438.

Appendix

Story completion stimuli (Experiment 1)

The stimuli were all of the form *Name - IC/non-IC verb - complex NP - who.*

1. Carl admires/works with the agent of the rockstars who...
2. Greg adores/smiles at the secretaries of the lawyer who...
3. Jared blamed/noticed the friends of the athlete who...
4. Frank complimented/met the guests of the bride who...
5. Bill congratulated/visited the teacher of the second-graders who...
6. Candice criticized/talked to the leader of the activists who...
7. Beth despises/babysits the children of the jazz musician who...
8. Casey detests/looks like the father of the students who...
9. Melissa dislikes/watches the little girls of the neighbor who...
10. Sandra insulted/chatted with the gardeners of the millionaire who...
11. Ryan likes/resembles the captain of the old sailors who...
12. Joel pities/hires the bodyguards of the celebrity who...
13. Ken praised/videotaped the assistants of the CEO who...
14. Alan punished/saw the accountant of the businessmen who...
15. Tina resents/knows the doctors of the supermodel who...
16. Luis recognized/scolded the landlady of the actors who...
17. Craig rewarded/inspected the servants of the dictator who...
18. Scott ridiculed/counted the fans of the singer who...
19. George thanked/interviewed the representative of the employees who...
20. Alice values/lives next to the surgeon of the soldiers who...
21. Paul worships/listens to the coach of the cheerleaders who...

Sample story completions (taken from participants' completions from Experiment 1)

1. Non-IC verb: Beth babysits the children of the jazz musician who lives in La Jolla. [mod - low]
2. Non-IC verb: Frank met the guests of the bride who were her friends from high school. [mod - high]
3. Non-IC verb: Melissa watches the little girls of the neighbor who works evening shifts. [exp - low]
4. Non-IC verb: Craig inspected the servants of the dictator who were suspected of stealing. [exp - high]

5. IC verb: Melissa dislikes the little girls of the neighbor who lives on her right. [mod - low]
6. IC verb: Frank complimented the guests of the bride who were sitting in the front row. [mod - high]
7. IC verb: Bill congratulated the teacher of the second-graders who had all learned their times tables. [exp - low]
8. IC verb: Alan punished the accountant of the businessmen who was notorious for IRS fraud. [exp - high]

Verb biases observed in Experiment 1

Verbs differed in the proportion of explanation RCs produced.

Verb	Class	% Expl	Verb	Class	% Expl
chat-with	Non-IC	0%	hire	Non-IC	34.8%
count	Non-IC	0%	adore	IC	36.0%
interview	Non-IC	0%	insult	IC	36.4%
know	Non-IC	0%	watch	Non-IC	46.2%
live-next-to	Non-IC	0%	compliment	IC	50.0%
look-like	Non-IC	0%	praise	IC	50.0%
meet	Non-IC	0%	admire	IC	52.6%
recognize	Non-IC	0%	reward	IC	54.5%
resemble	Non-IC	0%	scold	IC	60.0%
see	Non-IC	0%	videotape	Non-IC	61.1%
talk-to	Non-IC	0%	blame	IC	64.3%
work-with	Non-IC	0%	criticize	IC	66.7%
visit	Non-IC	5.3%	dislike	IC	76.2%
babysit	Non-IC	8.0%	worship	IC	80.0%
smile-at	Non-IC	8.7%	congratulate	IC	81.0%
inspect	Non-IC	13.0%	despise	IC	82.6%
notice	Non-IC	14.3%	pity	IC	82.6%
listen-to	Non-IC	18.2%	resent	IC	84.0%
value	IC	22.7%	punish	IC	85.7%
like	IC	26.7%	ridicule	IC	91.3%
thank	IC	33.3%	detest	IC	95.7%

Reading time stimuli (Experiment 2)

The stimuli were all of the form *Name - IC/non-IC verb - complex NP - who - singular/plural RC verb - adverb - continuation.* The forward slash ('/') separates alternatives that differed between conditions (IC/non-IC; singular/plural). An underscore ('_') connects words that were revealed together in one region. Comprehension questions are listed in brackets.

1. Anna scolded/studied_with the_chef of the_aristocrats who was/were routinely letting food go to_waste. [Did food go to waste?]

2. John stared_at/lived_next_to the_teacher of the_second_graders who was/were definitely smartest in the school. [Was the teacher/the second graders smart?]

3. Jenny assisted/joked_with the_maid of the_executives who was/were regularly late to work. [Were the executives/was the maid late to work?]

4. Nick trusted/stood_near the_captain of the_sailors who has/have consistently weathered big storms. [Did the captain have Nick's confidence? Was Nick near the captain?]

5. Angela corrected/gossiped_with the_secretary of the_lawyers who has/have occasionally made small mistakes. [Have there been occasional errors?]

6. Bob comforted/greeted the_leader of the_activists who was/were deeply disappointed by the_court's decision. [Was Bob disappointed with the court's decision?]

7. Laura envies/knows the_manager of the_cashiers who has/have supposedly received a_huge raise. [Did the manager/cashiers get a huge raise?]

8. Zack valued/recognized the_daughter of the_shopkeepers who was/were usually willing to_spot him a_few_dollars. [Did Zack lend money to the daughter?]

9. Sarah fears/jogs_with the_uncle of the_toddlers who is/are often heard yelling and screaming. [Are toddlers known for being well behaved?]

10. Adam noticed/resembled the_representative of the_employees who was/were always wearing safety goggles. [Were the employees / Was the representative wearing safety goggles?]

11. Tina praised/met the_gardeners of the_millionaire who has/have recently installed a_solar powered sprinkler. [Has the millionaire / Have the gardeners put in a new sprinkler system?]

12. Justin hates/carpools_with the_cousins of the_accountant who is/are forever telling the_same tasteless jokes. [Is the accountant / Are the cousins likeable?]

13. Emily blamed/waited_with the_nieces of the_florist who has/have repeatedly ruined expensive orchids. [Did some flowers get damaged?]

14. Joe helped/ran_into the_brothers of the_athlete who is/are perpetually failing math class. [Are the brothers / Is the athlete failing math?]

15. Jessica reproached/worked_with the_doctors of the_supermodel who was/were adamantly in_favor of plastic surgery. [Did the supermodel/doctors advocate plastic surgery?]

16. Brian pacified/visited the_associates of the_businessman who was/were nearly bankrupted by the_new tax_policy. [Did the new tax policy benefit businesses?]

17. Melissa detests/babysits the_children of the_musician who is/are generally arrogant and rude. [Does Melissa get frustrated with the children? / Could Melissa be a teenager?]

18. Frank thanked/talked_to the_servants of the_dictator who has/have lately been helping the_poor. [Does Frank admire altruism? / Did Frank talk to the dictator's staff?]

19. Tracy congratulated/chatted_with the_bodyguards of the_celebrity who was/were constantly fighting off the_paparazzi. [Does the paparazzi ignore celebrities?]

20. Kevin mocked/counted the_fans of the_singer who was/were continually stagediving and getting hurt. [Were the fans diving off the stage? Is the singer someone who dives off the stage?]