

Another Context Effect in Sentence Processing: Implications for the Principle of Referential Support

Michael J. Spivey-Knowlton

Department of Psychology
University of Rochester
Rochester, NY 14627
spivey@psych.rochester.edu

Abstract

A major goal of psycholinguistics is to determine what sources of information are used immediately in language comprehension, and what sources come into play at later stages. Prepositional phrase attachment ambiguities were used in a self-paced reading task to compare contexts that contained one or two possible referents for the verb phrase (VP) in the target sentence. With one set of sentences, a VP-attachment preference was observed in the 2-VP-referent context, but not in the 1-VP-referent context. With another set of sentences, no effect of context was observed. This result falls outside of the scope of the principle of referential support (Altmann & Steedman, 1988) as currently formulated. It suggests that a similar but more broadly-based theory is required.

Introduction

The strong claim for modularity in language processing (Fodor, 1983) has inspired a wealth of research occasionally supporting the existence of an informationally encapsulated syntactic processor (e.g. Ferreira & Clifton, 1986; Frazier, Clifton & Randall, 1983; Rayner, Carlson & Frazier, 1983) and occasionally questioning it (e.g. McDonald, 1992; Spivey-Knowlton, Trueswell & Tanenhaus, 1992; Taraban & McClelland, 1988; Trueswell & Tanenhaus, 1991, this volume; Trueswell, Tanenhaus & Garnsey, 1989). See Altmann (1989) for a review of some of the work on this issue.

The vast majority of this research examines subjects' reading times¹ in sentences that contain temporary syntactic ambiguities. If a subject makes an incorrect syntactic commitment at the ambiguity,

¹ Reading times are measured either by a self-paced reading task in which the subject presses a button to present each successive word or phrase in the sentence, or by monitoring eye-movements and fixation durations while the subject reads text on a computer screen.

reading time will be slow when she encounters a region that resolves that ambiguity.²

A frequently studied type of syntactic ambiguity is the prepositional phrase (PP) attachment ambiguity. Consider the sentence, "Johnny attacked the cat with the rubber mouse." The PP "with the rubber mouse" can be syntactically attached either to the verb phrase (VP) "attacked" or to the noun phrase (NP) "the cat". The former attachment corresponds to an interpretation in which the attacking was done with an instrument called a rubber mouse. The latter attachment would mean that the cat was somehow distinguishable by its association with a rubber mouse -- perhaps it was playing with it.

Some studies have shown that, upon encountering the ambiguously-attaching preposition, readers tend to prefer the VP-attachment (Altmann, 1986; Frazier, 1978; Rayner, et al., 1983). That is, when the sentence is more plausibly an NP-attachment (i.e., "Johnny attacked the cat with the short hair."), readers find themselves "garden-pathed" when they reach the disambiguating region ("short hair").

The predominance of the VP-attachment bias has been interpreted as evidence for the minimal attachment principle (Frazier, 1978, 1987). This explanation rests on the fact that the constituent structure of the NP-attached sentence contains more nodes and more depth of branching than the constituent structure of the VP-attached version. It is claimed that, regardless of context, the sentence processor will initially commit to the structurally less complex attachment.

In response to this, Altmann & Steedman (1988) proposed what is now generally known as the Referential Theory of syntactic disambiguation. Their principle of referential support (derived from Crain & Steedman's (1985) principle of referential success) is based on the assumptions that a definite

² This is called a "garden-path". It is assumed that the reader must then reinterpret the syntactic and thematic relations of the sentence, which takes additional processing time.

NP (e.g., "the cat") presupposes the existence of a unique referent in the discourse model, and when a definite NP is encountered, the reader attempts to link it to the appropriate entity in her discourse model. When the syntactically ambiguous preposition "with" is encountered, both attachments are considered in parallel and their presuppositions weighed. The VP-attachment, or simple NP analysis, "attacked the cat with (instrument)" maintains the previous presupposition of a unique referent. However, the NP-attachment, or complex NP analysis, "attacked the cat with (attribute)" presupposes more than one cat in the discourse model, to one of which a distinguishing reference is being made. Thus when the context contains zero cats³ or one cat, the VP-attachment upholds the correct presupposition. On the other hand, when the context contains two or more cats, only the NP-attachment upholds the correct presupposition. Hence, the reader's attachment preference is modulated by referential pragmatics, not by syntactic complexity.

To experimentally test this theory, Altmann & Steedman (1988) manipulated discourse contexts so that they had either one or two referents for the definite NP preceding the ambiguous PP ("the cat" in "Johnny attacked the cat with..."). For example, one version of a context would contain two possible referents for the definite NP: a cat that has short hair and a cat that has long hair, while the other version would contain one such referent: a cat that has short hair and a dog that has long hair. It was demonstrated that, while a 1-NP-Referent context did not change the VP-attachment bias in the target sentence, a 2-NP-Referent context produced a clear preference for NP-attachment in the target sentence (Altmann & Steedman, 1988; Spivey-Knowlton, 1991).

However, Altmann & Steedman's referential explanation of their findings, though more explicit than broader accounts, may be too narrowly focused. A less restrictive, albeit more vague, explanation of this result is simply that they have set up a minimal pair of entities in context between which only an NP-attachment can discriminate. Their contexts systematically begin by introducing a character carrying an instrument or tool ("A burglar broke into a bank carrying some dynamite."), his intention is stated ("He planned to blow open a safe."), and then the one or two NP referents are introduced ("He saw a safe which had a new lock and a strongbox/safe which had an old lock."). Aside from setting up two referents, such a context may also set up a kind of conceptual uncertainty in which the reader anticipates

³ In a context that contains no cats at all [or in the absence of context, as in Rayner et al. (1983)] it is assumed that the reader will, upon reading the definite NP, "create" a single unique referent in her discourse model, thus biasing her toward the simple NP analysis.

that, since the burglar planned to blow open a safe, he's going to have to pick one to blow open first. The reader will expect to be told *which safe* the burglar decided to blow up. Assuming that the subject begins the target sentence with the specific goal of discriminating between the entities of this suspended minimal pair, immediate effects of an NP-attachment expectation should be observed. This idea that discourse may produce expectations for greater specificity in particular aspects of upcoming information makes unnecessary the distinction between presuppositions of the simple and complex definite NP.⁴ This proposal also makes the prediction that setting up a minimal pair of *events* should increase the preference for VP-attachment.⁵ Extending the logic above, a context that introduces two possible, but yet-to-occur, events should create an expectation that they will be distinguished in the description of subsequent related events. Attaching a PP to the VP is a common way to convey more detail of the event.

The Experiment

Recent findings by Taraban & McClelland (1988) have also warranted a theory of syntactic disambiguation that is very different from the single encapsulated rule (Minimal Attachment) proposed by Frazier and colleagues. Taraban & McClelland (1988) argue that, in conjunction with certain prepositions, some verbs may produce strong expectations for particular thematic roles that would be violated by attaching the PP to the verb. The result would be an indirect preference for NP-attachment, even in the absence of a context. Unlike Minimal Attachment, however, this hypothesis does not explicitly exclude simultaneous influences from context. To experimentally support this hypothesis, Taraban & McClelland (1988) demonstrated that the thematic role biases of several verb - noun - preposition combinations, as indicated by sentence completion and rating tasks, accurately predicted attachment biases in self-paced reading. They constructed a small

⁴ In fact, some preliminary findings (Spivey-Knowlton & Sedivy, in preparation) indicate that, in the absence of context, Altmann & Steedman's stimuli have a significant, though smaller, VP-attachment bias *even when the NP is indefinite*. As indefinite NPs carry no presuppositions, Referential Theory predicts *no attachment preference* for such sentences.

⁵ Because it has been argued that tense of the verb can be treated as referential (Webber, 1988), it is conceivable that the VP referents in context may act as individuals to one of which the VP is referring. This may be the basis upon which a reformulation of the principle of referential support may account for a VP-attachment bias due to 2 VP referents in context.

set of novel sentences to compare with those of Rayner et al. (1983), and showed that the attachment preferences were nearly equal and opposite.

With the opposition of these two groups of sentences in mind, this experiment compared contexts that contained one NP referent and one VP referent with contexts that contained one NP referent and *two VP referents*. According to discourse-driven structural expectations, the latter context should bias the reader toward a VP-attachment. To examine effects of these contexts on sentences that have context-free VP-attachment biases and those that have context-free NP-attachment biases, stimuli from Rayner, et al. (1983) and Taraban & McClelland (1988) were used as target sentences. Because Taraban & McClelland's stimuli were specifically constructed to have thematic role expectations that bias the reader toward an NP-attachment, and Rayner et al. stimuli were not constructed with such thematic factors in mind, it was possible that the contexts might have different effects on the two groups of stimuli. Nonetheless, evidence for an increase in preference for VP-attachment that is due to having two VP referents in context would, at the very least, require an extension of the principle of referential support.

Method

Subjects. Twenty-four undergraduates of the University of California, Santa Cruz participated in the experiment for course credit. All subjects were native English speakers.

Stimuli and Design. Sixteen context pairs were constructed to accompany sixteen target sentence pairs. Eight of the target pairs were taken from Rayner, et al. (1983) and eight were taken from Taraban & McClelland (1988). Referent order was counterbalanced across stimuli to avoid a referent recency effect (Clifton & Ferreira, 1989). Sixteen filler stimuli were constructed with contexts and targets that were superficially similar to the experimental stimuli. To minimize the difference between the two contexts, the 1-VP-referent context was formed by substituting the second candidate referent with an alternate verb. To avoid lexical priming effects, the structurally ambiguous preposition in a target sentence was not included in the corresponding context paragraph. See Figure 1.

Procedure. Contexts were presented in full on an 80-column computer screen by an IBM PC. Target sentences were subsequently presented in a non-cumulative word-by-word self-paced moving window fashion (Just, Carpenter & Woolley, 1982). Order of

CONTEXT:

One day on the subway, a kid got on carrying a weapon in each hand. He almost hit someone by swinging a whip and pretended to *threaten/hit* someone else using a baseball bat. Then he started to approach the people sitting next to me. There was a girl who had a wart and a boy who had a scar.

TARGET:

That kid hit the girl with a *wart/whip* before he got off the subway.

Figure 1. The 2 VP referent context contains *hit* and *hit*, while the 1 VP referent context contains *hit* and *threaten*. The NP attached target had *wart*, and the VP attached version had *whip*. This target sentence was taken from Rayner et al. (1983).

experimental and filler trials was randomized with the first 6 trials being fillers, to give the subject practice. An open-ended comprehension question followed each trial. Reading times were recorded for the disambiguating noun (*wart* or *whip*) and for the next four words.

Results

A repeated-measures analysis of variance (ANOVA) was computed for the answers to the comprehension questions. A main effect of Target Attachment was observed such that NP attachments produced more errors (9.4%) than did VP attachments (1.6%): $F(1,16)=10.99$, $p<.005$; $F(1,8)=11.67$, $p<.005$. This off-line result is consistent with the minimal attachment principle (Frazier, 1978, 1987).

Data for trials in which the subject gave an incorrect answer to the comprehension question (5.5%) were excluded from the reading time analysis. Trials in which reading time was greater than 2.5 standard deviations from the mean for a given word position (2.1%) were also excluded.

A repeated-measures ANOVA collapsing across recorded word positions revealed a hint of a three-way interaction of Sentence Source (Rayner et al. / Taraban & McClelland) X Context Bias (1-VP-Referent / 2-VP-Referents) X Target Attachment (NP-attached / VP-attached). It appeared that an increase in the VP attachment bias due to 2 VP referents in context was evident in Rayner et al.'s sentences but not in Taraban & McClelland's sentences. However, planned comparisons between means showed no significant differences in reading time.

Discussion

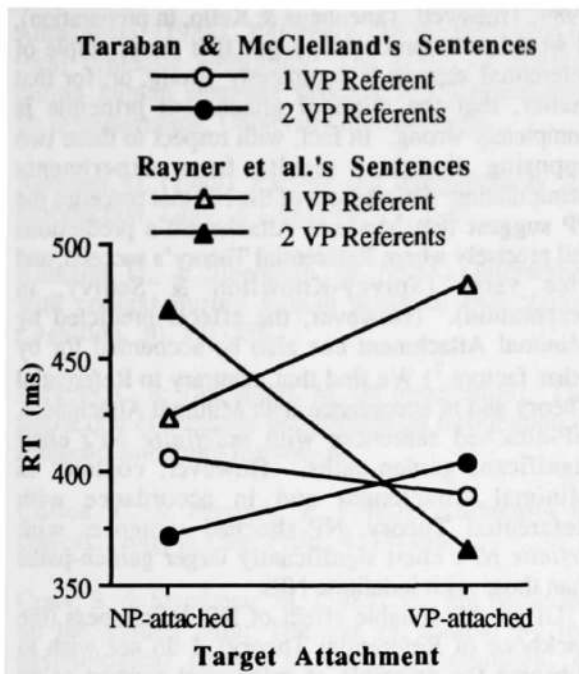


Figure 2. The effect of a second VP referent in context on reading times of the disambiguating noun in NP-attached and VP-attached sentences. Only small differences are seen with Taraban & McClelland's sentences. But with Rayner et al.'s sentences, a robust interaction between context and attachment is observed. This difference in context effect between the stimulus groups is evidenced in a significant three-way interaction (see text).

To test for an immediate effect of context, I conducted an analysis of reading times for the disambiguating noun. This showed a more robust interaction of Sentence Source X Context Bias X Target Attachment: $F(1,16)=6.62$, $p=.02$; $F(1,8)=4.34$, $p=.07$ (see Figure 2). Planned comparison protected t-tests indicated that: 1) with Rayner et al.'s sentences, a significant VP attachment bias (NP-attached vs. VP-attached) was observed in the 2 VP referent context ($p<.01$) but no significant attachment bias was observed in the 1-VP-referent context, and 2) Rayner et al.'s VP-attached targets were read faster when preceded by a 2-VP-Referent context than when they were preceded by a 1-VP-Referent context ($p<.01$).

The nature of the differential effects between stimulus groups complicates interpretation of the data. I initially expected that the effect of the second VP referent would be most visible in Taraban & McClelland's stimuli because, in the absence of context, they have an NP-attachment bias. Curiously, context had no effect on the reading times of these sentences. In contrast, reading times for Rayner et al.'s sentences, that have a VP-attachment bias outside of context, were modulated by the number of VP referents in context.

A possible, though post hoc, interpretation of this interaction between context effect and sentence source derives from a difference in thematic relations that characterize the VP-attached versions of the two groups of sentences. In six of the eight sentences from Rayner et al., the VP-attachment entails the thematic role type, *instrument*. Thus, the corresponding contexts contained two like events that were distinguished by their instruments (see Figure 1). Conversely, of the several thematic roles that characterize the VP-attachments of Taraban & McClelland's stimuli, only one sentence involved an instrument role. It is possible that, in context, minimal pairs of yet-to-occur events that are distinguishable by their *instruments*, rather than by location, manner or time, are more effective in producing a conceptual expectation for greater specification of a related VP. Whether one accepts this post hoc interpretation of the differences observed between sentence sources or not, the fact remains that an increase in VP-attachment preference is observed in Rayner et al.'s stimuli due solely to introducing a second VP referent in context.

As stated before, a contextual influence from having two yet-to-occur events in the discourse does not falsify the principle of referential support. By expanding its notions of definite reference and pragmatic presupposition to allow verb phrases to act in much the same way as noun phrases, Referential Theory might be able to account for this effect. Such expansion, however, is just a smaller version (though evolving in the same direction) of exactly what I am proposing. Gradually expanding a highly-restricted theory each time an effect is observed that it previously ruled out is, arguably, a too narrowly focused approach to the problem that may prevent us from considering other fruitful interpretations.

Alternatively, a much broader theory may better characterize the solution space with which we have to work. The less restrictive theory of discourse-driven structural expectations readily predicts appropriate effects from 2-NP-referent contexts and 2-VP-referent contexts. This theory posits that when an explicit conceptual uncertainty (such as a minimal pair of entities or events) is introduced into context, a reader

may develop an expectation for conceptual disambiguation that would select, on pragmatic grounds, between alternatives of a subsequent ambiguity. If the reader encounters a syntactic ambiguity, of which only one alternative promises to disambiguate the conceptual uncertainty, that syntactic alternative will be the preferred structural assignment.

As an example of when modification of the VP would be preferred on grounds of conceptual uncertainty, consider a context in which there is a little boy, Jimmy, and he is having dinner with the Queen of England. Jimmy's mother has instilled in him an overwhelming fear of the consequences of him using the wrong utensil at the wrong time in the presence of royalty. As desert comes around, he breaks into a cold sweat because he can't remember whether to eat his cake using the fork farthest left from the plate or to eat it using the fork above the plate. (Here, we have the two possible events between which a distinction is crucial, at least to Jimmy and, perhaps, the Queen.) If we later read, "*Finally, Jimmy ate some cake...*", we clearly expect to be told the important discriminating information, not which cake he ate, but which fork he used.

Each of these minimal pair contexts no doubt has a different degree of influence on attachment preference. The principle of referential support, however, has no obvious way of accounting for this variability. The degree of mutual exclusivity between two entities or events is likely, according to the theory of discourse-driven structural expectations, to modulate the strength of the context effect.⁶ For an extreme example, in a 2-NP-Referent context that contains two guns and someone intends to pull the trigger on someone else, a case of one gun having bullets and the other having blanks would produce a stronger expectation for entity (or referent) distinction (hence, modification of the NP "*the gun*") than would a case where one gun has .32 caliber bullets and the other has .38 caliber bullets. "*John picked up two guns; one that had bullets and one that had blanks. He pointed them both at me. Then, he pulled the trigger of the gun with...*"

In light of very recent evidence for numerous simultaneous constraints on syntactic ambiguity resolution, such as availability of syntactic alternatives, frequency of co-occurrence, semantic plausibility, NP definiteness, referential and temporal contexts, noun animacy, and verb subcategorization information (Burgess & Tanenhaus, in preparation; Hindle & Rooth, 1990; McDonald, 1992; Pearlmutter & McDonald, 1992; Spivey-Knowlton & Sedivy, in

preparation; Spivey-Knowlton et al., 1992; Trueswell & Tanenhaus, 1991, this volume; Trueswell, et al., 1989; Trueswell, Tanenhaus & Kello, in preparation), it would be imprudent to argue that the principle of referential support is completely wrong, or, for that matter, that the minimal attachment principle is completely wrong. In fact, with respect to these two opposing theories, results from experiments manipulating definiteness of the NP that precedes the PP suggest that Minimal Attachment's predictions fail precisely where Referential Theory's succeed, and vice versa (Spivey-Knowlton & Sedivy, in preparation). (However, the effects predicted by Minimal Attachment can also be accounted for by other factors.⁷) We find that, contrary to Referential Theory and in accordance with Minimal Attachment, NP-attached sentences with *indefinite NPs* elicit significant garden-paths. However, contrary to Minimal Attachment and in accordance with Referential Theory, NP-attached sentences with *definite NPs* elicit significantly larger garden-paths than those with indefinite NPs.

Given this reliable effect of NP definiteness (the backbone of Referential Theory), I do not wish to subsume the principle of referential support as an epiphenomenon of discourse-driven structural expectations. The two theories, however, are by no means completely separable. Finding and testing those places where the theories clearly diverge (e.g., NP definiteness, continuous degree of constraint) will be a challenge for the future.

This examination of discourse-driven structural expectations is in its initial stages and needs a great deal more refinement. Future work will require independent measures of the specificity of the expectation set, such as sentence completion and rating tasks. The resulting normed range of strongly-constraining to weakly-constraining contexts must then map appropriately onto degree of context effect in on-line reading time results. Moreover, demonstrating context effects with target sentences that contain *indefinite NPs* would be compelling evidence for a theory that is broader than the principle of referential support. The present result is a first step in finding the applicability of one such "broader theory" and developing ways to test it. This line of work has eventual implications for the issue of modularity in language processing and for the general goal of finding what sources of information are important in sentence processing at any stage, early or late.

⁶ The greater effect from events that are distinguished by instrument roles, mentioned above, may be due to their mutual exclusivity being more salient than that of events that are distinguished by location, manner or time roles.

⁷ Minimal Attachment's predictions for PP-ambiguities will often coincide with other locally-determined accounts of attachment preference, such as thematic role expectations (Taraban & McClelland, 1988) and frequency of co-occurrence (Hindle & Rooth, 1990).

References

- Altmann, G. (1986). Reference and the resolution of local syntactic ambiguity: The effect of context during human sentence processing. Unpublished Ph.D. dissertation, University of Edinburgh.
- Altmann, G. (1989). Parsing and interpretation: An introduction. *Lang. & Cog. Proc.*, 4, 1-19.
- Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, 30, 191-238.
- Burgess, C. & Tanenhaus, M. (in preparation). The interaction of semantic and parafoveal information in syntactic ambiguity resolution.
- Clifton, C. & Ferreira, F. (1989). Ambiguity in context. *Lang. & Cog. Proc.*, 4, 77-104.
- Crain, S. & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological parser. In Dowty, Karttunen, & Zwicky (Eds.) *Natural language parsing*. Cambridge: Cambridge U. Press.
- Ferreira, F. & Clifton, C. (1986). The independence of syntactic processing. *JML*, 25, 348-368.
- Fodor, J. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge: MIT Press.
- Frazier, L. (1978). On comprehending sentences: syntactic parsing strategies. Unpublished Ph.D. dissertation, University of Connecticut.
- Frazier, L. (1987). Theories of syntactic processing. In Garfield (Ed.) *Modularity in knowledge representation and natural language processing*. Cambridge: MIT Press.
- Frazier, L., Clifton, C. & Randall, J. (1983). Filling gaps: Decision principles and structure in sentence comprehension. *Cognition*, 13, 187-222.
- Hindle, A. & Rooth, M. (1990). Structural ambiguity and lexical relations. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*.
- Just, M., Carpenter, P. & Woolley, J. (1982). Paradigms and processes in reading comprehension. *JEP:General*, 111, 228-238.
- McDonald, M. (1992). Probabilistic constraints and syntactic ambiguity resolution. Manuscript submitted for publication.
- Pearlmutter, N. & McDonald, M. (1992). Garden-paths in "simple" sentences. Presented at the 5th Annual CUNY Sentence Processing Conference.
- Rayner, K., Carlson, M. & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *JVLVB*, 22, 358-374.
- Spivey-Knowlton, M. (1991). Discourse-driven expectations in the on-line resolution of syntactic ambiguity. Unpublished B.A. thesis, UC Santa Cruz.
- Spivey-Knowlton, M. & Sedivy, J. (in preparation). The effect of noun phrase definiteness on parsing preferences: An interactive constraints account.
- Spivey-Knowlton, M., Trueswell, J. & Tanenhaus, M. (1992). Context effects in parsing reduced relative clauses. To appear in *Canadian Journal of Psychology: Special Issue*.
- Taraban, R. & McClelland, J. (1988). Constituent attachment and thematic role expectations, *JML*, 27, 597-632.
- Trueswell, J. & Tanenhaus, M. (1991). Tense, temporal context and syntactic ambiguity resolution. *Lang. & Cog. Proc.*, 6, 303-338.
- Trueswell, J. & Tanenhaus, M. (this volume). Consulting temporal context in sentence comprehension: Evidence from the monitoring of eye movements in reading. Presented at the 14th Annual Conference of the Cognitive Science Society. Bloomington, IN.
- Trueswell, J., Tanenhaus, M. & Garnsey, S. (1989). Semantic influences on parsing: Use of thematic role information in syntactic disambiguation. Presented at the 2nd Annual CUNY Sentence Processing Conference.
- Trueswell, J., Tanenhaus, M. & Kello, C. (in preparation). Evidence of the immediate use of verb subcategorization information in sentence processing.
- Webber, B. (1988). Tense as a discourse anaphor. *Computational Linguistics*, 14, 61-73.

Acknowledgements: Part of this work was supported by a Sproull Graduate Fellowship from the University of Rochester. I would like to thank Bill Farrar, Ray Gibbs, Ken McRae, Mark Steedman, Mike Tanenhaus, and John Trueswell for helpful comments and discussions. The experiment described was part of a B.A. honors thesis in Psychology at UC Santa Cruz.