Contingent frequency effects in syntactic ambiguity resolution

Cornell Juliano
University of Rochester
Department of Psychology
Rochester, NY 14627
cornell@psych.rochester.edu

Michael K. Tanenhaus
University of Rochester
Department of Psychology
Rochester, NY 14627
mtan@psych.rochester.edu

Abstract[1]

We investigated contingent frequency effects in syntactic ambiguity resolution in three self-paced reading experiments. Experiment 1 demonstrated that the frequency with which *that* occurs as a determiner or a complementizer in different syntactic environments predicts readers initial parsing preferences. Experiments 2 and 3 demonstrated frequency and regularity effects for reading *the* and *that* after different types of verbs that are similar to effects that have been well-documented in word recognition

Introduction

Although the frequency of occurrence of linguistic forms clearly affects how those forms are processed, there has been relatively little exploration of frequency effects in syntactic processing. However, several recent computational and psycholinguistic studies have shown clear effects of frequency on syntactic ambiguity resolution (Hindle & Rooth, 1990; Mitchell & Cuetos, 1991; MacDonald, 1993; Trueswell, Tanenhaus & Kello, in press). These studies suggest that the frequency with which structures occur in different environments and the frequency of lexical co-occurrence patterns each influence ambiguity resolution. These frequency effects have important implications for understanding the mechanisms underlying syntactic ambiguity resolution. In particular, they suggest that parsing preferences that have been attributed to structurally based-decision principles (e.g., Minimal Attachment) may instead be due to the frequencies with which words and structures co-occur in the language. An explanation along these lines is generally consistent with constraint-based approaches to ambiguity resolution (cf., MacDonald, 1992; Trueswell, Tanenhaus & Kello, in press; Spivey-Knowlton, Trueswell & Tanenhaus, in press).

The work presented here is a preliminary report on a project in which we have been using the word *that* to investigate the role of different types of frequency in syntactic ambiguity resolution. We show that preferences to interpret *that* as a complementizer or as a determiner are determined by its frequency of occurrence in different syntactic environments. This has general implications for studies of syntactic ambiguity resolution since many syntactic ambiguities hinge upon categorically ambiguous words. We also demonstrate a frequency by regularity interaction involving *that* that is similar to those that have been well-documented in word recognition research This result supports an alternative type of explanation for data patterns that are often interpreted as support for parsing models that ignore lexically-specific constraints in making initial syntactic commitments

The word *that* has several properties which make it useful for exploring the role of frequency in syntactic ambiguity resolution. *That* is a relatively frequent word and thus statistics based on even a small corpus are likely to be reliable. *That* is also multiply ambiguous as the examples in (1-5) illustrate.

(1) *John believed* ***that.*** (Pronoun)
(2) *John believed* ***that*** *man.* (Determiner/NP-comp)
(3) *John believed* ***that*** *man was lying.*
(Determiner/S-complement)
(4) *John believed* ***that*** *Bill was lying.*
(Complementizer/S-complement)
(5) *John told the man* ***that*** *Sue was dating .*
(Relative pronoun).

Most importantly, the frequency of occurence for the different readings of *that* varies as a function of syntactic and lexical environment. As a result *that* provides a useful tool for investigating contingent frequency effects. For example, an analysis of the Brown Corpus reveals that at the beginning of a sentence *that* is most often a pronoun (54%), then a determiner (35%) and finally a complementizer (11%). But when it immediately follows a verb *that* is most often a complementizer (93%), then a determiner (6%) and least frequently a pronoun (1%).

Our studies focused on the ambiguity between *that* as a determiner and *that* as a complementizer. The first experiment examined whether readers' preferences to parse *that* as a determiner or a complementizer are contingent on the syntactic environment in which it

[1]This research was partially supported by NIH grant HD-27206.

Table 1. Average reading time in milliseconds for each scored word position in Experiment 1.

		THAT	Adjective	Noun	Auxiliary
Condition		*that*	cheap	hotel(s)	is/are
Verb-*that*	Sg. Noun	414	453	484	492
	Pl. Noun	409	470	479	433
Initial-*that*	Sg. Noun	399	456	454	438
	Pl. Noun	381	461	479	563

occurs. The second two experiments examined the interaction of verb frequency and lexical co-occurrence in syntactic ambiguity resolution.

Experiment 1

We conducted a self-paced reading study using sentences such as those illustrated in the examples in (6). In each of the four sentence conditions, *that* is temporarily ambiguous between a determiner and a complementizer. It is disambiguated by the following noun.

(6) a. *The lawyer insisted that experienced diplomat would be very helpful.* (Det/S-comp)
b. *The lawyer insisted that experienced diplomats would be very helpful.* (Comp/S-comp)
c. *That experienced diplomat would be very helpful to the lawyer.* (Initial Det/NP)
d. *That experienced diplomats would be very helpful made the lawyer confident.* (Initial Complementizer/S-comp)

When the noun is singular, *that* is a determiner in an NP such as *that experienced diplomat....* When the noun is plural, *that* is disambiguated as a complementizer that introduces a sentential complement, as in *that experienced diplomats....* Reading times to the noun and the immediately following words can be used to determine whether *that* was initially processed as a determiner or as a complementizer. Long reading times should result when the number of the noun is inconsistent with the reader's initial preference.

If readers take into account the fact that *that* is usually a complementizer when it follows a verb (contingent frequency) but usually a determiner when it begins a sentence, then there should be an interaction between the type of the sentence and the number of the noun. Plural nouns should be read faster than singular nouns for *that* after a verb (6a>6b), whereas singular nouns should be read faster than plural nouns for sentence initial *that* (6c<6d). In contrast, a delay strategy (Frazier & Rayner, 1987) or a strategy based upon simple frequency would not predict this interaction.

Method

Thirty two subjects from the University of Rochester participated for course credit. All were native speakers of English.

Twenty sets of four sentences each were generated from twenty *that-adjective-noun* phrases. Each critical trial consisted of the target sentence followed by a sentence that was a natural continuation of the idea expressed in the first. Filler trials were created in the same fashion. Four lists were constructed using the twenty target trials, which were pseudo-randomly combined with 42 distractor trials, for a total of 62 trials in each list. At least one distractor appeared before each experimental trial.

Stimuli were presented on a color monitor attached to an IBM PC or an IBM compatible. The monitor was set to display 80 characters per line in a moving window presentation (Just, Carpenter & Wooley, 1982). Each trial began when the subject pressed the **START** button on the response box. This caused the entire text to be displayed with each alpha-numeric character replaced by a dash(-) but spaces and punctuation left the same. Subjects controlled the word-by-word presentation of the stimuli by pressing the **SCROLL** button on the response box. Each button press caused the next word of the text to appear and the previous word to be replaced by dashes. Comprehension questions were asked after a quarter of the trials. Subjects were given a practice session of seven trials before the actual experiment. The experiment lasted about 35 minutes.

Results.

Only data from subjects who answered 80% or more of the questions correctly were analyzed. Separate ANOVAs with four factors were performed on reading times for each subject and item. The factors were List(four lists) or Item Group(four groups), Sentence type(verb followed by "that" or initial-that), Noun type (singular or plural) and Word position(that, adjective, noun, verb). Table 1 presents the reading times to the ambiguous word *that* and each of the three words that followed.

Reading times to the word after the disambiguating noun (the auxiliary verb) showed a clear interaction

between the number of the noun and the position of *that* ($F1(1,32)=18.06$, $MSe=6017$, $p<0.01$; $F2(1,16)=5.86$, $MSe=10352$, $p<0.05$). When *that* followed a verb, reading times at the auxiliary verb were longer in the singular noun condition compared to the plural noun condition ($F1(1,32)=6.24$, $MSe=9857$, $p<.05$; $F2(1,16)=3.04$, $MSe=11225$, $p=0.1$). This indicates that subjects initially took *that* to be a complementizer. When they read a singular noun, like *hotel*, the word *that* had to be reanalyzed as a determiner. The opposite pattern occurred when a sentence began with *that*. Reading times following the singular noun were faster than those after the plural noun ($F1(1,32)=22.28$, $MSe=12747$, $p<0.01$; $F2(1,16)=19.26$, $MSe=8211$, $p<0.01$). This demonstrates that subjects initially assumed *that* was a determiner which introduced an NP in a main clause. In the plural noun condition, subjects had to reanalyze *that* as a complementizer and reanalyze the noun phrase as being within a sentence complement. Thus, both the parse of *that* and the sentence structure had to be revised for an initial *that,* whereas only the parse of *that* had to be corrected in cases where *that* followed a verb. This provides a likely explanation for why the magnitude of the revision effects were larger in the *Initial-that* condition than in the *Verb-that* condition.

In sum, this experiment demonstrated that readers' initial preferences for interpreting the lexically ambiguous word *that* are contingent on the syntactic environment in which the word occurred, with the ambiguity resolved in accordance with the most common pattern in the language.

Experiments 2 and 3

It is well-documented that readers and listeners often have clear preferences for interpreting locally ambiguous sentences. For example, most people would complete a fragment such as *The lawyer remembered the...* so that *the* is part of an NP complement of the verb (e.g., *The lawyer remembered the trial this morning*) rather than as part of an NP that is the subject of a S-complement (e.g., *The lawyer remembered the trial was delayed*). These preferences are reflected in temporary reading difficulties in sentences with sentence complements. One explanation is that readers are following a structurally-based decision principle at the point of the ambiguity. However an alternative is that these preferences reflect statistical patterns in the language, e.g., *the* following a verb is more often the start of an NP (direct object) complement.

Readers also experience difficulty in reading a sentence complement even when it is preceded by a verb that does not allow an NP complement, e.g., *insisted* (Ferreira & Henderson, 1990). This result has been interpreted as evidence that verb-specific syntactic constraints are ignored in initial parsing decisions, in favor of general category based parsing principles. However, there is an alternative explanation based upon the type of frequency by regularity interactions found in word recognition research.

The generalization that emerges from the word recognition literature is that there is a processing penalty associated with processing a low frequency exception word but not a high frequency exception.

For example, the mapping between spelling and pronunciation in English follows clear patterns, although there are many exceptions. For example, the sequence **C**(onsonant)**+AVE** is usually pronounced with a long *a* as in *gave,* though there are words such as *have,* which violate the common pattern or "rule". **Regular** words which conform to the most frequent pattern in the language are read aloud (named) more rapidly than **exception** words which do not conform to the common pattern. However, the penalty for naming an exception word that violates the pattern is modulated by its frequency of occurrence. There is little or no penalty for high frequency words, whereas there is a large penalty for low frequency exception words. This frequency by regularity interaction has been replicated numerous times and it has been convincingly modeled by constraint-based models of word recognition (e.g., Seidenberg & McClelland, 1989). Similar frequency by regularity interactions also occur when subjects are timed as they generate the past tense of regular (weak) and irregular (strong) verbs. In general, subjects can generate the past tense for a regular verb more rapidly than the past tense for an irregular verb. However the penalty for being an exception is again modulated by frequency (Prasada, Pinker & Snyder, 1990; Daugherty & Seidenberg, 1992).

Let's return now to the difficulty that readers have with sentence complements. Trueswell, Tanenhaus & Kello (in press) found two results that support an explanation based on contingent frequencies. First, processing difficulty for sentence complements with sentence-complement (S-bias) verbs such as *insisted* begin immediately at the following NP (e.g., *insisted the trial*), as opposed to at the disambiguating verb, as is the case with verbs such as *remembered.* Secondly, the magnitude of the difficulty was correlated with whether the verb was typically followed by the complementizer *that* when it was used in a sentence complement construction, as determined by completion norms. They labeled this as the verb's *that*-preference. If *that*-preference is correlated with verb frequency then these results can be seen as directly analogous to frequency effects in word recognition. The sequence *verb-determiner* typically occurs when the determiner begins an NP complement. For verbs that do not allow an NP complement but do allow a sentence complement, the determiner begins an NP that is the subject of a sentence complement. This pattern occurs frequently only for verbs that have low *that*-preferences.

Table 2. Reading time predictions at the determiner and at the adjective following the verb.

Phrase	Type	Penalty
insisted that (Comp)*	Regular	None
insisted the (Det. in S-Comp)	Exception	Correlated with Verb Freq.
visited that (Det)	Low Frequency Exception	Large
visited the (Det. in NP-Comp)	Regular	None

* Prior to the noun, *that* is taken to be a complementizer

Thus, these verbs should behave like high frequency exceptions, i.e., they should be relatively immune from interference effects due to the regularity pattern in the language. In contrast, verbs with high *that*-preferences should behave like low frequency exceptions. This analysis makes a novel prediction. **NP-only** verbs like *visit* require a following *that* to be a determiner, which runs counter to the regular pattern in the language. In addition, verbs like *visit* are rarely followed by the word *that*. Thus, sequences like *visit that* are analogous to low frequency exceptions and should be difficult to process. There should be strong interference from the fact that *that* is typically a complementizer.

Experiments 2 and 3 were conducted to explore the predictions that we have outlined above. Experiment 2 investigated whether readers would have difficulty processing *that* after a verb when the verb required *that* to be a determiner. In addition, we included a condition where each verb-type was followed by the determiner *those* (e.g., *The lawyer visited those cheap hotels...*)to rule out the possibility that processing *that* would be difficult because the context did not introduce appropriate referential material, i.e. a referent for *that* to refer to. *Those*, like *that*, also has referential properties. If subjects expect a referent then reading times to the area after both *that* and *those* should be elevated.

Experiment 3 replicated and extended the results of Trueswell et al (in press). In order to explore the possibility of a frequency by regularity interaction, we included sentences in which a noun phrase beginning with the determiner *the* followed the verb (Table 3, (8)). After an S-bias verb, like *insisted* , the NP would be the subject of a sentence complement. After an NP-only verb, such as *visited*, it would be the object of the verb. The determiner *the* following a verb most commonly begins an object complement (i.e., NP complement) in English. Thus, *visited the* follows the regular pattern, whereas *insisted the* is the exception. Therefore, we would predict a penalty for sentence-complement verbs followed by a *the*. However, as we pointed out earlier, this penalty is correlated with a verb's *that*-preference. It turns out that there is a strong, negative correlation between the *that*-preference of an S-bias verb and the verb's frequency. Thus, high frequency sentence complement verbs typically are not followed by *that*, whereas lower frequency verbs are. *That*-preference effects can be seen as an example of a frequency by regularity interaction. For **high frequency** S-biased verbs the sequence *verb-the*, where *the* begins a sentence complement, will be a high frequency exception and thus will show little or no penalty for violating the regular pattern in the language. In contrast, **low frequency** sentence complement verbs will show a penalty effect. In other words the exception penalty should be correlated with *that*-preference for individual verbs. These predictions are summarized in Table 2.

We evaluated these predictions in two self-paced reading experiments using the same procedure as was used in the first experiment. Because of space limitations we will describe the experiments together along with an abbreviated presentation of the statistics. All of the differences that we discuss are reliable at $p<.05$ unless noted otherwise. Twenty-eight subjects participated in each experiment. Sample materials for each experiment are presented in Table 3.

Table 3. Sample materials for experiments 2 and 3.

Experiment 2

(7) a *The lawyer insisted that cheap hotel was clean and comfortable.*
b *The lawyer insisted those cheap hotels were clean and comfortable.*
c *The lawyer visited that cheap hotel to stay for the night.*
d *The lawyer visited those cheap hotels to stay for the night.*

Experiment 3

(8) a *The lawyer insisted that cheap hotel was clean and comfortable.*
b *The lawyer insisted the cheap hotel was clean and comfortable.*
c *The lawyer visited that cheap hotel to stay for the night.*
d *The lawyer visited the cheap hotel to stay for the night.*

Results

The results of Experiment 2 are presented in Table 4. As predicted, readers had difficulty with sentences that contained an NP-only verb followed by *that* (low frequency exception condition). This was reflected in extremely long reading times to the determiner *that*

Table 4. Average reading times in milliseconds for each scored word position (Experiment 2).

Condition		Verb	Determiner	Adjective	Noun	Auxiliary
		insisted	that/those	cheap	hotel(s)	was/were
S-bias	that	494	462	548	587	524
	those	497	466	543	573	486
		visited	that/those	cheap	hotel(s)	to
NP-only	that	489	555	619	579	478
	those	483	475	513	558	486

Table 5. Average reading times in milliseconds for each scored word position (Experiment 3).

Condition		Verb	Determiner	Adjective	Noun	Auxiliary
		insisted	that/the	cheap	hotel	was
S-bias	that	368	364	382	453	416
	the	377	378	372	400	368
		visited	that/the	cheap	hotel	to
NP-only	that	372	349	464	453	386
	the	372	353	349	401	374

compared to when it followed an S-bias verb ($F1(1,24)=8.0$, $MSe=14922$, $p<0.01$; $F2(1,16)=6.15$, $MSe=13951$, $p<0.05$). The reading times at the noun and the following auxiliary in the S-bias condition replicated the results presented in Experiment 1. Readers had difficulty in the singular noun region when it was preceded by an S-bias verb followed by *that,* indicating that they had initially parsed *that* as a complementizer. A hint of this difficulty began at the noun, but it did not become reliable until the next word, the auxiliary verb ($F1(1,24)=4.03$, $MSe=5055$, $p=0.05$; $F2(1,16)=5.16$, $MSe=2836$, $p<0.05$). In contrast readers did not have difficulty with the singular noun region following an NP-only verb. Readers apparently used the information provided by the verb to correctly parse *that* as a determiner. Also, a comparison of reading times between the *that* and *those* conditions confirms that processing difficulties are not due to the lack of an appropriate referential context for *that.*

These results clearly establish that readers are rapidly using verb-specific information, yet they are paying a penalty when the word that follows the verb requires a parse that runs against the regular pattern in the language. Let's now consider the results of Experiment 3, which directly examined whether this exception penalty would be modulated by frequency. The results are presented in Table 5.

In Experiment 3 subjects again had difficulty when *that* followed an NP-only verb (a low frequency exception phrase) compared to when it followed an S-bias verb. A large difference in reading time occurred at the adjective ($F1(1,24)=9.84$, $MSe=9535$, $p<0.01$; $F2(1,16)=19.25$, $MSe=3496$, $p<0.01$). This effect is the same one that we saw in Experiment 2 except that it began one word earlier in that experiment. Readers also had more difficulty after an S-bias verb followed by the determiner *that* compared to S-bias verbs followed by the determiner *the.* This was expected because as in Experiment 1, readers should have initially parsed *that* as a complementizer and then had to revise their analysis when the noun was singular.

The data of primary interest are the reading times to the determiner after the S-bias verbs. Recall that we predicted that reading times for the S-bias condition should vary according to the verb's frequency which is negatively correlated with its *that*-preference. Thus, a frequency by regularity interaction would manifest itself as a positive correlation with *that*-preference. To factor out effects of individual lexical items we correlated *that*-preference with the difference between reading times to *the* and *that* when they followed an S-bias verb. There was a strong correlation between the difference in reading times and *that*-preference ($r^2=0.82048$, $F=31.99$). The correlation with the log frequency of the verb was also highly reliable $r^2=0.70242$, $F=14.16$). There were also reliable correlations at the adjective and for the determiner and the adjective combined.

In order to see graphically that these correlations represent a frequency by regularity interaction, we grouped the verbs into high and low frequency categories. Figure 1 plots the average reading times at *the* for low and high frequency verbs of each type. Here, NP-only verbs represent the common pattern and S-bias verbs represent the exception pattern. We see that for the low frequency verbs, the word *the* is read faster after NP-only verbs (regular pattern) than for S-bias verbs (exception pattern). For high frequency verbs the reading times for *the* are the same, regardless of whether the pattern is regular or an exception.

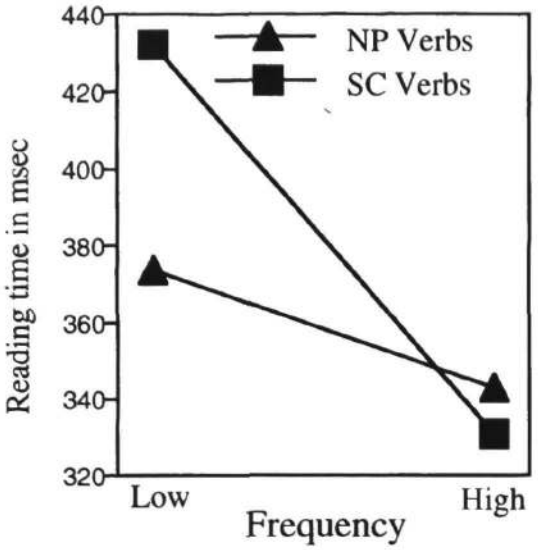

Figure 1. Reading time to the word *the* as a function of the frequency and type of the preceding verb.

General Discussion

Though the work presented here focused primarily on a single word and was limited to investigations of a few structures, the results are likely to have broad implications for research in sentence processing. Many structural ambiguities are like the sentence complement/NP-complement ambiguity, in that a structural ambiguity hinges upon a lexical category ambiguity. Experiment 1 demonstrates that a complete account for parsing preferences for these types of ambiguities will require taking into account contingent frequencies. Likewise Experiment 2 demonstrates that processing difficulties that are often interpreted as due to local garden-paths and rapid reanalysis may be better understood as contingent frequency effects similar to those that are well established within the word recognition literature. Processing difficulties are likely to be observed whenever the structure being processed runs counter to the regular pattern in the language. It is important to note that we are not proposing a set of frequency-defined strategies to replace structurally defined decision principles. We expect that contingent frequency effects like the ones that we have identified will emerge as a natural by-product of learning models with a distributed representation. An investigation of this hypothesis is in progress.

Acknowledgments

We are grateful to John Trueswell for his invaluable contributions to this work.

References

Daugherty, K. & Seidenberg, M.S. (1992). Rules or connections? The past tense revisited. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 259-264. Bloomington, IN.

Ferreira, F. & Henderson, J.M. (1990). The use of verb information in syntactic parsing: A comparison of evidence from eye movements and word-by-word self-paced reading. *Journal of Experimental Psychology: Learning, Memory and Cognition,* 16, 555-568.

Frazier, L. & Rayner, K. (1987). Resolution of syntactic category ambiguities: Eye movements in parsing lexically ambiguous sentences. *Journal of Memory and Language* 26, 505-526.

1.9

0.9

Hindle M. & Rooth M. (1990). Structural ambiguity and lexical relations. In *Proceedings of the 28th Annual Meeting of the Association of Computational Linguistics.*

Just, M.A., Carpenter, P. A. & Woolley, J. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General,* 111, 228-238.

MacDonald, M. (1992) Probabilistic constraints and syntactic ambiguity resolution. Submitted for publication.

MacDonald, M.C. (1993) The interaction of lexical and syntactic ambiguity. Submitted for publication.

Mitchell, D. C. & Cuetos, F. (1991). The origins of parsing strategies. Unpublished manuscript.

Prasada, S., Pinker, S., & Snyder, W. (1990). Some evidence that irregular forms are retrieved from memory but regular forms are rule generated. *31st Annual Meeting. The Psychonomic Society.*

Seidenberg, M.S. & McClelland, J.L. (1989). A distributed developmental model of word recognition and naming. *Psychological Review*, 96, 523-568.

Spivey-Knowlton, M.J., Trueswell, J.C., & Tanenhaus, M.K. (in press). Context effects in syntactic ambiguity resolution: Discourse and semantic influences in parsing reduced relative clauses. *Canadian Journal of Psychology: Special Issue.*

Trueswell, J.C., Tanenhaus, M.K., & Kello, C. (in press). Verb-specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory & Cognition.*