

Misinterpretations of Garden-Path Sentences: Implications for Models of Sentence Processing and Reanalysis

**Fernanda Ferreira,^{1,4} Kiel Christianson,^{1,3} and
Andrew Hollingworth^{2,3}**

Theories of sentence comprehension have addressed both initial parsing processes and mechanisms responsible for reanalysis. Three experiments are summarized that were designed to investigate the reanalysis and interpretation of relatively difficult garden-path sentences (e.g., While Anna dressed the baby spit up on the bed). After reading such sentences, participants correctly believed that the baby spit up on the bed; however, they often confidently, yet incorrectly, believed that Anna dressed the baby. These results demonstrate that garden-path reanalysis is not an all-or-nothing process and that thematic roles initially assigned for the subordinate clause verb are not consistently revised. The implications of the partial reanalysis phenomenon for Fodor and Inoue's (1998) model of reanalysis and sentence processing are discussed. In addition, we discuss the possibility that language processing often creates "good enough" structures rather than ideal structures.

KEY WORDS: parsing; reanalysis; semantics; syntactic ambiguity.

INTRODUCTION

The goal of research in sentence comprehension is to discover how people understand language. It has proved useful to investigate this question using sentences that contain a temporary syntactic ambiguity. For example, consider the sentence *While Anna dressed the baby spit up on the bed*. A large number of studies have shown that people have trouble understanding this sentence when it is presented visually and without any internal punctuation.

¹ Department of Linguistics, Michigan State University, East Lansing, Michigan 48824.

² Department of Psychology, Michigan State University, East Lansing, Michigan 48824.

³ Cognitive Science Program, Michigan State University, East Lansing, Michigan 48824.

⁴ To whom all correspondence should be addressed. email: fernanda@eyelab.msu.edu

They initially assume that *the baby* is the object of *dressed*, and when they encounter *spit up*, they realize that they have made some sort of error. This initial parse, then, must somehow be revised. Models of reanalysis attempt to explain how the revisions take place. All such models try to explain how it is that the comprehender eventually locates a subject for the main clause verb *spit up*. More generally, the goal of all models of reanalysis is to describe and motivate the mechanisms used by the parser to detect errors, deduce useful information about the nature of the necessary repair from those errors, and ultimately create a successful parse. Of course, reanalysis is not always successful (indeed, the experiments we summarize in this article demonstrate that reanalysis processes succeed even less often than most of us working in the field supposed). Another goal of reanalysis models, then, is to explain why it is that successful revision is possible for some sentences but apparently impossible for others.

One characteristic of garden-path sentences that seems to strongly influence ease of reanalysis concerns the syntactic relationship between the error signal and the head of the phrase that has been misanalyzed (Ferreira & Henderson, 1991, 1998). Consider a slightly modified version of our sentence above: *While Anna dressed the baby that was small and cute spit up on the bed*. The error signal is the word *spit up*, because it tells the parser that the analysis it has computed up to that point is flawed somehow. In a large number of experiments, Ferreira and Henderson demonstrated that reanalysis is more difficult when the head of the misanalyzed phrase (*baby in the baby that was small and cute*) is distant from the error signal. Earlier work (Warner & Glass, 1987) had shown that the *length* of the ambiguous phrase was important; reanalysis was more successful with a short phrase such as *the baby* than a long phrase like *the baby that was small and cute*. Ferreira and Henderson showed that length itself is not the important variable, because reanalysis was as successful for sentences with an ambiguous phrase such as *the baby* as for those with a longer phrase such as *the small and cute baby*. This is because, for both of these phrases, the head of the misanalyzed phrase is in the same position with respect to the error signal—immediately adjacent. What is important, then, is the distance between the head of the phrase and the error signal. Ferreira and Henderson argued that this effect arises because the sentence comprehension system assigns a thematic role to a phrase only upon encountering its head. For garden-path sentences in which the head and error signal are right next to each other, the wrong thematic role has been assigned only momentarily and so can easily be revised; for sentences in which those two elements are far apart, the sentence comprehension system has been committed to the wrong thematic role for a relatively long period of time and, therefore, has more trouble relinquishing it.

The Ferreira and Henderson experiments are important because they are among the few done specifically to investigate reanalysis processes systematically. The experiments we will summarize here briefly (for a more full report, see Christianson, Hollingworth, Halliwell, & Ferreira, in press) add to this body of work by exploring a somewhat different, but ultimately, related question: To what extent does reanalysis lead to the successful interpretation of garden-path sentences? Most work (including Ferreira and Henderson's) has assumed that reanalysis is also an all-or-none phenomenon: People either successfully parse a sentence or they do not (this assumption is revealed, for instance, in the use of dependent measures, such as grammaticality judgments, which allow only one of two responses from participants). The experiments we will describe next focused on the possibility that people might only partially reanalyze garden-path sentences. The experiments take advantage of the head-position effect by demonstrating that the longer people have been committed to the wrong syntactic analysis, the less likely they are to engage in complete and successful reanalysis. After we have highlighted the main results from these experiments, we will focus on their implications for the Fodor and Inoue (1998) model of reanalysis. The last parts of the article explore broader issues in sentence processing and cognitive science and mention some future directions for research on the topic of how and why people misinterpret garden-path sentences.

Experiments Demonstrating that Thematic Roles Assigned Along the Garden-Path Linger

The experiments we conducted were motivated by our concern that previous work on how garden-path sentences are processed had barely addressed how these sentences are actually understood. Reading-time measures yield a record of the processing profile for the constituent words as they are encountered and reread and grammaticality judgments indicate whether a person believed the sentence could be assigned a legitimate syntactic structure overall. Virtually no data existed on the question of whether people understand these sentences the way we presume they should: For instance, for *While Anna dressed the baby spit up on the bed*, people ultimately should end up believing that a baby spit up on some bed and that Anna dressed herself. But do they?

To address this issue, Christianson *et al.* (in press) designed three experiments using the same basic paradigm: Participants were asked to read garden-path sentences or corresponding control sentences. Following each, they answered a YES/NO question and indicated their confidence in that answer. The primary dependent measures, then, are question-answering accuracy and confidence. For most experiments the method of computer

presentation allowed us to collect reading times for the sentences and the questions, but these measures yielded results that are of secondary interest and so will not be described here (see Christianson *et al.* for details).

The first experiment was designed to investigate the fundamental question whether people end up with an appropriate understanding of garden-path sentences. The stimuli are given in (1).

- (1a) While Bill hunted the deer (that was brown and graceful) ran into the woods.
- (1b) While Bill hunted the deer (that was brown and graceful) paced in the zoo.
- (1c) While Bill hunted the pheasant the deer (that was brown and graceful) ran into the woods.

Sentences (a) and (b) are the garden-path versions; (c) is a non-garden-path control (*the pheasant* provides an object for the verb *hunted*, and so there should be little temptation to make *the deer* an object of that same verb). The material in parentheses in (1) was included in the long ambiguous phrase conditions, and excluded in the short ambiguous phrase conditions. The difference between the (a) and (b) versions concerns the plausibility of the misinterpretation. In (a), the misinterpretation is plausible: People hunt deer that are in woods. In (b), the misinterpretation is implausible: People rarely hunt animals that are housed in a zoo. Thus, we can evaluate whether reanalysis processes are more likely to run to completion if the misinterpretation based on the incorrect structure violates the comprehender's pragmatic knowledge.

Any given participant saw only one of these six possible versions. However, there were a large number of sentences such as (1) in total (along with numerous fillers designed to disguise the main purposes of the study) and, therefore, over the course of the experiment, participants experienced each of the six conditions several times. Following each sentence, a YES/NO question was asked. For (1), the question was *Did Bill hunt the deer?* A NO response is technically what is required and so we will label this response "correct" [while acknowledging that a person may legitimately *infer* that Bill hunted a deer even in conditions (a) and (c); we will discuss this issue in more detail below]. A YES response is technically incorrect—the garden-path sentences state only that Bill hunted something. Thus, if participants answer YES when given sentences such as (a) and (b), we have evidence that they did not end up with an appropriate interpretation for the sentences. Furthermore, if they are more likely to incorrectly say YES when the ambiguous phrase is long [that is, when the material in parentheses in (1) is included], then we have further evidence that the tendency to say YES is attributable to the initial syntactic misanalysis.

The sentences were presented one at a time on a computer monitor. Participants read each sentence at their own pace and then pushed a button that caused the sentence to disappear. A question then appeared on the monitor, which the participant answered YES or NO by pushing one of two buttons on a response panel. Then the participant made his or her confidence judgment.

The results from this experiment are shown in Figure 1. First, consider the two rightmost bars, representing the control conditions in which the verb *hunt* has an object (*the pheasant*) and so the misinterpretation should not be tempting. As can be seen, error rates in these conditions are below 15% and provide a baseline measure of performance in the absence of garden-pathing. [The small effect of region length is significant even for these nongarden-path sentences, but smaller than for garden-path sentences. A similar pattern was obtained by Ferreira Henderson (1991)]. Next, consider the leftmost two bars. When the deer is in the woods, people incorrectly say YES almost 30% of the time in the short-phrase condition and over 40% of the time in the long-phrase condition. Finally, the middle two bars represent the data for the condition in which the deer is in the zoo: If the region is short, people rarely think the man is hunting that captive deer; but if it is long, then error rates are over 30%. Another important aspect of these results is the confidence data: People overall were highly confident in their responses, providing average ratings of about 3.2 on a scale of 1 to 4 (with “4”




Fig. 1. Percentage incorrect YES responses for Experiment 1b of Christianson *et al.* (in press). Error bars are 95% confidence intervals.

indicating complete confidence in the answers). In addition, in the garden-path conditions, participants were as confident about their incorrect responses as they were about their correct responses (see Christianson *et al.*, in press, for details). Therefore, it does not appear that comprehenders are simply muddled about the garden-path sentences. The misinterpretations they compute for them are beliefs they hold with fairly strong confidence.

The next experiment was designed to explore the effect further by introducing two innovations. First, rather than using the control condition that was used in Experiment 1, we chose to create non-garden-path conditions by simply reversing the order of the subordinate and main clauses. This control provides an even better baseline measure of the extent to which people might simply think that a man who is hunting might hunt a deer. That is, we expect that when participants are given *The deer ran into the woods while the man hunted*, they might infer that the man was hunting that deer. Therefore, if the tendency to draw that same inference is greater when the clauses are reversed so that the sentence creates a garden-path, then we have an estimate of people's tendency to think the man is hunting the deer based just on their having considered *the deer* as a syntactic object of the verb *hunted*. The second innovation was to ask people two kinds of questions: Not only did we ask whether the man hunted the deer, we also asked whether the deer ran into the woods. This latter question is important, because if people can answer it correctly, then we have evidence that reanalysis processes were initiated, at least far enough to allow the parser to locate a subject for the disambiguating verb.

The stimuli for this experiment, then, are as shown below:

- (2a) While Bill hunted the brown and graceful deer/the deer that was brown and graceful ran into the woods.
- (2b) The brown and graceful deer/the deer that was brown and graceful ran into the woods while Bill hunted.

Notice that we also have changed our phrase length manipulation: Instead of short vs. long, we made both phrases long but positioned the nominal modifiers so that the head noun ended up either right next to the disambiguating verb or not. Recall from the first Section that the long, pronominal adjective conditions (*the brown and graceful deer*) should work just like the short conditions—reanalysis should be easy, compared to the long, postnominal conditions.

The participants read each sentence one at a time at their own pace and then answered one of two questions: either *Did Bill hunt the deer?* or *Did the deer run into the woods?* After answering the question, the participants supplied a rating of their confidence in the answer.

The results from this experiment are shown in Figure 2. Examine the rightmost four bars first. Clearly, people can answer the question whether the




Fig. 2. Percentage incorrect responses for Experiment 2 of Christianson *et al.* (in press). For questions about the subordinate clause, the correct response was NO. For questions about the matrix clause, the correct response was YES. Error bars are 95% confidence intervals.

deer ran into the woods with great accuracy. Importantly, accuracy was affected little by head position and not at all by the garden-path manipulation. Now look at the leftmost four bars. Error rates for the question about whether Bill hunted the deer are high—over 50% across the board. This finding indicates that, regardless of whether people are garden pathed, they have a strong tendency to think that the man hunted the deer. Nevertheless, the tendency is much greater in the garden-path condition, particularly when the head of the misanalyzed phrase is located far from the disambiguating word (that is, the head occurs early in the phrase rather than phrase-finally). As in Experiment 1, we also found that confidence in all responses was high and that in the garden-path conditions confidence for incorrect responses was as high as for correct responses.

This experiment, then, establishes that reanalysis processes get off the ground to some extent. When the parser discovers that it has no subject for the main clause verb in the garden-path structures, it manages to locate one (it steals a noun phrase from the previous clause, as will be detailed in the next section). Yet, having stolen the object of the first verb so as to make it the subject of the second, it does not appear that the initial interpretation is also revised: People answer the question about the man hunting the deer as if that NP were still object of the verb.

Up to now, we have described the misinterpretation effect as any tendency to think that the man is hunting the deer that is greater in the garden-path than in the non-garden-path condition. This definition acknowledges that when someone is intransitively hunting, he or she may be hunting a deer. We set up the last experiment to prevent this sort of inference. To do so, we used a class of verbs sometimes called Reflexive Absolute Transitive (RAT) verbs. These include verbs such as *dress* and *bathe*. Unlike regular optionally transitive verbs, such as *hunt*, the intransitive form of *dress* does not mean that the agent dressed something or other; instead, it means that the agent dressed him- or herself. Thus, if people are given the sentence *While Anna dressed the baby spit up on the bed* and if they are asked whether Anna dressed the baby, they have little option but to say NO, if they correctly reanalyzed the sentence. They must say no because the sentence states that Anna dressed herself, not that she dressed some unspecified thing that could well be the baby.

This experiment consisted of two subexperiments, which differed in the way that the non-garden-path control condition was established. The first of the two was like the experiment described previously: The non-garden-path condition was created by reversing the order of the subordinate and main clauses. The experiment thus had a 2×2 design: The sentence was either garden-pathing or not and the verb was either a normal, optionally transitive verb (such as *hunted*) or a RAT verb. The second of the two subexperiments differed only in that rather than varying clause order, we varied whether or not a comma separated the subordinate and main clauses (always in that order). Thus, the sentence *While Anna dressed, the baby spit up on the bed* was the non-garden-path control, which clearly differs minimally from the garden-path version. This subexperiment had the same 2×2 design.

The results are shown in Fig. 3. As can be seen, the pattern is identical in both subexperiments, so they can be discussed together. Consider the bars representing the conditions in which optionally transitive verbs were used. These data replicate the previous experiments: People tend to think the man hunted the deer, but they do so much more when they have been garden-pathed. Now consider the bars representing the conditions in which RAT verbs were used. In the non-garden-path condition, participants rarely believed that Anna dressed the baby. But in the garden-path conditions, performance is the same as for the optionally transitive verbs. These results, then, provide strong support for our main contention: When people are garden-pathed, they do not give up on the initial interpretation that they built based on an ultimately incorrect syntactic analysis.

To summarize, the experiments demonstrated the following. First, the interpretation based on the initial misparse lingers. Second, the misinterpretation is more likely to linger the longer the parser was committed to the




Fig. 3. Percentage incorrect responses for Experiments 3a and b of Christianson *et al.* (in press). Garden-path sentence for RAT and optionally transitive verbs were compared to two control conditions, a clause order control (matrix-subordinate) and a comma control. Error bars are 95% confidence intervals.

incorrect syntactic structure supporting it (the phrase length/position effect). Third, this misinterpretation effect is somewhat selective. It does not survive if the misinterpretation is implausible (recall that participants are not as likely to think Bill is hunting a deer if that deer is in a zoo). In addition, it is selective in that it is not the entire sentence that is misunderstood. People behave as if they started reanalysis processes and therefore managed to obtain a subject for the main clause; but they often appear not to allow the processes to run to completion, thus producing an improper interpretation of the subordinate clause. In other words, they appear to treat the noun phrase as both the subject of the main clause and the object of the subordinate clause. In the next section, we discuss potential explanations of the misinterpretation effect.

SYNTACTIC LOCUS OF THE MISINTERPRETATION EFFECT

As described above, the experiments from Christianson *et al.* (in press) were designed to rule out the possibility that the tendency to answer questions incorrectly could be attributed to pragmatic knowledge. The general reasoning argument is essentially that full reanalysis takes place and the high numbers of incorrect answers to the comprehension questions stem

from pragmatic inference: The man is hunting something, and the deer, which is a common quarry for hunters, is running away; therefore, the man is probably hunting the deer. The first piece of evidence against a general reasoning story is the strong head position effect in Experiment 2. There is no obvious reason why “the deer that was brown and furry” should be more likely to be inferred to be the quarry of the hunt than “the brown and furry deer,” yet participants were significantly more likely to derive an incorrect interpretation from the former construction than from the latter.

The second test of the general reasoning explanation was to eliminate the syntactic ambiguity in the sentences without changing their informational load. Two independent manipulations were used to this end: reversing the clause order (from Subordinate-Matrix to Matrix-Subordinate order) and separating the clauses with a comma in the Subordinate-Matrix order sentences. Figure 3 shows that both sentence types resulted in significantly fewer incorrect “YES” responses than the garden-path sentences. While we assume that conversational implicature and the Gricean maxim of relation (relevance) (Grice, 1975) bear on people’s final interpretations of such sentences, the data show syntactic effects over and above the baseline inferences.

The final test of the general reasoning explanation was the use of the small class of verbs termed Reflexive Absolute Transitive (RAT) verbs. As described above, the syntax of these verbs is such that, even in the intransitive form, they have a specified object/theme—the subject (i.e., *Anna dressed* means that Anna dressed herself). Therefore, it is much more difficult to argue that comprehenders misinterpret the subordinate clause in the garden-path conditions simply because the subject of the main clause can be pragmatically inferred to be the target of the action. As Figure 3 shows, participants also had trouble reaching a correct interpretation of RAT verbs in the garden-path conditions, despite the fact that full reanalysis would yield an unambiguous syntactic structure that would block incorrect inferences. These varied sources of data taken together make it clear that pragmatic inferencing cannot account for the misinterpretation effect in the garden-path conditions.

Another potential nonsyntactic explanation of our results, which also assumes full reanalysis has taken place, is one in which the reanalyzed syntactic structure (especially that of the subordinate clause), although fully computed, is never mapped onto the semantic representation to replace the original incorrect interpretation. Such an explanation begs the question of why it would not, especially since part of the syntactic structure (the matrix clause) does proceed on to the semantics, as shown by the results of our second experiment. Similarly, one might argue that the reanalyzed structure does in fact get passed on to the semantic representation, but when it comes

time to compare the original incorrect interpretation with the new one, they are merged, since the original (incorrect) proposition “the man hunted the deer” entails the (revised, correct) proposition “the man hunted something.” Crucially, however, this explanation fails for the RAT verbs, as their syntactic properties block this entailment relation: “Anna dressed the baby” does not entail “Anna dressed herself.”

The final alternative explanation that we will consider is that some garden-path sentences were completely reanalyzed, but for others, no reanalysis occurred. The significant proportion of incorrect YES responses in the garden-path conditions could then be attributed to guessing on the trials on which reanalysis did not take place. Fortunately, we have clear evidence against this possibility. Recall that in one of our experiments, participants responded either to a question about the subordinate clause or a question about the matrix clause. Performance in the latter condition was close to perfect and much lower in the former condition, suggesting just what we have argued for: partial reanalysis. Some degree of reanalysis must have taken place for the NP initially parsed as the object of the subordinate clause to be recruited as the subject of the matrix clause.

Given the inability of nonsyntactic factors to account for the data we obtained, we locate the source of the misinterpretation in the initial parse and the subsequent failure to fully reanalyze the incorrect structure in the syntactic representation. Within this account of the data, we must choose again between two alternative explanations. The first is the most radical, namely, that reanalysis terminates with and the parser settles on a structure that is not licensed by the principles of human syntax. In effect, two incompatible structures are merged into one—a sort of tree splicing—in which one NP serves as both object/theme of the subordinate clause and subject/agent of the matrix clause, violating, at the very least, Case Theory and the Theta Criterion, two independently motivated modules of Chomsky’s theory of human language syntax (Chomsky 1986). If this could be shown to be the case, it would be a remarkable finding, because it would undermine the most basic assumptions about the nature of the human sentence processor—that it obeys the fundamental principles of the grammar when constructing interpretations.

We do not believe that it is necessary to adopt such a radical explanation. Instead, we favor an explanation in which the NP that is necessary for the main clause is stolen from the subordinate clause, yet is not erased from its original position. As a result, the syntactic representation ends up with the same NP in two locations and is thus no longer consistent with the input string. This explanation is consistent with a model of parsing and reanalysis proposed by Fodor and Inoué (1998). We turn next to a description of how our results provide support for their approach.

IMPLICATIONS FOR REANALYSIS: FODOR AND INOUE (1998)

Fodor and Inoue proposed the principles of Attach Anyway and Adjust to explain how reanalysis processes operate. Attach Anyway directs the parser simply to attach an incoming element even if it does not fit into the current phrase marker, temporarily ignoring the ungrammaticality that results. Adjust resolves the grammatical conflict minimally, so that the adjustment may result in yet another conflict elsewhere in the tree. Adjust then moves on to correct that conflict, and so on. Adjust is normally constrained by the Grammatical Dependency Principle (GDP), which states that Adjust operations apply between nodes that are in some sort of a grammatically defined relationship: a head and its argument, for instance, or a verb and its subject. Reanalysis of the sentences used in these experiments does not involve elements in such a relationship: The verb in the matrix clause and the NP that the parser would like to steal from the subordinate clause in order to create a grammatical sequence are not grammatically dependent on one another. In these circumstances, Fodor and Inoue propose that the parser engages in a process they term “theft,” whereby the lexical string *the deer* (in our example sentence) is stolen from the subordinate clause and assigned a position in the matrix clause, in violation of the GDP. In the case of the experimental sentences used here, theft could result in the dual thematic role assignment of Theme and Agent to *deer* that we observed with our participants, because when the GDP is violated, the backward repair operations (Adjust) that are normally set in motion by Attach Anyway do not take place. As a result, the parser would not end up being able to return to the subordinate verb, reaccess the lexicon, and locate its intransitive argument structure. The subordinate verb would, therefore, remain transitive, yielding a structure that could support the dual Theme/Agent (mis)interpretation.

Our data suggest that for the plausible sentences in our first experiment (e.g., *While the man hunted the deer ran into the woods*), the overall interpretation for the entire sentence is checked according to the Minimal Revisions Principle (MRP) (Frazier & Clifton, 1998), or something like it (see Fodor & Inoue’s “Minimal Everything,” 1998). The MRP states that any revisions deemed necessary and consistent with the error signal should maintain as much of the initially assigned structure *and interpretation* as possible. As Frazier and Clifton suggest, it seems clear that the only way that such interpretation-preserving revision could be accomplished is for there to be a successive cyclic interpretation check after each new piece of syntactic structure is built. It seems unreasonable to claim that no interpretation is considered until the input string ends, so such a cyclic evaluation of interpretations relative to syntactic structure must occur. This process of evaluation would, it seems, result in extremely fast reanalysis in cases where the

NP is semantically disfavored by the verb, which would be consistent with, for example, Clifton (1993) and Rayner, Garrod, and Perfetti (1992), because it is plausible that a deer is being hunted and also running into woods and because the occurrence of theft means that Adjust revisions do not take place, reanalysis may terminate before an entirely legal syntactic tree is established. The MRP predicts that the proposition stating that the man hunted the deer would not necessarily be erased, because that interpretation is plausible and minimal revision states that interpretations should be maintained to the extent possible. In contrast, for the implausible cases (the sentences in which the man is hunting and the deer is pacing in a zoo), the process of semantic evaluation would cause the comprehender to detect that something is wrong with the global interpretation of the entire sentence. The outcome of the evaluation, then, could trigger the reanalysis routines to continue until a correct structure (and grammatical equilibrium) is achieved.

Fodor and Inoue's (1998) model also suggests a possible explanation for the results observed with garden-path sentences containing RAT verbs. Fodor and Inoue argue that the parser resists "demot[ing] a syntactic argument role [that has already been successfully assigned] from overt (lexically realized) to implicit (unrealized)" (p. 118). What appears to be crucial is the parser's preference to assign thematic roles rather than to leave them unassigned and unrealized. If a verb is optionally transitive, the principle of Late Closure (Frazier & Fodor, 1978) causes the parser to assume the maximal number of argument positions, so a postverbal NP is made into the verb's Theme. In garden-path structures, such as the ones we examined in these experiments, the parser receives an error signal when the main clause verb is encountered, which then triggers theft of the NP that had been made into a Theme of the preceding, subordinate clause. The Theme role that the verb could have assigned must then be left unrealized. If the verb in the subordinate clause is a RAT verb (as in *While Anna dressed the baby spit up on the bed*), reanalysis differs in some important ways. In a successful reanalysis, the NP *the baby* must be changed from Theme of *dressed* to Agent of *spit up*. Subsequently, the Theme role is reassigned by the RAT verb to an empty position in its argument structure, resulting in the only syntactically licensed interpretation: *While Anna dressed herself the baby spit up on the bed*. The parser must reaccess the lexical entry for *dressed* before taking this step of postulating an empty category and assigning a thematic role to it, in order to know that the verb licenses such a category and allows this particular co-indexing. Our finding that garden-path sentences containing RAT verbs in the main clause were easier to reanalyze than those containing regular optionally transitive verbs suggests that the parser is happier when it can discharge a verb's thematic role to a constituent (even a null one) rather than adopting an interpretation which does not allow the role to

be discharged to any constituent. That is, the correct analysis of the sentence *While Anna dressed the baby spit up on the bed* had the parser assigning the role of Theme to an empty category after *dressed*, resulting in a specific theme for the verb *dressing* (Anna herself). This situation contrasts with the state of affairs for a sentence such as *While Bill hunted the deer ran into the woods*, because *hunt's* Theme role cannot be assigned to any syntactic constituent at all. Therefore, the interpretation that must be created is generic: Bill hunted something or other. Consistent with Pritchett (1992), then, as well as Fodor and Inoue, the somewhat better performance with garden-path sentences containing RAT verbs suggests that the parser prefers to discharge a verb's thematic roles rather than to leave them unrealized.

Finally, our study supports Fodor and Inoue (1998)'s suggestion that Adjust does not always yield a "stable state" in which "grammatical equilibrium" is achieved (p. 106). Our data show that Adjust may indeed terminate before a stable, grammatical state is reached. Specifically, it seems that once the critical NP is analyzed as object/Theme of the subordinate clause and then reanalyzed as subject/Agent of the matrix clause, reanalysis operations may terminate. A novel implication of our results is that comprehenders are not disturbed to end up in this state of disequilibrium. Their level of confidence, for instance, suggests that they are relatively insensitive to any violations that might result from the early termination of Adjust operations.

MORE GENERAL IMPLICATIONS: IS LANGUAGE COMPREHENSION JUST "GOOD ENOUGH"?

The results from the experiments we have described here have implications that go beyond models of sentence processing and reanalysis. We believe that the results call into question some of the fundamental assumptions that our field has made about the nature of comprehension. Theories in psycholinguistics (and in cognitive science more generally) have assumed that input systems are designed to create full, complete, and detailed representations for the stimuli with which they are presented. In the field of language, we have taken it as given that when people receive a sentence made up of a set of words in a particular syntactic configuration, successful comprehension requires them to recover those words and the exact structure. Yet, the experiments reported here suggest that people are often satisfied with inaccurate representations based on incomplete processing of the sentence. These findings join a body of work (much of it discussed outside of psycholinguistics) showing that comprehenders can be remarkably insensitive to discrepancies between the interpretation they obtain and the one that

is appropriate given the content of a sentence (see Ferreira & Henderson, 1999). Consider the Moses illusion, for instance (Erickson & Mattson, 1981): If people read the question, *How many of each kind of animal did Moses take on the ark*, they typically respond “two” rather than objecting to the presupposition behind the question (Noah, not Moses, saved creatures from the great flood). Barton and Sanford (1993) also showed that people overlook the anomaly in a sentence such as *The authorities were trying to decide where to bury the survivors*. More recently, Ferreira and Stacey (submitted) demonstrated that people systematically misinterpret simple passive sentences such as *the dog was bitten by the man*. Their error is to flip the thematic roles so as to create a more plausible sentence.

Taken together, these studies present a challenge for the fundamental assumption in psycholinguistics that comprehension is based on the creation of full, accurate, and detailed representations. It appears, instead, that people work on sentences until they reach a point where it subjectively makes sense to them and then processing may cease. The criterion that an individual sets may vary depending on the particular circumstances in which the linguistic communication is taking place. If a casual conversation is happening in a noisy bar or restaurant, then the criterion will likely be set quite low; at the other extreme, if the comprehender is a participant in an experiment in which she knows that her ability to read and understand sentences is being measured, then the criterion will be set much higher. An implication of this view is that the participants in our experiments were likely setting quite a stringent criterion for what they would consider adequate comprehension and yet they still failed to understand these garden-path sentences completely. This observation suggests that these garden-path sentences somehow produce an illusion of comprehension in our participants. These sentences, then, might form a class with the items used in Moses illusion studies and perhaps even the stimuli that elicit visual illusions. They will be misinterpreted despite the best attempts of the comprehender to come up with a correct analysis.

FUTURE DIRECTIONS

The research we have described in this article offers an exciting and novel way to think about sentence comprehension. Research could proceed from here in a variety of different directions. We will highlight just a few themes that we will be exploring in future studies.

First, the misinterpretation effect we have discovered should be evident for other structures involving temporary syntactic ambiguities. Work is currently underway to test whether in a sentence such as *Mary bumped into the*

busboy and the waiter told her to be careful people believe that Mary bumped into the waiter. Or, consider the famous reduced relative ambiguity: If people read *the soldiers warned about the land mines changed their locations*, will they say YES if asked whether the soldier warned someone about some land mines (i.e., more than for the same sentence with disambiguating words that eliminate the ambiguity)? Testing this possibility also allows us to examine a slightly less related question, but one that is important for models of sentence comprehension and reanalysis: Does the head-position effect identified by Ferreira and Henderson (1991) hold for other structures besides the subordinate/main clause ambiguity? This question can be answered as we study misinterpretations, because one of the tools we use to show that the effect is due to syntactic misanalysis is the head-position effect. Therefore, to establish that (for example) participants think that soldiers warned others to avoid land mines because of a syntactic misanalysis and not because of their confidence in the good character of military personnel, we need to show that the belief is held more strongly the more severe the misanalysis.

Second, we intend to examine more carefully people's memories for these sentences. Consider a simple recall task: If people are asked to recall *While Anna dressed the baby spit up on the bed*, will they recall the sentence in a way that suggests their memory representation for it is distorted—for example, the representation contains two occurrences of *the baby*, once as object of *dressed* and once as subject of *spit up*? Another strategy is to ask people to answer the question *Did Anna dress the baby*, as we did before, but then to follow up that question with one that might expose to the comprehender the problem with his interpretation: What if people were asked, *Did Anna dress herself*? This investigation would reveal whether the misinterpretation effect is based on participants making a copy of the ambiguous NP so that it can serve in both syntactic positions.

Finally, we are intrigued by the possibility that successful comprehension of these garden-path sentences (which, after all, was achieved on a fair proportion of the experimental trials) requires inhibition of the inappropriate proposition. That is, when people properly understand these garden-path sentences, they have inhibited the proposition that (for instance) Anna dressed the baby. This hypothesis suggests an important role for inhibition processes in sentence comprehension, a topic about which little is known. One interesting question to explore is whether, having inhibited a given proposition on trial n , they would then have more difficulty creating it on trial $n + 1$. This possibility would follow from standard views of inhibition, which assume that a representation that has been inhibited is more difficult to retrieve or use later.

CONCLUSIONS

The work we have described here and in Christianson *et al.* (in press) has yielded an important empirical result: Comprehenders often fail to understand certain garden-path structures. This finding is particularly striking given that the subordinate-main clause ambiguity is not among the most difficult ones to resolve; the reduced relative structure, for instance, is generally thought to be more challenging. We have argued that this discovery has important implications for models of sentence comprehension and reanalysis. We also believe that the research is of interest to those interested in architectural issues in cognitive science more generally, because it seems to suggest that comprehension may be based on superficial and somewhat distorted representations. Clearly, the field of sentence comprehension would benefit if more work were conducted that focuses directly on trying to uncover people's interpretations of language. As we have seen, there is much to be learned from studies designed to reveal the content of people's representations for sentences.

REFERENCES

- Barton, S. B., & Sanford, A. J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory and Cognition*, *21*, 477–487.
- Chomsky, N. (1986). *Knowledge of Language*. Dordrecht: Foris.
- Christianson, K., Hollingworth A., Halliwell, J., & Ferreira, F. (in press). Thematic roles assigned along the garden path linger. *Cognitive Psychology*.
- Clifton, C., Jr. (1993). Thematic roles in sentence parsing. *Canadian Journal of Experimental Psychology*, *47*, 222–246.
- Erickson, T. D., & Mattson, M. E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behavior*, *20*, 540–551.
- Ferreira, F., & Stacey, J. (submitted). The misinterpretation of passive sentences.
- Ferreira, F., & Henderson, J. M. (1991). Recovery from misanalyses of garden-path sentences. *Journal of Memory and Language*, *25*, 725–745.
- Ferreira, F., & Henderson, J. M. (1998). Syntactic reanalysis, thematic processing, and sentence comprehension. In J. D. Fodor & F. Ferreira (Eds.), *Reanalysis in sentence processing* (pp. 73–100). Dordrecht: Kluwer.
- Ferreira, F., & Henderson, J. M. (1999). Good enough representations in visual cognition and language. Paper presented at *Architectures and Mechanisms of Language Processing Conference*, Edinburgh, Scotland.
- Fodor, J. D., & Inoue, A. (1998). Attach anyway. In J. D. Fodor & F. Ferreira (Eds.), *Reanalysis in sentence processing* (pp. 101–141). Dordrecht: Kluwer.
- Frazier, L., & Clifton, C., Jr. (1998). Sentence reanalysis and visibility. In J. D. Fodor & F. Ferreira (Eds.), *Reanalysis in sentence processing* (pp. 143–176). Dordrecht: Kluwer.
- Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, *6*, 1–34.
- Grice, P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics, Vol. III: Speech acts* (pp. 41–58). New York: Seminar Press.

- Pritchett, B. L. (1992). *Grammatical competence and parsing performance*. Chicago: University of Chicago Press.
- Rayner, K., Garrod, S., & Perfetti, C. A. (1992). Discourse influences during parsing are delayed. *Cognition*, *45*, 109–139.
- Warner, J., & Glass, A. L. (1987). Context and distance-to-disambiguation effects in ambiguity resolution: Evidence from grammaticality judgments of garden-path sentences. *Journal of Memory and Language*, *26*, 714–738.