INFERENCE GENERATION DURING AUDITORY LANGUAGE COMPREHENSION

David A. Swinney
Lee Osterhout

I. Introduction

It has become nearly axiomatic (and even likely true) to hold that inference generation during language comprehension is a function of two things: (1) the nature of the inference under consideration and (2) the conditions under which the inference is examined. Thus, in the first case, we find the literature rife with observations that while there is little evidence of "automatic" generation of certain "forward-looking" or "elaborative" inferences (e.g., Alba & Hasher, 1983; McKoon & Ratcliff, 1981; Singer & Ferreira, 1983), inferences involving certain likely causal consequences may be "weakly generated" (e.g., Keenan, Baillet, & Brown, 1984; McKoon & Ratcliff,1986). Further, it has been observed that certain "backward-looking" inferences such as bridging inferences or the inference linking an overt anaphor with its antecedent may be reliably and automatically generated (Clark & Sengul, 1979; Corbett & Chang, 1983; McKoon & Ratcliff, 1981, 1986). Indeed, the literature on text processing details an amazingly variable array of different elaborative functions for different "types" of inferences.

However, it has also been amply demonstrated that inference elaboration appears to be a function of the conditions under which these inferences were examined. Roughly speaking, this latter issue has three aspects: the beliefs the listener or reader generates about the task to be

THE PSYCHOLOGY OF LEARNING
AND MOTIVATION, VOL. 23

performed in the experiment, the mode of perceptual presentation (roughly, the encoding conditions), and the mode of response (recall, recognition, automatic priming, etc).

A few of these aspects deserve some brief discussion. The basic issue to be determined in the study of inference elaboration in natural language is not one of whether or not some inference can be elaborated. Indeed, if one can imagine an inference, then it is possible to generate it. Rather, the basic issues are *when* (in the temporal course of comprehension) that inference is elaborated, and *what* the processing results of that elaboration are. Regarding the latter, one might be interested in the memorial consequences of a generated inferences, or in the immediacy of availability of that inference for use in problem solving, or some other such processing result. Thus, it is important to consider these parameters in choosing the response task. Probably the most obvious and important distinction to be made in this regard is between tasks aimed at elucidating the perceptual processing of inferences and those involved in demonstrating memorial representation of inferences. The former situation calls for on-line, temporally immediate tasks whereas the latter requires tests of memory at appropriate temporal distances from the perceptual elaboration of the inference.

Similarly, just as issues concerning response mode are critical in the study of inference processing, so is the issue of presentation (stimulus) mode. Each presentation situation imposes its own strategies and limitations on the examination of the inference process. As simply a single example of the issue, a great deal of the current literature on inferences involves the self-paced visual presentation of words (reading). This task obviously has important contributions to make to the literature. However, self-paced reading is a task that invites a number of special strategies on the part of the subject. For example, because there is pressure in this task to keep a rhythm going in button pressing, subjects in this task are often pressing the "next-presentation" button before they have actually finished reading a particular word or sentence. Similarly, this task allows for the subject to pause in order to "mull" things over. While this behavior is, of course, part of what can occur in reading, it is a process that mingles problem-solving behavior with perceptual language processing. Certainly, we know that one can elaborate an inference through problem-solving techniques (by querying oneself about what something might mean, for example). However, it is obviously important to consider whether such cognitive, problem-solving processes are those one wishes to examine, or whether something more focused on early "perceptual" processing is the target of enquiry in one's study of inference generation.[1]

[1]Here, the term *perceptual* is used in the sense taken by Garrett (1979) and by Fodor (1983). The term is not synonymous with traditional perceptual psychological usage which implies that perceptual analysis is limited to the results of early, peripheral processing. Nor

Finally, it is certainly a defensible position to hold that *reading* is a process that is derivatively based upon the processing underlying normal *listening* (see, e.g., Liberman, Shankweiler, Camp, Blackman, & Werfelman, 1980). Thus, if one is interested in the early stages of inference processing (stages prior to cognitive problem solving, etc.), it will be worthwhile to examine inference processing during auditory comprehension as well as during reading (which is the modality that essentially all of the past and current literature has used in examining inference).

II. Perceptual and Cognitive Inferences: Toward a Typology of Inference Generation

As argued previously, there are a wide variety of considerations and issues to be dealt with in attempting to understand when and how we make inferential elaborations during language comprehension. In what follows, we will present a model aimed at elucidating some preliminary classes of inference generation, classes that we feel differentiate the above issues in important ways. Following that, we will provide an array of empirical support for this classification.

The model identifies closely with Fodor's (1983) theory of mental modularity. Based on this theory, we propose to make a first-cut differentiation of inferences into two categories: perceptual and cognitive. Such a dis-

---

does it refer to processing that is noninferential, in the broad sense. Rather, it refers to initial processing of a stimulus array (auditory or visual) that is both peripheral and inferential and distinguished from other levels of processing in terms of the depth of cognitive abilities that can be brought to bear on the process. In terms of sentence processing, for example, perceptual processing is that early, superficial (yet inferential, in part) analysis of the sentence that derives from lexical access, structural analysis, and combination of these into simple logical form. The perceptual result of sentence processing is, thus, a superficial analysis, not a fully elaborated one. See footnote 2.

[2] Use of the terms *perceptual* and *cognitive* inference is, undoubtedly, going to generate annoyance on the part of some readers. It is not done with that intent, however. The terms have held different meanings throughout the history of psychology. Particularly, in the last decade or so, much effort has gone into demonstrating that processes traditionally thought of as "perceptual" have cognitive components. Our use here is not intended to refute that position; indeed, much of what was classified as perceptual is influenced by cognitive factors (whatever is meant by the two terms). However, we use the terms here in keeping with the definitions that Fodor (1983) has introduced in distinguishing modular and nonmodular processes. In that light, the terms do reflect definitionally and conceptually distinct processes. We take that distinction to be a potentially profound one for psychological examination, and hence we use the term *perceptual* instead of terms such as *modular* or *automatic* or *early*, and the term *cognitive* instead of *nonmodular* or *controlled* or *late*, which might be seen as alternative definitions. We feel these terms best capture the behavior we are attempting to describe; namely, a first-pass, superficial, automatic, sentential analysis that is equivalent to perceptual processing and a later, more "cognitive," problem-solving type analysis that follows the earlier superficial one.

tinction allows for a differentiation of those inferences which are predicted to be *automatically, mandatorily,* and *immediately* derived during the on-line perceptual processing of language (perceptual inferences) and those inferences which are never automatically or immediately derived and which are under cognitive control (nonperceptual or cognitive inferences). This distinction, then, is essentially that encompassed by Fodor's (1983) distinction between encapsulated and nonencapsulated processes. It is argued here that the class of perceptual inferences is necessarily a part of the early perceptual, encapsulated processing of language. Such inferences will have the qualities of (1) being independent of world knowledge, intentions, and pragmatics and (2) being drawn mandatorily and immediately. By *immediately* we mean that the inference will be drawn as soon as the licensing conditions for that inference are met (i.e., occurrence of the perceptual stimuli which make the inference possible). Thus, even knowledge that tells the processor that the inference is likely to be inappropriate in a certain discourse condition is hypothesized to be unable to change the immediate and mandatory elaboration of these perceptual inferences. Therefore, perceptual inference elaboration is argued to occur *immediately during* comprehension, and not merely at the point of the test of comprehension or at recall from memory.

Cognitive inferences, on the other hand, are those inferences which are neither mandatory, automatic, or immediate. In fact, by definition, cognitive inferences cannot be generated or elaborated immediately upon occurrence of potential licensing conditions. That is, even though information capable of supporting a particular cognitive inference may appear in the initial superficial perceptual analysis of a sentence, cognitive inferences are not automatically elaborated. Rather, they are under the "cognitive" control of strategies tied to world knowledge, statistical bias, and pragmatic knowledge. Further, in the sense of Posner and Snyder (1975), cognitive inferences are controlled as opposed to automatic. Thus, criterially, cognitive inferences cannot be elaborated immediately. They are taken to be entities beyond initial, first-order sentence analysis.

The distinctions we are proposing between perceptual and cognitive inferences are distinguished theoretically by virtue of their amenability to influence by world knowledge (or, alternatively, their information encapsulation) and the automaticity with which the inference may be generated during ongoing comprehension. Empirically, the two classes of inference differentiate on the basis of the temporal characteristics of their elaboration. Perceptual inferences are those that will be elaborated as soon as the licensing conditions are encountered by the listener or reader, regardless of whether certain pragmatic conditions counterindicate elaboration of these inferences (see following sections for examples). Cognitive inferences are those that will be shown not to be capable of being immediately

elaborated upon occurrence of licensing conditions, regardless of how favorable the pragmatic conditions (world knowledge, likelihood, etc.) are for the immediate elaboration of these inferences.

We propose that the class of inferences we have labeled as perceptual constitutes a limited set. Specifically, we hypothesize that perceptual inferences constitute a subset of those inferences involved in establishing coreference during language comprehension: explicit and implicit anaphoric elements. Future research and theory may eventually support the existence of other members of this class, but at this point we know of no others. Is the following section we will present empirical evidence that shows these, but not other, inferences to be mandatory, immediate, and uninfluenced by world knowledge.

Cognitive inferences, on the other hand, involve the majority of those inference types which have been studied in the literature under the labels predictive, causal, schema-related, metaphorical, and so on. By and large, the reason that researchers have found these inferences to be unreliably generated in various studies is, we argue, because these inferences are not mandatory, automatic, or immediate, but rather are under cognitive control, are constrained by world knowledge, pragmatics, and estimates of probability, and thus cannot be immediately, or universally, generated. Again, in what follows, we provide several pieces of empirical evidence examining these inferential processes.

III. Empirical Evidence for the Distinction between Perceptual and Cognitive Inferences

A. PERCEPTUAL INFERENCES

In what follows, we sketch a number of experiments demonstrating the properties of perceptual inferences, providing more detailed evidence for a few of the more surprising examples of this inference type.

First, as was suggested, an obvious candidate for a perceptual inference is the process involved in establishing coreference for explicit pronominal anaphors. A large number of investigators have looked at antecedent assignment for pronouns and other anaphoric elements (e.g., Bever & McElree, 1988; Corbett & Chang, 1983; Dell, McKoon, & Ratcliff, 1983). In these studies, subjects were visually presented material containing pronouns (or other anaphoric elements) with coreferents in an antecedent position (earlier in the sentence). The findings of all studies indicated that some antecedents (either one or several) are activated *relatively* quickly upon the reader seeing the pronoun. In both the Corbett and Chang (1983) and Bever and McElree (1988) studies, subjects appeared to have acti-

vated an antecedent to the pronoun by the end of the sentence containing the pronoun. Dell *et al. (1983)* showed that 250 msec after reading an anaphor (not a pronoun in this case) in a paced reading condition, some sort of activation for the correct antecedent was evident. Thus, it appears that the inference binding the antecedent to anaphoric elements (and, in particular, pronouns) is made fairly quickly in controlled reading situations. However, while this certainly points strongly toward a claim that a coreferent is bound (assigned) to the pronoun rapidly, these studies do not demonstrate true immediately of coreference assignment. For this to be demonstrated, it must be shown that such assignment is made immediately upon occurrence of the licensing condition (the pronoun) in the sentence, not just temporally "downstream" from it. A study reported in Swinney, Ford, and Bresnan (in press) does, however, provide some finer-grained evidence about the immediacy of inference generation in the case of pronominal anaphora. We will describe it here in a fair amount of detail, as the general experimental technique will be relevant to subsequent studies in this article.

The Swinney *et al.* study involved the use of the cross-modal lexical priming technique (see Swinney, Onifer, Prather, & Hirshkowitz, 1979) to examine activation of potential coreferents for overt pronouns during sentence comprehension. Subjects listen to auditorily presented sentences containing overt pronouns with several prior possible antecedents. For example, in this study, subjects heard sentences such as:

> *The boxer visited the doctor that the swimmer at the competition had* $\overset{1}{*}$ *advised him* $\overset{2}{*}$ *to see about his injury*.

Subjects were told that one of their tasks was to listen carefully to the sentence and to understand it, and that they would be tested on their comprehension during the experiment (and they were). Subjects were told that they had a second task which they had to perform while listening to the sentences: they would see a letter string appear on a CRT in front of them and their job was to decide as quickly as they could whether or not the letter string was a word (a lexical decision task). In the case of the experimental (as opposed to control or filler) sentences, the letter strings were presented visually at the numbered test points indicated in the example sentence by asterisks during auditory presentation of the materials. The words that were visually presented were the possible coreferents of the pronoun (e.g., BOXER, DOCTOR, SWIMMER) or control words that were matched for frequency, length, and a priori lexical decision time to these targets but which were semantically unrelated to them. Thus, for the experimental sentences, identity priming for each of these possible co-

referents was examined. In all such experiments, any given subject was only tested at one display point and with one target or control word. Further, there were always sufficient filler materials (containing unrelated targets) so that experimentally related target words occurred less than 20% of the time throughout the experiment.

Thus, identity priming for lexical decision to the coreferent target (compared to matched controls) was used as a measure of activation of the antecedent at the particular point in the sentence at which the probe was presented.[3] The first target presentation point in the sentence (after the auxiliary verb) was used as a baseline, to determine the activation of each of the possible coreferents prior to the pronoun. Comparison of activation at the baseline to amount of activation (priming) at the second target presentation point (immediately after the pronoun) provides a measure of immediate activation of appropriate and inappropriate coreferents by the pronoun. Thus, this design provides a more direct test of the immediacy of inference generation for pronouns than some of the prior studies have been able to, as it tests at the point where such inference is first licensed-that is, immediately after the pronoun. The results for this study can be seen in Table I. The appropriate referent to the pronoun (BOXER) was the only item that demonstrated significant priming compared to its control at test point 2 (after the pronoun) and a significant increase in priming between test point 1 (baseline) and test point 2. Thus, this study provides direct evidence that the pronoun immediately activates its antecedent during comprehension. It also turns out that this evidence suggests that only the correct antecedent is activated and that this process may be guided by linguistic knowledge. That is, it was not simply the most recent possible antecedent or even all possible antecedents that were activated. Instead, only the one that was syntactically free to be an antecedent of the pronoun was activated. However, this is a complicated argument and one not directly relevant to the issue at hand, In all, overt pronouns act in a manner in keeping with the definition of a perceptual inference; they immediately activated their coreferents via inference generation.

[3] It is worth noting briefly that although there has been some interesting discussion of the effects of backward priming in the use of cross-modal lexical decision tasks (see Glucksberg, Kreuz, & Rho, 1987) this argument cannot apply to the case here. In the first place, research by Burgess, Tanenhaus, and Seidenberg (1989) and by Prather and Swinney (1988) has demonstrated that the backward priming effect, as argued by Glucksberg *et at.*, does not hold up in accounting for cross-modal priming work. Further, in this particular case, backward priming, even if it could have accounted for prior work, cannot account for results under this current paradigm, because there is nothing to backward prime from in these cases. That is, the pronoun is not a source of information as to what to activate. Quite to the contrary, the pronoun is itself seeking information about what to activate.

TABLE I

INFERENCE PROCESSING: OVERT PRONOUNS[a]

Visual targets	Test point 1	Test point 2
Boxer	23	51*
Doctor	32	42*
Swimmer	9	20

[a]Lexical decision results are given as priming (control minus experimental) reaction times in msec to items presented at different sentential test points.

* $p < .05$.

As suggested earlier, this result may not seem very surprising given the prior literature. It simply validates the claim for immediacy in this type of inference generation. However, in the studies we present immediately below, we provide evidence that supports a more surprising type of anaphoric relation also to be a case of perceptual inference generation: those involving the implicit (nonovert) anaphoric conditions found in filler-gap constructions.

Linguistic "gaps" are phonetically empty linguistic constructs created by movement (in the cases we will examine here) of object noun phrases in relative clauses. Thus, for example, in the sentence

That is the cat that the dog chased (t).

the (*t*) marks a type of gap that is called a *trace.* This trace is phonologically empty (that is, it is not phonologically realized) and it indicates that the direct object of the verb *chased* was moved from this *t* position forward to the matrix (main) clause in the sentence. (The underlying relative clause is taken to be, in its "deep" form, *The dog chased the cat;* thus, when *cat* was moved to the matrix clause, it left behind a trace.) This is simply one example of a sentence with a linguistic gap. Such gap constructions are interesting because, by linguistic standards, they are cases of implicit anaphoric relationships. Any processing account of language comprehension must account for how the antecedent of this implicit anaphor (in this example, *cat*) is linked to the trace. This gap-filling process is, thus, a case that is similar in some respects to the case of overt pronouns. However, here the marker of the anaphoric relationship *t* (what we have called the licensing condition for the inference) is simply not overtly produced for the

comprehender to see or hear. However, it is a function of the verb preceding it. That is, the verb (in our example, *chased*) requires a direct object. It turns out that the direct object has occurred early in the sentence (and is so marked by the relative pronoun *that*). Thus, the question relevant to inference generation is: Will the verb license an inferential process in which the antecedent to the implicit anaphoric element (the trace or gap) is immediately determined? Or, put another way, will any verb requiring a direct object (transitive verbs) license perceptual inferential processing to determine that direct object?

Two experiments by Swinney, Ford, Frauenfelder, and Bresnan reported in Swinney *et al.* (in press) directly addressed that question. In both, the results support the argument that appropriate perceptual inferences are made immediately upon occurrence of an inference licensing condition for a gap in a sentence. In these studies, subjects heard sentences such as:

> *The policeman saw the boy that the crowd at the party* ^{*1} *accused (t)* ^{*2} *of the crime.*

Lexical decisions were made to words related to either *boy* (e.g., GIRL) or *crowd* (e.g. GROUP) or appropriate, semantically unrelated control words. Priming for these items was examined at test point I (baseline, before the verb) and test point 2 (at the gap). The results are provided in Table II. As can be seen, although *group* (which had just been heard in the sentence) was significantly primed before the verb, all priming disappeared after the verb had been heard and the gap was reached. This, suggests that *crowd* was not accessed or activated as the antecedent of the trace in this sentence. However, even though a target word related to the correct antecedent of the gap (*boy*) was not significantly primed at test point 1 before the verb, it was significantly primed at the gap. This evidence strongly supports the notion that a mandatory, immediate, perceptual inference has been made at this point as to the appropriate antecedent of the trace (*boy*).

The same materials were used in a second study using a naming technique rather than the lexical-decision technique in a cross-modal priming paradigm.[4] The experiment was essentially identical to the previous one, except that a naming task was used. The words POLICEMAN, BOY, or CROWD appeared at the two test points and the time to say these words at each test point was measured. Priming was measured by subtracting reac-

[4] The reasons for this address claims in the literature concerning the relative sensitivity of lexical decision and naming paradigms to various processing properties.

TABLE II

INFERENTIAL PROCESSING: COREFERENCE IN FILLER-GAP CONDITIONS

	Experiment 1[a]		
Target related to	1	2	Experiment 2[b]
Policeman	—	—	6
Boy	12	27*	-29*
Crowd	44*	8	-5

[a]Lexical decision values are priming (control minus experimental) reaction times in msec to semantically related words at sentential test points. I and 2 refer to test points.

[b]Naming values are the difference in msec between reaction time to naming the following words at sentential test points.

*$p < .01$.

Lion time to name a target word at test point I from that to name it at test point 2 (averaged across subjects, of course, since no one subject saw the same word at more than one point). These priming scores are displayed in the right side of Table II. As can be seen, these results replicate those of the first study. Priming was immediately evident for (only) the appropriate antecedent of the trace, again supporting the conclusion that this gap-filling process for these implicit anaphoric elements is a case of perceptual inference processing.

These studies provide evidence of the immediate nature of these inferences. Additionally, they point up the mandatory nature of these inferences, given that there is nothing overt forcing the inference to occur in the case of the gap-filling studies. However, there is another property of perceptual inference processing that we have proposed-independence from effects of plausibility, pragmatics, or world knowledge. Although the range of studies necessary to fully examine this issue is enormous, we have begun a number of experiments aimed at precisely this issue. We will present just the first in this line of research. In this study, the likelihood of a particular agent (actor) performing an act was set in strong opposition to the appropriate antecedent for a trace in a gap-filling situation. Thus, we were looking to see if pragmatic information would cause the gap to be filled with the pragmatically appropriate but syntactically inappropriate

antecedent, or whether the gap-filling procedure truly takes place independently of world knowledge. For example, subjects heard sentences such as:

> *Everyone watched the enormous heavyweight boxer that the small 12-year old boy on the corner had* $\overset{1}{*}$ *beaten (t)* $\overset{2}{*}$ *so brutally.*

Here, the syntactically appropriate antecedent of the gap is *boxer* (i.e., it is the boxer who is being beaten), although pragmatics suggests that this is an unlikely event. Naming scores for BOXER and BOY were collected at both the point immediately before the critical main verb (e.g., after the auxiliary verb *had* in the above example) and the point after the critical main verb. The preliminary results are uniform and unequivocal in their interpretation. Significant priming is always found for naming times to BOXER after the verb, but those for BOY show nonsignificant effects. Thus, only *boxer* is inferentially generated as the appropriate antecedent to the trace. In short, at least in this case, the likelihood of who would be beaten in a contest between a 12-year-old boy and a boxer is ignored by the inferential process which is determining the antecedent to the trace. While it is certainly true that this preliminary experiment does not prove that all perceptual inferences are immune to all world-knowledge or plausibility effects, it does strongly suggest that they may be independent in precisely the manner argued previously.

B. COGNITIVE INFERENCES

However interesting or compelling one finds the case for perceptual inferences to be, in order to make the case for the perceptual/cognitive inference distinction it is necessary to demonstrate evidence that some inferences have the characteristics claimed to be associated with cognitive inferences. The major such characteristic is that they cannot be immediately derived.

Even though the discourse-inference literature is filled with claims of inference elaboration for an enormous array of inference types, there have been very few claims of elaboration of these inferences prior to experimental query about these inferences (see, e.g., Alba & Hasher, 1983, McKoon & Ratcliff, 1986, for reviews of this literature) and even fewer experimental results that could be taken as support for immediate inference generation. Indeed, even in the cases for which researchers have suggested that automatic inference generation is most likely (e.g., those involving schema, forward inferences, and causal inferences) there is no evidence

that unequivocally establishes that they are elaborated immediately in the sense of the perceptual inferences discussed previously. Part of the problem is methodological. Much of the literature has not utilized on-line tasks which can determine immediate activation. Further, those studies employing tasks that do have the ability to detect immediate activation (e.g., McKoon & Ratcliff's deadline procedure) seem to have failed to demonstrate immediate inference generation (that is, generation of an inference immediately upon occurrence of licensing conditions for that inference). For example, McKoon and Ratcliff (1986) have argued that even in cases involving extremely high likelihood (highly predictable outcomes) in causal events, they find that the inferences are, at best, only weakly encoded into memory.

In what follows, we briefly present a number of studies involving a range of inferences that have been thought to be good candidates for immediate, automatic activation of inferences. In each case we have used sensitive on-line experimental tasks that are capable of detecting activation of the inference and we have made conditions as favorable as possible for occurrence of the inference. In each case, we find that the inference is *not* activated immediately upon occurrence of the licensing conditions, but is only activated a short time later (if at all), depending on the nature of the context and processing conditions. These inferential types include those involved in metaphor processing, presuppositional processing, and schema-based instrument activation.

In the first, the processing of metaphors was examined with a cross-modal priming task to determine, when activation of literal and metaphoric meanings takes place during normal sentence processing (Swinney, 1981). In this experiment, subjects heard definitional metaphors (e.g., *The sky was a fire.)* in a sentential context. Thus, for example, subjects heard the sentence:

Jeff told his friend that Mary's hair was $\overset{1}{*}$ *honey* $\overset{2}{*}$ *and her* $\overset{3}{*}$ *smile gleamed.*

At each of the three test points, subjects were primed for words semantically related to the literal or the metaphoric meaning of the metaphor *(hair was honey).* In this case the words BEE and BLOND were used, along with appropriate control words. The basic issue is one of whether the metaphoric interpretation (the inference) was activated immediately upon encountering the licensing condition for the inference (the end of the metaphor). Table III presents the data from this experiment. As shown in the table, there is no priming (inferential activation) of the *metaphorical* interpretation until nearly 500 msec (the third test point) after the end of

TABLE III

INFERENTIAL PROCESSING:
DEFINITIONAL METAPHORS[a]

	Test point		
Target	1	2	3
Bee	-4	32*	15
Blond	26*	9	35*

[a]Lexical decision results are given as amount of priming (control minus experimental) reaction times in msec.
* $p < .05$.

the metaphor. Importantly, this pattern of results was replicated in a situation in which all materials in the experiment were metaphors. Thus, in this condition, subjects knew that they would be getting tested on metaphors and that they should make metaphor inferences in each of these sentences. However, this pragmatic situational knowledge did not change the basic results: no immediate activation of the inference took place. (It should be noted that a 500-msec delay in processing is an enormous amount of processing time and safely takes this out of the realm of being a perceptual inference.)

The second case we will describe is one involving the perceptual consequences of presuppositions. In this work (Cutler & Swinney, 1978), factive verbs were contrasted with nonfactive verbs to examine the elaboration or activation of the (inferentially) presupposed information. (Factive verbs, such as the verb *to know* in the sentence *I know how old 1 am* presuppose something to be a fact.) Thus, subjects heard sentences of the following general nature:

> *The death-penalty requirement that the young lawyer knew applied to this case was discussed thoroughly.*

In this study, the subjects' task was to detect when some word in the sentence began with the "uh" sound (as in the word *applied).* This phoneme monitoring task has been shown (Foss, 1970) to be sensitive to increases or decreases in processing load. If the occurrence of the factive verb *knew* causes activation of the presupposed information that was implied by that verb, then it should cause increased processing load and hence increased reaction time for the detection of the target phoneme.

(This condition in the example above is compared to that for the same sentence with the control verb *thought,* which is not a presuppositional verb.) The results demonstrated no immediate increase in processing load. Any effect of the factive verb (presuppositional inference) must have occurred following the inference licensing condition.

Finally we present evidence concerning the activation of instrumental inferences in schema contexts. Despite early claims in the literature'that instruments are routinely inferred during reading (e.g., Paris & Lindauer, 1976), work by Corbett and Dosher (1978) convincingly demonstrated that this earlier work was not adequate evidence of inference elaboration during comprehension, must less at point of recall. In a careful study they demonstrated that implicit (unstated) instruments are "not routinely inferred in comprehension." Similarly, work by McKoon and Ratcliff (1981) did not find routine instrumental inference generation. Their results suggested that an inference might be generated in the case of instruments highly semantically related to the verb that appeared in the sentence, but this evidence did not support a claim that they were immediately activated, in the sense of immediacy we have been talking about, even for semantically related instruments. Thus, in order to gather data that could better address the immediacy hypothesis, we used a cross-modal priming task in which we tested for the generation of inferences about instruments in the presence of a well-established schema. In this experiment, subjects heard either the schema-inducing context sentence or its control, followed by the test sentence. For example, subjects heard either (1 a) or (Ib):

(1) *a. John sat down to eat his meal that evening.* $^{1}*$
[schema-inducing context]
b. John sat down in his normal chair that evening. $^{1}*$
[control context]

Following that they heard the continuation sentence (2):

(2) He cut the juicy meat $^{2}*$ *and began eating his dinner.*

Sentences (la) [or (lb)] and (2) were presented together, as in normal discourse. Priming for the possible instrument (in this case, KNIFE, a basic instrument in eating a meal) or the control for that word (in this case the word FRAME) was examined at the two points indicated by the asterisk. The first test point was used to determine whether the instrument would be activated by the licensing condition (in this case, the schema-inducing context *eat his meal).* The second test point was used to determine whether the schema was activated later, perhaps as a result of

TABLE IV

EXAMINATION OF INSTRUMENT ACTIVATION IN STRONG SCHEMA CONTEXTS[a]

	Test point	
Context	1	2
Schema	18	43*
Control	5	17

[a]Lexical decisions (magnitude of priming in msec) to inferred instrument or control word (e.g.. *knife, frame*).

* p <.03.

cognitive inference generation in the case involving a schema-inducing context. The results are presented in Table IV. As can be seen by inspection, there is no evidence of instrument inference generation at the licensing point (test point 1). However, in the schema-generating context condition *only*, there is evidence of instrument activation considerably downstream (later in time) from that context (test point 2) Recently, we have replicated these results with a naming test.[5] In sum, there is evidence that an instrumental inference might be eventually generated in a sentence with appropriate plausibility conditions (here, the schema condition), but it is not immediately, mandatorily, or autonomously activated by the inference-licensing conditions. Thus, again, we have evidence of a contextually malleable, cognitive inference process.

IV. Summary

In summary, we have presented evidence which supports a distinction between two classes of inferences, those that are generated mandatorily, automatically, and immediately (which we have called perceptual inferences) and those which cannot be generated immediately during perceptual processing, but which are generable under the influence of plausibility, world knowledge, and pragmatics (cognitive inferences). We have

[5] It is worth noting that, in this respect, our results differ from those of Potts, Keenan, and Golding (1988), who found differences in sensitivity between lexical decision and naming tests for highly predictable events.

identified the former type of inference with early, perceptual, encapsulated, modular stimulus processing and the latter with later-occurring, cognitively driven analysis. The challenge for future empirical work on this issue is to identify and delineate further the types of inferences and the conditions of generation for those inferences that may further distinguish the classes of processes we have discussed here.

ACKNOWLEDGMENTS

The authors thank Susanne Shulman for her help in collecting the data discussed in this article. This work was supported in part by NIH grant NS21806.

REFERENCES

Alba, J. W., & Hasher, L. (1983). Is memory schematic? *Psychological Bulletin,* 93,203-231.

Bever, T. G., & McElree, B. (1988). Empty categories access their antecedents during comprehension. *Linguistic Inquiry,* **19,** 35-44.

Burgess, C., Tanenhaus, M., & Seidenberg, M. (1989). Context and lexical access: Implications of nonword inference for lexical ambiguity resolution. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 620-632.

Clark, H., & Sengul, C. (1979). In search of referents for nouns and pronouns. *Memory and Cognition, 7,* 58-64.

Corbett, A. T., & Chang, F. R. (1983). Pronoun disambiguation: Accessing potential antecedents. *Memory and Cognition, 11,* 283-294.

Corbett, A. T., & Dosher. B. A. (1978). Instrument inferences in sentence encoding. *Journal of Verbal Learning and Verbal Behavior, 17,* 479-491.

Cutler, A., & Swinney, D. (1978). The processing of presuppositions during sentence comprehension. *Tufts University Papers in Cognitive Science, 7.*

Dell, G. S., McKoon, G., & Ratcliff, R. (1983). The activation of antecedent information during the processing of anaphoric reference in reading. *Journal of Verbal Learning and Verbal Behavior, 22,* 121-132.

Fodor, J. A. (1983). *Modularity of mind.* Cambridge, MA: MIT press.

Foss, D. J. (1970). Some effects of ambiguity upon sentence comprehension. *Journal of Verbal Learning and Verbal Behavior,* 9, 699-706.

Garrett, M. (1979). Word and sentence perception. In R. Held, R. Teuber, & L. Leibowiz (Eds.), *Handbook of sensory physiology (Vol.* VIII). New York: Springer-Verlag.

Glucksberg, S., Kreuz, R. J., & Rho, S. H. (1987). Context can constrain lexical access: Implications for models of language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 323-335.

Keenan, J. M., Baillet, S. D., & Brown, P. (1984). The effects of causal cohesion on comprehension and memory. *Journal of Verbal Learning and Verbal Behavior,* 23, 115-126.

Liberman, A. L., Shankweiler, D., Camp, L., Blackman, D., & Werfelman. (1980). Steps toward literacy. In P. Levinson & L. Sloan (Eds.), *Auditory processing in language: Clinical research perspectives.* New York: Grune & Stratton.

McKoon, G., & Ratcliff, R. (1981). The comprehension processes and memory structures in instrumental inference. *Journal of Verbal Learning and Verbal Behavior, 20,* 671-682.

McKoon, G.. & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimenial Psychology: Learning, Memory, and Cognition, 12,* 82-91.

Paris, S. G., & Lindauer, B. K. (1976). The role of inference in children's comprehension and memory for sentences. *Cognitive Psychology, 8,* 217-227.

Posner, M. I., & Snyder, C. R. R. (1975). Facilitation and inhibition in the processing of signals. In P. M. A,. Rabbit & Dornic (Eds.), *Attention and performance V.* New York: Academic Press.

Potts, G. R., Keenan, J. M. (1988). Assessing the occurrence of elaborative inferencing: Lexical decision versus naming. *Journal of Memory and Language, 27,* 399-415.

Prather, P., & Swinney, D. (1988). Lexical processing and ambiguity resolution: An autonomous process in an interactive box. In S. L. Small, G. W. Cottrell, & M. H. Tanenhaus (Eds.), *Lexical ambiguity resolution in the comprehension of human language.* San Francisco: Morgan Kaufman.

Singer, M., & Ferreira, F. (1983). Inferring consequences in story comprehension. *Journal of Verbal Learning and Verbal Behavior, 22,* 437-448.

Swinney, D. (1981). *The processing of literal and non-literal language.* Paper presented at the 22nd annual meeting of the Psychonomic Society, Philadelphia.

Swinney, D., Ford, M., & Bresnan, J. (in press). On the temporal course of gap-filling and antecedent assignment during sentence comprehension. In M. Mackin & R. Kaplan (Eds.), *Language processing.* Stanford, CA: CSLI.

Swinney, D., Onifer, W., Prather, P., & Hirshkowitz, M. (1979). Semantic facilitation across sensory modalities in the processing of individual words and sentences. *Memory and Cognition, 7,* 159-165.