

Chapter 14

The Resolution of Indeterminacy During Language Comprehension: Perspectives on Modularity in Lexical, Structural and Pragmatic Process

David A. Swinney

The Graduate Center of the City University of New York
New York, New York U.S.A.

The resolution of indeterminacy is the hallmark of the language comprehension process. Indeterminacy exists at essentially every descriptive level of word, sentence, and discourse analysis. It is found in conditions ranging from overt ambiguity (such as for polysemous words and structural ambiguity) which have been the focus of work in psycholinguistics since the inception of the field, to less obvious, but equally vexatious, problems involving for such things as phonetic identity, segmentation assignment, thematic role assignment, quantifier scope, and co-reference assignment, to name but a few. The problem is so pervasive that about the only generalization that appears possible is that the language processing device seems to resolve these uncertainties.

Interestingly, models of language theory (such as linguistic grammars) are traditionally not expressed in terms that capture many of the problems of indeterminacy; uncertainties of this type are considered to be a problem to be dealt with by processing models, something viewed as performance perturbations over which generalizations are to be made, not facts which are to be captured in models of 'universal grammar'.¹ The immediate focus of this chapter, then, will be on an issue central to *processing* models: that of how the language processor resolves several types of indeterminacy during comprehension. In what follows, evidence concerning the resolution of co-reference assignment (an issue that involves structural and semantic information processing) and the resolution of lexical ambiguity will be examined. In each case, recent evidence examining the role of different types of 'context' on the resolution of these types of indeterminacy will be examined in detail.

Prior to presentation of such evidence, it is important to briefly discuss a fundamental issue that underlies the concern over indeterminacy resolution, the issue of the general nature of the cognitive architecture that supports language processing. One of the central contributions of the information processing ap-

proach in cognition to the study of language has been the development of models of processes (and representations over which they operate) that are couched in terms of the nature of the temporal interaction of different information types during comprehension and production. Out of this approach have emerged two general models which, in their extreme, are totally apposite. In one, the Interactive Model (see, e.g., Marslen-Wilson & Tyler, 1987, for a modified version), any one type of information (say, for example, world knowledge) is deemed available to any other information type (say, for example, lexical representation) at any time during processing. That is, any type of information can be used to help constrain the processing of other types of information as soon as it is useful. The other, the Modularity Model (Fodor, 1983), holds that there are principled points of interaction between information sources and *only* at those points may different information types have access to each other. This modularity model has as an important component the concept that many processes are autonomous - that is, their *internal* operations are never affected by information from other sources. Rather, it is only upon completion of the internal operations of such an autonomous 'module' that the results of that process are available for use by other processes.

In the study of language processing, examination of the effects of one putative information source upon the processing of another has become the empirical testing ground for the interactivity vs. modularity debate. An important issue related to the facts of methodology in this field follows from this: It has become axiomatic that only on-line (sometimes called real-time) methodologies are capable of making the fine-grain temporal distinctions that have come to represent acceptable tests capable of distinguishing these theories. The problem is, of course, that 'on-line' is not a categorical notion. Rather, tasks are only relatively more or less sensitive to the temporal course of information interaction. Two tasks that appear to 'tap' processing of a sentence at the same point in time may actually differ by a few hundredths of a second in terms of what their own task demands add to the process, and that difference may be sufficient to (in one case) capture a process or (in the other) just 'miss' capturing the process. The point here is that all 'on-line' tasks are not equal, and much of whether one believes that a task has 'demonstrated' interactivity or modularity relies on careful evaluation of whether or not that particular task is demonstrably sensitive to both the processing effects under investigation and their (exceedingly rapid) time-courses.

In what follows, two processes - co-reference assignment and lexical access - are examined for evidence of the effects of a number of types of semantic and pragmatic contexts upon their operation. A task is used that has been demonstrated to be sensitive to the results of such processes and, as will be seen, the evidence leads to an overall conclusion that both of these processes are autonomous, independent subprocesses of the language comprehension system.

ASSIGNMENT OF CO-REFERENCE DURING COMPREHENSION

One area of indeterminacy resolution that has particular significance for both 'competence' and 'processing' models of language lies in the area of co-reference. This is the manner in which a referentially-dependent item (such as a pronoun, reflexive, or 'empty category' element is linked to its referent). For example, in the sentence:

When the butcher saw the baker he was quite intoxicated.

the exact antecedent (referent) assigned to the word 'he' is indeterminate. However in order to understand the sentence (i.e., to have a meaning assigned to it) the listener must eventually assign one or the other of the two possible antecedents to the word 'he'. Now, in this particular case, the assignment must be made on pragmatic or strategic grounds: If the listener knows something about this butcher or this baker, or butchers or bakers in general, he/she may be able to decide which one is more likely to be intoxicated. However, if such pragmatic knowledge is unavailable, the listener may simply use strategies (such as: The first-mentioned person is the likely sentence topic, and hence the likely subject of the main clause). Indeterminacy resolution for cases such as this involving overt pronouns is an interesting and important topic for study. However, there is an even slightly more intriguing (and structurally parallel) case of co-reference resolution, that associated with empty categories, which will serve as the focus of inquiry here.

Empty categories are sentence constituents which have no phonological, acoustic or orthographic realization. They are theoretical entities which have roughly been a part of generative grammar at least since the appearance of *Aspects of a Theory of Syntax* (Chomsky, 1965; see Fodor, 1989, for an excellent extensive review of the empty category issue). They are entities postulated to exist when deletions or movement cause some piece of a sentence to be missing from its canonical position. Thus, for example, according to some approaches, the sentence:

Which book_i did Ed read (e_i)?

derives from movement of a wh-phrase ('which book') that is the object of the verb 'read' in the canonical (underlying) sentence

Ed read *which book?*

According to the theory, the wh-phrase is moved to the initial position in the sentence, and a 'trace' is left in the original position it was moved from. This trace is an empty category (e). It is important to add that, again according to

some theories, co-indexation between the moved wh-phrase and the trace is what allows the correct assignment of the thematic role of object (patient) to the moved phrase; the thematic role is assigned to the position the trace occupies, and it must be somehow linked to the moved phrase - hence co-indexation.

Note that it is not critical that this analysis be described in terms of current generative theory, as has been done above; many other theories could in principle accommodate an account of the apparent displacement of an object from its language-canonical position to some other place in a sentence. However, it is only the current generative models which have made the issue of empty categories a primary focus of linguistic inquiry. Hence, given that it is in terms of 'empty categories' that this concept has been most comprehensively elucidated, it makes sense to utilize this descriptive mechanism here.

The processing question that is raised by empty categories (and, for now, the focus will be on wh-trace), is how the antecedent to that trace is linked to the empty category during comprehension; how co-reference is accomplished between the antecedent (which book) and the trace (e) during actual processing.

There are a number of simple, preliminary hypotheses that one might make about how such relationships are established. To begin, one could imagine that while co-reference is established for overt pronouns, it may *not* be for empty categories; they may just be convenient tools of a linguist to keep grammatical theory consistent. Thus, under this model, there is no indeterminacy and thus one would expect to find no co-reference processing necessary. A second hypothesis is that the processor does nothing when some reference-seeking item is first encountered (be it a pronoun or an empty category), but rather simply waits until such relationships are clearly and uniquely disambiguated by later information. This Delayed Determination model has the apparent advantage of fitting with the fact that we aren't typically aware of linking a pronoun to some antecedent when we first hear it, and to the possibility that if we wait long enough, context may provide some help in the assignment process. The final hypothesis that will be considered is one which holds that the comprehension device cannot delay establishing co-referential relationships, because they help constrain other on-going processes, and hence the system will do everything it can to establish co-reference immediately. There are several reasonable variants of this Immediate Determination Model. One, something that might be called the weakly-constrained variant, holds that the device activates all possible antecedents (e.g., all prior noun phrases in the immediate discourse) in order to then sift through them to find the correct antecedent to a pronoun or an empty category. A more strongly constrained variant would allow linguistic principles to constrain the choice of the co-referent 'on-line'; failure of this to narrow the field to a single possibility might require guessing based on psychological processing principles (such as, for example, Frazier, Clifton & Randall's, 1983, "most recent filler" hypothesis).

Happily, there is at this point sufficient evidence to begin to choose among

these alternatives. A number of researchers, using different techniques, have provided evidence that.. for object-relative wh-traces the processing device appears to immediately re-activate the co-referent at the point where the trace appears; indeterminacy is resolved immediately. For example, Bever and McElree (1988) have used a probe-recognition task whereby they found that recognition of a probe word from an antecedent noun phrase (NP) is faster (with lower error rates) in cases theoretically containing these 'empty' wh-traces than in control conditions with no wh-trace. In another study, Tanenhaus, Carlson, and Seidenberg (1985) have demonstrated that lexical decisions for words phonologically related to the antecedent are faster than those to a control at a point shortly after the 'gap' is encountered in a sentence. (See also recent papers by Tanenhaus, Boland, Garnsey & Carlson, 1989; Garnsey, Tanenhaus & Chapman, 1989; Nicol & Swinney, 1989, among others, for relevant findings.)

One on-line study of co-reference resolution in wh-trace conditions by Ford, Frauenfelder, Bresnan and Swinney (1984; reported in Swinney, Ford, Bresnan, & Frauenfelder, 1988) will be presented here in a bit more detail, both because it employs one of the more temporally sensitive on-line tasks for examining this issue (cross-modal lexical priming), and because it leads directly to our current focus - the examination of the effects of contextual information on this process. In the Ford et al. study, subjects were presented auditorily (only the *words* were heard) with sentences such as:

The policeman saw *the boy* that the crowd at the party *¹ accused (*e*)
*² of the crime.

At the points indicated by the '*' a string of letters appeared on a computer screen in front of the subjects and they made lexical decisions to those letter strings. In the critical conditions, those letter strings were semantic associates to a potential antecedent from the sentence. All conditions were counterbalanced, so that any one subject saw only one possible associate in only one of the two test positions. The logic of the experiment is that, if reaction time to make a lexical decision to a word related to the appropriate antecedent (co-referent) to the trace is faster than reaction time to an unrelated, but otherwise equivalent, control word, that 'priming' result constitutes evidence that reactivation of that antecedent must have taken place at that point in time. Note that because there is no actual physical stimulus present at the 'trace' (this is also called a 'gap') such an outcome would be a particularly striking example of evidence for the psychological reality of these 'traces'. The results are easily summarized. When the letter string was related to one of the inappropriate potential antecedents (e.g., CROWD), reaction time to decide it was a word was no faster at test point *² than at test point *^{1,2}. However, when the probe word was related to the appropriate antecedent to the trace (in this case: BOY), priming was significantly larger at test point *² (precisely where the trace is posited to exist) than after test point *¹. In a follow-up experiment (Swinney,

Ford, Bresnan, Frauenfelder, 1986; reported in Swinney et al., 1988), the same materials were used with a so-called 'naming' task wherein reaction time for subjects to 'say out loud' the word that appears on the computer screen was made. Identical results were obtained with this naming paradigm as in the lexical decision paradigm. Thus, it was concluded that the appropriate co-referent was re-activated immediately after the verb (at the point where the trace is hypothesized to exist). Interestingly, there was no evidence that the appropriate co-referent was simply activated and kept active until the gap (trace) occurred in the sentence; the associate to that co-referent was not primed at test point *¹. Equally interesting is the finding that other 'potential' co-referents to the trace (such as CROWD) are not activated at the gap. Rather, it appears that the strongly-constrained version of the Immediate Determination Hypothesis described above seems to hold for wh-trace processing; at the point the trace is posited (after the verb) the appropriate antecedent is calculated by the processing device.³

Under any interpretation, the immediate reactivation of the co-referent of the trace in these constructions constitutes one of the few documented examples of an on-line structural process in operation. That is, the reactivation of the antecedent of a trace can only be needed for structural processing (or, perhaps, the vague shadowy area combining structural and semantic processing), and it is one of the few such automatic, perceptual process in structural processing for which we have any evidence. As such, it makes an interesting testing ground for the issues of indeterminacy raised above, those concerning precisely what conditions control the discovery and activation of the appropriate co-referent to the empty category. Couched in terms of the interactivity/modularity hypotheses, the question becomes one of whether information that is external to structural processing can affect the gap-filling procedure. Under the modularity hypothesis, information external to a putative structural processor will not be able to affect the gap-filling process. However, the interactionists would hold that, quite the opposite, any pragmatic or world-knowledge information that indicates the identity of the appropriate co-referent will aid in gap-filling.

The two experiments reported below examine this controversy in this unique structural processing condition. In the first, plausibility of the potential co-referents is manipulated to determine whether world-knowledge (in the form of plausibility) can control the gap-filling process. In this, 86 subjects from Rutgers University were subjected to 106 auditory presented sentences, 48 of which represented sentence versions schematized in the following example:

The crowd looked at the enormous heavyweight boxer
that the small 12-year old *boy* on the corner had
[beaten (*e*)* so badly / hugged (*e*)* so intensely]
a few minutes earlier.

Subjects were required to understand the sentences and to make a naming responses to words that were visually displayed at the test points (*). Again, no subject saw or heard more than one variant of the materials; complete counterbalancing was performed. The visual words were either the pragmatically likely antecedent (BOXER for the verb 'BEATEN') or the syntactically correct antecedent (BOY), or 'control' words matched for a priori naming reaction times with each of the experimental words, but semantically unrelated to anything in the sentence. Note that a control verb was used (HUGGED in this example) for which either antecedent could have been pragmatically plausible. The logic of this experiment is that if plausibility can "penetrate" (control the internal operations of) a putatively modular syntactic process such as co-reference assignment, then we will find that following the verb 'BEATEN' there is 'priming' for the word BOXER even though it is not the structurally appropriate antecedent. If, on the other hand, modularity does exist for the structural processing device, then such plausibility information should have no such effect; one should only find priming for the correct co-referent (BOY) in both the context of the implausible verb (BEATEN) and the neutral verb (HUGGED).

Reaction time to name the words in the conditions described above were collected, and are reported here in terms of 'priming scores' (reaction time to control minus related word). When the control verb HUGGED was used, there was a 31 millisecond priming effect for BOXER ($p < .05$), but only a non-significant 2 milliseconds effect for BOY at the test point. When the verb BEATEN was in the sentence, there was a 29 millisecond priming effect for BOXER ($p < .05$), but only a non-significant -6 millisecond effect for BOY at the test point. As is evident from these data, there is no evidence of plausibility information 'penetrating' the structure-based co-reference process in these conditions. Apparently, information about the plausibility of the co-referent as a possible object for the verb cannot direct or interact with the co-reference assignment process.

It might seem, however, that the relative plausibility of the structurally appropriate antecedent simply might not be sufficiently strong to be used by the comprehension device to guide co-reference assignment (although there is certainly no a priori reason to believe this to be true). Thus it might be argued that a stronger test of whether world knowledge can affect the internal operations of a structural processor in indeterminacy resolution would be to make the verb-object relationship one of impossibility rather than implausibility. To that end an experiment was run in which subjects were presented sentences (auditorily) in which the structurally appropriate antecedent was not only implausible, but was actually not a possible object of the verb such as:

The police captain said that the *cop* from his precinct that
the soup in the bowl had [eaten / splashed] (*e*)* was
going to give a talk on public service.

In this sentence, while soup can *splash* the cop, soup cannot *eat* the cop. Thus, world knowledge about what things can eat and be eaten should allow for a prediction that it must be the cop doing the eating not being eaten. If the co-reference processor can take advantage of this knowledge, then one should find that the structurally correct antecedent (COP) for the trace should not be activated at the gap after the verb EATEN (but it should be activated after SPLASHED). This experiment involved 32 experimental items given to 74 subjects (distributed across 8 materials conditions), in a naming response paradigm. In this, subjects saw either the word SOUP or BOY (or control words for each of these) at the test point (*) and were required to 'name' the words as quickly as possible. Reaction times to name the words were collected. In the condition involving the control verb (SPLASHED) the priming score for the structurally appropriate antecedent (COP) was 48 millisecond (significant at $p < .05$). The priming score for the structurally inappropriate antecedent (SOUP) was a non-significant -2 milliseconds. Importantly, the same pattern held for conditions with the verb producing the impossible reading (EATEN): The priming score for the structurally appropriate, but impossible antecedent (COP) was 41 milliseconds (significant at $p < .05$), while there was a non-significant 7 millisecond priming score for the structurally inappropriate, but semantically possible, antecedent (SOUP).

Thus, it appears that the structural processor can not take into account information about the plausibility of the potential co-referent or even information indicating that certain nouns cannot possibly be arguments of verbs in initially assigning co-reference to a wh-trace. This evidence fits with work by Tanenhaus, Stowe and Carlson, 1985, which used an embedded anomaly technique in which they demonstrated that sentences containing a potentially anomalous (but, eventually not) antecedent-trace relationship were judged to not make sense more frequently than those without an anomaly. Similarly, N400 studies of evoked brain potentials for such embedded anomalies also demonstrate evidence supporting the view that an antecedent will be immediately assigned to a gap position, even when that antecedent is anomalous in that position (Gamsey et al., 1989). Thus, it appears that arguments for interactivity of information sources are not upheld in this co-reference indeterminacy-resolution situation, and that the best interpretation of the data is that co-reference assignment is an autonomous modular process.

Of course, no one doubts that, at some point in processing, information such as plausibility may have some effect upon the final interpretation of an utterance. Thus, in an effort to discover *if and when* the plausibility information in the experiment described above has any effect on processing, a follow-up study was embedded in the original cross-modal paradigm experiment. In this, subjects were asked, after half of the experimental trials, who did what to whom. For example, for the sample sentence given above ("The crowd looked at the enormous heavyweight boxer that the small 12-year old boy on the corner

had [beaten so badly ... /hugged so intensely...]), subjects were asked: Who was hugged? or Who was beaten? There were only a total of 1% errors in responses for the HUGGED condition. However, 27% of the time, subjects reported the incorrect (but more plausible) object in the BEATEN condition. Thus, while the autonomous, modular process that provides for immediate co-reference assignment in wh-trace conditions does so *independently* of plausibility information, the final interpretation that becomes consciously available to the listener is affected by the plausibility information in about 1/3 of the cases. However, such a plausibility effect is not uniform, and takes place far after the initial co-reference assignment, as predicted by the modularity model.

There are, of course, other types of empty categories, some of which might be deemed to have different properties than these wh-trace object-relative constructions we have just examined. For example, unlike the wh-trace condition, in which an empty category is posited to exist when an NP is moved, an empty category called PRO has been generally deemed to be a base-generated component. It exists in English in untensed constructions involving infinitives or gerunds, in which the semantic subject is always omitted. For example(s):

*The boy*_i decided PRO_i to go to the store.

or

The boy decided that PRO running home was what he should do.

In both cases, the subject of the infinitive (to go) and the gerund (running) are phonologically 'empty', but, according to certain linguistic theories, must be represented in the underlying structure, as done here with the term 'PRO'.

In a recent publication examining the role of pragmatic information on assignment of antecedents to such constructions, Marslen-Wilson and Tyler (1987) presented subjects with short paragraphs such as:

As Philip was walking back from the shop he saw an old woman trip and fall flat on her face in the street.
She seemed to be unable to get up again.
(He ran toward ...// Running toward)

Once subjects heard either the "He ran toward.." or the "Running toward.." endings to the paragraph, they were presented with either the word HIM or the word HER and were asked to make a lexical decision to it. The logic of this experiment was that, since in the first ending (He ran toward..), the word HE (Philip) is assigned as the subject, the object (the person HE was running toward) can be anticipated as being the old woman (HER), and hence reaction time was predicted to be faster to the word HER than to the word HIM, just in

the case that the context had been used predictively. That is precisely the reaction time result that was obtained. The argument regarding the gerundive ending (Running toward...) is roughly the same, with the additional assumption that, if pragmatic context from the paragraph CAN be used to penetrate (predict) co-reference assignment, then even when there is no explicit subject of 'running' provided, as in the gerundive case, pragmatics determine that it is only Philip who can run in this scenario, and thus, Philip will be immediately assigned as the co-referent of the empty category PRO. Then, once that assignment is made, subsequent lexical decisions to HER will be facilitated over those for HIM at the test point, for precisely the reason in the infinitival case; namely, that since the subject position of run is filled by Philip, then Philip must be running toward the old woman (HER). This, too, is how their data came out. At first blush, this result appears to strongly support a highly interactionist account of co-reference assignment for the empty category 'PRO'. It would seem that the discourse context allows prediction of both the fact that HER is the likely object of who is being RUN TOWARD and prediction of who the likely subject of RUNNING must be. However, such a conclusion would be premature. The issue here is a methodological one - an issue related to the relative sensitivity of on-line tasks, as mentioned in the introduction. The task employed in the Marslen-Wilson and Tyler study, while fairly immediate, does make its test of co-reference assignment at a point considerably downstream (temporally) from occurrence of PRO (at least 1 second later - a long period in terms of sentence processing operations). And, it is at least possible that assignment of antecedents in these materials doesn't actually take place until the pronoun HER or HIM is seen by the subject. In short, the finding that Marslen-Wilson and Tyler report could be a function of a process taking place at the point of the test (when the words are presented on the screen) - and not as a result of a discourse or pragmatics 'predicting' the assignment of the appropriate co-referent during normal sentence processing.

In order to examine that possibility, Fodor, Garrett and Swinney (1990) have recently undertaken a series of studies which used the cross-modal priming technique to determine precisely *when* assignment of an antecedent to the PRO in this construction takes place, and whether pragmatically biasing discourse contexts such as those used by Marslen-Wilson and Tyler (1987) can be used by the comprehension device to predict/speed such assignment. The materials were identical to the structure of the example given above, with the exception that the last sentence had a modifying adverb placed between the verb and the preposition and the final sentence was presented in completed form. So, following the example given above, the last line was changed to read as either:

*¹ He ran rapidly *² toward her *³ and tried to help her get up.

or

*¹ Running rapidly *² toward her *³, he tried to help her get up.

As in prior studies with this technique, subjects heard one or the other version of the paragraph auditorily, and were given visual probes at the points indicated by the (*) consisting of either the word PHILIP or an unrelated control word matched for naming time in a pre-test, to which they were to make a naming response. Reaction times to name PHILIP or the control word (in this example, DAVID) were recorded. The logic of the experiment is that if and when the co-referent of PRO is assigned as PHILIP, reaction time to name PHILIP would be facilitated compared to that to a matched control word.

For the 'control' sentence condition containing an overt pronoun (*He ran rapidly toward her...*), response times were as predicted from previous studies: There was a significant 32 millisecond ($p < .05$) effect of priming for PHILIP at *¹ (the probe appeared coincident with onset of the word HE in this sentence), indicating that PHILIP has been assigned as co-referent of HE immediately. This priming effect was still active at test point *², where a significant 50 millisecond priming effect was found. Finally, there was no significant priming for PHILIP at test point *³ (coincident with onset of the word HER). In the critical experimental condition with the empty category (*Running rapidly toward her...*), it was found that there was a nonsignificant 8 millisecond priming effect at test point *¹, indicating that, coincident with the onset of the gerund, no co-referent for the PRO had been assigned. Similarly at test point *², a nonsignificant 22 millisecond effect was found for PHILIP vs. the control word. Thus, at this point, no assignment of an antecedent for the empty category PRO had yet been made. Finally, at test point *³ a significant priming effect for PHILIP of 37 milliseconds was found. This suggests that it wasn't until subjects actually heard the word HER (or, in the case of the Marslen-Wilson and Tyler experiment, saw the word HER) and processing was begun on finding the antecedent to it, that assignment of the co-referent to the PRO element was undertaken. Indeed, when this took place, it appears that the discourse context may have aided such final assignment, but it is clear that the discourse context in no way predicted or allowed for prior assignment of the co-referent in this condition. In short it appears that in this case of an empty category, as in the one for wh-trace, there is no evidence that supports a highly interactive or predictive view of structural processing and co-reference assignment. Again, it is only through use of a sufficiently temporally sensitive technique that these conclusions can be reached.

To summarize the findings of this section that are relevant to the issue of indeterminacy resolution during comprehension: It appears that wh-trace co-reference assignment - a structural based process - is an autonomous, modular subsystem that is not affected by plausibility or impossibility (anomaly) in its operation. Resolution of the uncertainty over the identity of the co-referent is accomplished immediately on the basis of structural information alone. The

resolution of indeterminacy concerning the identity of the co-referent of a PRO (a construction which is not identifiable until after an infinitive or gerund has been heard) appears neither to be made immediately, nor is it predictively made through use of discourse pragmatics.

LEXICAL ACCESS DURING SENTENCE COMPREHENSION

Over the past two decades, psycholinguistic study of indeterminacy resolution during language processing (and its corollary issues concerning modularity and interactivity) have had a major testing ground in the area of lexical ambiguity resolution. Although the battle is far from definitively resolved, much on-line work in the past 10 years has led to acceptance of a general view that while context and meaning-dominance effects are generally used quickly to determine which interpretation of a lexical ambiguity is the correct one, there is strong reason to believe that *initial* access of meanings for these words *in sentences* involves exhaustive, modular, context-independent processing (see, e.g., Simpson, 1981; Swinney, 1979; Seidenberg, Tanenhaus, Leiman, and Bienenkowski, 1982; to name but a few of the studies that allow for this interpretation). Recently, however, work by Tabossi and her colleagues (Tabossi, 1988, Tabossi, Colombo, & Job, 1987) has challenged the view that even initial access in this process is context-independent. Tabossi performed a set of clever experiments designed to demonstrate that prior work had simply utilized inappropriate contexts in attempting to discover contextually driven lexical access. It was argued, for example that Onifer and Swinney (1981) used contexts that "biased no particular aspect of ...meaning" for a lexical ambiguity, and that, by contrast, she and her colleagues had provided a way of making "salient a very characteristic feature of either the dominant or the subordinate meaning of the ambiguous word" in such a way as to "establish constraints on the information provided the upcoming ambiguous word."⁴ Because this work stands as the only current on-line study that seriously challenges the modularity-of-initial-access view, this section will briefly examine her study and describe a preliminary attempt to replicate it in English. Its relevance to the issue of precisely how indeterminacy of meaning for an ambiguous word is resolved is self-evident; if Tabossi is correct in her analysis, then contextual information (of the right type) *can* direct the lexical access process and, thus, modularity of the putative lexical access process is not upheld.

The following facts about the Tabossi study appear to be relevant. First, Tabossi employed a total of nine ambiguous words in these studies. Second, all of the work was in Italian. Third, the critical contexts for these studies consisted of the following: "For each of the nine ambiguous words, a sentence was constructed so as to render a central aspect of its dominant meaning particularly salient." It is important to note here that in neither the Tabossi, 1988, or the Tabossi et al., 1987 paper are operational definitions provided concerning what

a priori criteria constitute achievement of these context conditions. Rather, the reader is only told that sentences constructed with this intent are submitted to "judges to point out what aspect of the ambiguous word each sentence makes them think of. For eight of the nine sentences there was 75% agreement on the intended aspect, while no such agreement was reached for the ninth sentence which was modified accordingly." Tabossi, admirably, publishes the nine experimental items, and the following English sentence is her translation of one of her Italian experimental sentences. It is provided here to give the reader a sense of what the special contexts were like. The ambiguous word is given here in Italian, so as not to bias the reader. The two interpretations of that word, in English, are given below the sentence:

The water in the bay was so calm that it seemed to be in a
STAGNO, rather than in the sea.

(STAGNO = pond or tin)

Tabossi, using a cross-modal lexical priming task (Swinney, Onifer, Prather & Hirshkowitz, 1979), presented probe target words at the offset of the ambiguous word in the sentence. These probe words were carefully chosen to be words that denoted very characteristic aspects of the meanings of the ambiguity, but were not highly associated to it (unlike the words of Onifer & Swinney, 1981, which were associates to the meanings of the ambiguity). In this, Tabossi derived these words by having 12 judges produce "relevant semantic aspects of both meanings for each ambiguity... the criterion used was that each aspect was mentioned by at least nine of the judges." Examples of the English translation of the Italian probe words for the above example are: FROG and LEAD). Tabossi then ran a cross-modal lexical decision task with these materials and reported that "lexical decision on the visual word associated with the dominant, contextually congruent meaning of the ambiguity was significantly faster than lexical decision on both the visual word associated to its subordinate, contextually incongruent meaning, and the control word" (Tabossi, 1988, p. 333). Thus, it was claimed that resolution of indeterminacy concerning the appropriate meaning of an ambiguous word can occur prior to lexical access.

In order to examine this claim in a bit more detail, an experiment by Swinney, Islewitz and McKinnon was performed in an attempt to replicate and extend the Tabossi findings in English. In this study the materials were developed as described in the Tabossi studies. Twenty judges were used to rate 2 versions each of 55 sentences containing ambiguous words that had been created with the intent of causing a central aspect of each meaning of the ambiguity to be particularly salient. In addition, although it will not be discussed in detail here, some effort was made to classify the nature of the relationship existing between the context and the ambiguous word, so that some objective

operational definition of the context-producing procedure could be developed. Out of this, four general classifications emerged: 1) materials where the context suggests one meaning is more likely than another (plausibility); 2) materials where the context renders one meaning very likely through associativity or likelihood relationships, but does not rule out the other interpretation; 3) materials where the context rules out one of the meanings as being possible (anomalous); 4) materials where the context rules out one interpretation, and also provides a basis for anticipation or prediction of the other interpretation (usually based on association-type links). Materials making up the 55 sentences to be judged fell into categories 2, 3, and 4. Once judges had followed the Tabossi procedure, the 55 sentences were analyzed and the Tabossi criterion of 75% of the judges having to agree on what aspect of the ambiguous word each sentence made them think of was used. Forty sentences were chosen that met this criterion. These sentences were mixed with 60 other 'filler' sentences in an experimental script. In the preliminary study reported here, only the dominant-meaning biased version of each of these sentences was examined (dominance had been judged on an *a priori* basis by a separate group of subjects). Probe words were chosen as in Tabossi, such that the probes were judged to be words that denoted very characteristic aspects of each of the meanings of the ambiguity, but were not highly associated to it. These items were each matched with a frequency and length matched control word. The 'related' word and its matched control were demonstrated not to differ in reaction time from each other in a pre-test involving 35 subjects.⁶ A sample item is as follows:

His favorite aspect of going to bed
was anticipation of slipping between the clean,
cool, crisply ironed white *sheets* that
were provided fresh each morning in the Inn.

related probe word: REFRESH
probe word for 'other' meaning: WRITE

In this study involving 61 subjects, the findings reported by Tabossi were partially replicated in that there was significant priming (43 milliseconds, $p < .01$) demonstrated for the word related to the central feature of the contextually biased ambiguity (REFRESH) compared to its control. However, there was also significant priming (37 milliseconds, $p < .01$) for the word related to the 'other' meaning of the ambiguity (WRITE), compared to its control. While it is true that there was less priming for the secondary interpretation of the ambiguity than for the dominant sense, it should be remembered that all the the words in the context were related to the dominant meaning, and undoubtedly contributed to the size of the effect on that meaning.

Thus to the degree that we have managed to re-create the conditions used

by Tabossi (and we have tried to do so accurately) we have thus far failed to find the contextually-driven access process Tabossi reported. Rather, we find evidence, as many others have, for initial context-independent, modular, lexical access. There are any number of reasons why we may have achieved this finding rather than replicating the Tabossi results. However, it seems to us that two of the major issues that must be considered are: 1) That there is simply something special about the processing of ambiguous words in Italian - something tied to how common ambiguities are, etc. However, because we are not native speakers of Italian we can offer only vague speculation here; 2) It may be that in using only nine items and a limited number of subjects the Tabossi experiment simply did not provide a sufficiently strong forum in which to test for statistically significant activation of the 'contextually irrelevant interpretation' of the ambiguity. The story is clearly not yet complete, in that our judges may not have been giving the same information in their ratings as the Tabossi judges, etc. However, until that information is quantified and described in sufficient detail, such speculation is impossible to either substantiate or dismiss. The most parsimonious interpretation of the evidence at this time, then, is that, even with the special 'featural' contexts of Tabossi, at least in English, the lexical ambiguity resolution process begins with exhaustive, modular access of multiple interpretations of the ambiguity, and identity of the intended meaning is resolved thereafter on the basis of contextual information and dominance (inherent bias).

Finally, an issue mentioned briefly in the introduction deserves a brief further discussion. Consider the role of evidence of the type we have examined here in relation to models of language theory - things such as linguistic grammars (e.g., Government and Binding Theory (Chomsky 1981) or Lexical Functional Grammar (Kaplan & Bresnan, 1982), or Generalized Phrase Structure Grammar (Gazdar, Klein, Pullum & Sag, 1985).

Such models are, of course, meant to be abstractions which capture significant generalizations about what it is that we know when we know language. This can be contrasted with performance models which are meant as implementation of that knowledge. (Marr & Poggio, 1976, have nicely captured the difference in terms of a contrast between a descriptive level at which the nature of computation is expressed and a level at which algorithms that implement that computation are characterized - sort of a "WHAT" vs "HOW" distinction.) However, in practice, even abstract "WHAT" models are grounded in expressions that capture empirical "HOW" facts.⁷ In short, models of linguistic theory, however abstract, are generalizations over some processing data base and, in practice, what is often really at stake in competence/performance distinction arguments is the type of evidence that the abstract (formal) competence model should abstract over.

The problem, then, in maintaining that there is an important distinction to be made between competence and performance models⁸ is to determine which

processing details constitute critical (primitive) facts to be captured by the abstract model and which processing details are merely alternate computational options for implementing the theory. For example, if one were to believe that there were a known physical hardware (brain) limitation on certain types of computations, it *might* be important for that to be captured in the abstract theory (one can't be certain about this, of course, until a complete, working abstract "theory of the computation" is in place). This same argument holds for functional-level (perceptual and cognitive) "facts:" If a particular arrangement of processing is known to hold (be a constant) at a functional level, it may well be that such a fact is one that should be captured by a abstract "theory of the computation" (linguistic grammar).

There are several things that lead to the reasonable belief that linguistic theory might do well to capture details of modularity of mind, as expressed in empirical work (and not simply as *a priori-based* assumptions of modularity). First, given that there is no monolithic linguistic theory - no uniformly 'correct' representation of language knowledge that captures all of the significant facts - there is no reason to think that linguistic grammars, as currently formulated, are going to work. Second, as suggested above, current linguistic grammars are far from 'pure' abstractions across performance characteristics, anyway. Many incorporate processing-like information in their configurations (everything from verb bias information to relatively arbitrary assumptions about what constitute distinctions among information types; e.g., morphological vs. lexical vs. phonological information type distinctions are often arbitrary divisions). Third, it should be pointed out that even those data that are traditionally allowed to form the basis of linguistic theories do not constitute a uniform data base. While the 'intuitions' of speakers/hearers of the language are typically taken to be the final arbiter of sentence acceptability, the fact is that, in practice, there are several types of such intuitions used in linguistic inquiry (each type calling on potentially quite different processing operations)⁹ and in many of the interesting cases it is very difficult to get concurrence of such judgments. That is, different degrees of training result in different levels of acceptability for sentences, and even equivalently trained linguists do not always agree.

The point here is simply that if we want our theoretical models of language knowledge to accurately capture significant generalizations about language, it *may* be necessary for them to capture aspects of modularity of processing. If so, it may be necessary for such theoretical models to accommodate to considerations of temporal events in language processing. That is, in order to accurately capture formal 'knowledge' divisions among information types, it may be necessary for such models to account for the relative temporal order in which such types are recruited in processing (for indeterminacy resolution, for example). Ultimately, of course, whether this suggestion holds to be a useful one will be a matter of determining if such information can improve such models, or whether correct knowledge representations in language can be made

without use of such information. One thing which is certain, however, is that, independently of whether models of language theory can best use such information, correct *processing* models cannot be built without precise detailing of facts about modularity, such as those we have begun to provide here.

Notes

¹This is a perspective which will be re-examined at the end of this paper, following presentation of empirical evidence about language processing.

²Actually, the data examined at each test point were priming scores - the difference in reaction times between the control word and the related word. Thus, the difference between test points 1 and 2 described here are actually differences in magnitude in priming at each of the two points - and interaction of the related vs control word for test points 1 vs. 2.

³Note that it is entirely possible that wh-trace represents a very special case for gap-filling. It is well marked by the relative marker 'that,' it is controlled by the verb (i.e., the verb either takes a direct object or not), and it is in an argument position. All of these lead to the possibility that the immediacy in co-reference assignment is a function of some structural issue other than the 'trace' aspect of this example. For example, it could be the case that argument positions (or the moral equivalent in anyone's favorite terminology) in the sentence (e.g., agent, theme, patient) must be immediately filled in order for both structural and semantic interpretation to proceed. It may well be that other empty categories are not dealt with in the same fashion.

⁴Quotes are taken from page 327 of Tabossi (1988).

⁵In order to avoid misinterpretation, details here and below are quoted from Tabossi (1988) and Tabossi et al. (1987).

⁶Tabossi reports a similar manipulation, but used only one control word for the two related words. Unfortunately, while she reports finding no differences in reaction time among the three words in an *a priori* isolated test, the reader is not told how strong a test looking for pre-existing differences was conducted (e.g., we don't know how many subjects were involved in this test). The only reason that this seems to be a potential problem is that some of the members of the triplets appear to be very different in frequency and length. For example, probes for the ambiguous word (one of the 9 she used) *coppa* (which means either a bowl or a type of salami) were: *spumante* (champagne), *salame* (salami), or *libro* (book, the control word).

⁷Relatedly, there is a well-known set of arguments which claim, for example, that there is little to distinguish so-called competence grammars from performance grammars in linguistics, particularly since the former are simple abstractions over a certain *type* of performance evidence - namely, linguistic intuitions. The term 'abstract performative grammars' was coined to capture

this fact (see, e.g., Watt, 1972).

⁸This is an 'in practice' argument; in the ideal they would be distinct.

⁹This can be seen, for example, in cases where linguistic evidence is gathered by simply asking 'native speakers' if a given sentence is 'acceptable' as compared with asking whether it is 'grammatical.' The former potentially recruits what you think is meant by the term 'acceptable' to affect decisions about whether proper form is used or not. Of course, native listeners may respond to the two types of questions the same way, which in a sense is worse for linguistic theory, as one might never know what is motivating acceptance or rejection of a sentence. Contrast these two (differing) conditions in which such judgements are gathered with those in which subjects are asked to decide "which of these sentences is better?" This latter is a common request of linguists in determining which of two possible filters or rules has precedence. However, as is well known, there are quite different *performance* characteristics that accompany forced choice relative decisions (which is better?) vs. absolute categorization (is this sentence good or bad?). The point is not to belabor which method is the better - neither is. The point is simply that these different processing tasks, with different performance criteria, are used in linguistic theory building as though they were uniform, unbiased windows on mental representation. Clearly they are not.

References

- Bever, T. G. & McElree, B. (1988). Empty categories access their antecedents during comprehension. *Linguistic Inquiry*, 19, 35-43.
- Chomsky, N. (1965). *Aspects of a theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1981). *Lectures on government and binding*. Dordrecht: Foris.
- Fodor, J. A., (1983). *Modularity of Mind*. Cambridge, MA: MIT Press.
- Fodor, J.A., Garrett, M.F., & Swinney, D.A., (1990). *The role of discourse contexts in constraining sentence processing*. Manuscript in preparation; paper presented at CUNY Third Annual Conference on Language Processing, New York.
- Fodor, J.D. (1989). Empty Categories in Sentence Processing. *Language and Cognitive Processes*, 4, 155-209.
- Garnsey, S. M., Tanenhaus, M. K. & Chapman, R. M. (1989). Evoked potentials and the study of comprehension. *Journal of Psycholinguistic Research*, 18, 51-60.
- Gazdar, G., Klein, E., Pullum, G. K. & Sag, I. A. (1985). *Generalized phrase structure grammar*. Cambridge, MA: Harvard University Press.
- Kaplan, R. & Bresnan, J. (1982). Lexical-functional grammar: A formal system for grammatical representation. In J. Bresnan (Ed.), *The mental representation of grammatical relations*. Cambridge, MA: MIT Press.

- Man, D. & Poggio, T. (1976). *From understanding computation to understanding neural circuitry*. MIT A.I. Memo 357, May, 1976.
- Nicol, J. & Swinney, D., (1989). The role of structure in coreference assignment during sentence comprehension. *Journal of Psycholinguistic Research*, 18, 5-19.
- Onifer, W. & Swinney, D. A. (1981). Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. *Memory & Cognition*, 9, 225-236.
- Seidenberg, M. S., Tanenhaus M.K., Leiman, J.M., & Bienkowski, M. (1982). Automatic access of the meaning of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology*, 14, 489-537.
- Simpson, G. B. (1981). Meaning dominance and semantic context in the processing of lexical ambiguity. *Journal of Verbal Learning and Verbal Behavior*, 20, 120-136.
- Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re) consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18, 645-660.
- Swinney, D., Ford, M. Bresnan, J., & Frauenfelder, U. (1988). Coreference assignment during sentence processing. In M. Macken (ed.) *Language structure and processing*. Stanford, CA: CSLI.
- Swinney, D. A., Onifer, W. Prather, P., & Hirshkowitz, M. (1979). Semantic facilitation across sensory modalities in the processing of individual words and sentences. *Memory & Cognition*, 7, 159-165.
- Tabossi, P. (1988). Accessing lexical ambiguity in different types of sentential contexts. *Journal of Memory and Language*, 27, 324-340.
- Tabossi, P. Colombo, L. & Job, R. (1987). Accessing lexical ambiguity: Effects of context and dominance. *Psychological Research*, 49, 161-167.
- Tanenhaus, M. K., Boland, J., Garnsey, S. M., & Carlson, G. N. (1989). Lexical structure in parsing long-distance dependencies. *Journal of Psycholinguistic Research*, 18, 37-50.
- Tanenhaus, M. K., Carlson, G. N., & Seidenberg, M. S. (1985). Do listeners compute linguistic representations? In D.R. Dowty, L. Karttunen, & A.M. Zwicky (Eds.) *Natural language parsing: Psychological, computational, and theoretical perspectives*. New York: Cambridge University Press.
- Tanenhaus, M., Stowe, L. & Carlson, G. (1985). The interaction of lexical expectation and pragmatics in parsing filler-gap constructions. In *Proceedings of the Seventh Annual Cognitive Science Society Meetings*, pp. 361-65.