

Lexical Ambiguity Resolution:

*Perspectives from Psycholinguistics,
Neuropsychology, and Artificial
Intelligence*

Contributors

Geert Adriaens
Curt Burgess
Claudia Brugman
Cristiano Castelfranchi
Garrison W. Cottrell
Kurt P. Eiselt
Dan Fass
Ilene M. France
Dedre Gentner
Helen Gigley
Richard H. Granger, Jr.
Graeme Hirst
Jennifer K. Holbrook
Alan H. Kawamoto

Marta Was
George Lakoff
Steven L. Lytinen
Edward H. Matthei
Robert Milne
Domenico Parisi
P. A. Prather
Mark Seidenberg
Greg B. Simpson
Steven L. Small
David A. Swinney
Patrizia Tabossi
Michael K. Tanenhaus
Cyma Van Petten

Edited by

Steven L. Small,
Department of Computer Science, University of Rochester
Garrison W. Cottrell,
Institute for Cognitive Science, University of California, San Diego
Michael K. Tanenhaus,
Department of Psychology, University of Rochester

Morgan Kaufmann Publishers, Inc.
San Mateo, California

Chapter

11

Lexical Processing and Ambiguity Resolution: An Autonomous Process in an Interactive Box

P. A. Prather
University of Illinois, Chicago

David A. Swinney
The Graduate Center
City University of New York

1 Introduction

Examinations of the nature of lexical ambiguity resolution have formed the major arena for study of a number of the basic questions concerning natural language understanding. Chief among these, and one which has been most clearly elaborated and detailed by Fodor [1983], is that concerning mental modularity. This issue, which revolves around the question of whether human cognition is, in general, comprised of a set of informationally encapsulated, autonomous subsystems, or whether it is a totally interactive system, lies at the heart of our basic conception of the nature of cognitive activity. There appears to be undeniable support for the modularity thesis at some level of processing

(see e.g., [Fodor, 1983; Garrett, 1985]). However, that portion of the argument specific to the modularity of *lexical* processing has proven extraordinarily vexatious. Examinations of lexical autonomy/modularity represent a field of endeavor that has produced seemingly conflicting evidence and a myriad of counterclaims throughout the short history of psycholinguistics.

In what follows, we organize old and present new evidence that bears on the issue of whether or not initial word perception is accomplished by an autonomous lexical processing subroutine—in short, a module. In the course of this analysis, we examine critical issues and relevant evidence concerning lexical ambiguity resolution.

Ambiguity is ubiquitous in language; it exists at every level of processing (from acoustic/phonetic to semantic to structural, etc.). An obvious, but important, point about this observation is that there is no reason to believe (nor has anyone seriously proposed) that these disparate types of ambiguity are each resolved in the same manner during language processing. We argue here that this same general point holds true for lexical ambiguity: there is no uniform, invariant solution to lexical ambiguity resolution (although such solutions may be monogenic). Rather, lexical ambiguity resolution is a complex function of at least four general types of issues: (1) The nature of the temporal processing requirements during perception (e.g., speed of utterance, speed of required response, etc.). (2) The nature of the available "contextual" information in which the ambiguity occurs (e.g., type of constraints, level of information, relationship to ambiguity). (3) The nature of the processing strategies recruited for analysis (e.g., automatized routines, intentional strategies, linguistic/perceptual vs. problem solving routines, etc.). (4) The nature of the response required of the processing system or, in experiments, of the experimental situation (e.g., passive comprehension, rapid recognition, use of the ambiguity in novel manner, etc.). The complexity and interactivity of these issues cannot be overrated.

This chapter will elaborate on certain aspects of each of these issues in ambiguity resolution, with a focus on the underlying question concerning mental modularity. In what follows, we detail evidence which we feel strongly supports the view that there is a brief period of time in the perception of a word during which the lexical processing routine cannot be interrupted or influenced by prior information or by other ongoing, non-lexical processes. We conceptualize this as a narrow temporal "window of autonomous lexical processing, a "window of contextual impenetrability," if you will. However, the "window" accounts for only a small part of what traditionally is considered to be lexical processing, and the temporal parameters of this window are not fixed. Whether or not the lexicon is considered to be an autonomous module will be based in large part, we argue, on how the boundaries (temporal and otherwise) of the lexicon are defined. We will also examine lexical ambiguity resolution evidence which has been argued to support the non-modularity view of lexical processing, and we provide an alternative explanation of these data, one fo-

cussed on issues of specialized processing strategies and erroneous initial assumptions underlying these experiments. In addition, we present some data to substantiate our arguments.

We begin with a brief consideration of the parameters of lexical processing, following which we present evidence supporting the view that there is a brief "window" of contextual invulnerability in lexical processing that meets the criteria for autonomy. We then consider the implications of this evidence for characterization of the lexicon as a truly modular component in language processing in the light of certain counterclaims. Finally, we briefly consider the merits of recent claims about what has been called backward priming, claims that have been argued to undermine the autonomy view.

2 Parameters of Lexical Processing

Current models of cognitive organization typically allow that there are several separable domains of cognitive processing, among them the language domain. Within the language domain, it is generally held that acoustic/phonetic, lexical, and semantic and syntactic processes are at least potentially separable information sources. Garrett [1984], Forster [1979], Fodor [1983], and Morton [1982], among others, have argued that in fact these are separate and autonomous subroutines or sub-modules within the language domain, each potentially operating at least somewhat in parallel with the others, in cases where interactions among these information sources are limited to the outputs of these autonomous modules. (Note, however, that there is much disagreement among those cited above concerning the degree of parallelism in the system.)

Within this conceptualization, the lexicon is generally characterized as that module within which either the acoustic/phonetic or orthographic code (depending on input modality) is connected with an informational data base for that code, a data base which contains information that allows assignment of word meaning, recognition that the code constitutes a word, or whatever else is required of the processor. Word understanding thus involves perception of the speech/orthographic code which provides access to this data base; word production involves the reverse of this process of assigning a code from a given meaning/data base (see e.g., [Garrett, 1984; Bowles and Poon, 1985]). One of the major areas of controversy in characterizing the lexicon is the question of whether or not the assignment of meaning in comprehension (the process typically labeled lexical access) is influenced by non-lexical knowledge. Simply put, can a prior context direct access in the lexicon? The implications with respect to claims of modularity are obvious: if lexical access is penetrated by prior context, then it is not an autonomous process. If, alternatively, access operates independent of contextual effects, then it is autonomous (at least with respect to semantic information). This question has been examined empirically

primarily by studying the assignment of meaning to ambiguous words—the resolution of lexical ambiguity. If a prior, disambiguating context serves to restrict access to just one meaning of an ambiguous word, the autonomy view of lexical access is undermined.

Results from studies examining effects of a prior context on disambiguation have appeared to sometimes support and sometimes contradict autonomy claims. Before examining this evidence, however, a brief consideration of the temporal characteristics of the lexical processing is in order. As is undoubtedly clear, our decision about whether or not to characterize lexical processing as autonomous will hinge on careful assessment of precisely what is to be considered as purely lexical processing and what is categorizable as being outside of (either pre- or post-) lexical processing.

2.1 *The Time Course of Automatic and Intentional Lexical Processes*

Inferences about the nature and time course of lexical processing have been derived primarily from various types of "on-line" studies, the most influential of which to date have been those using semantic priming effects. In initial demonstrations of the priming effect, it was shown that prior processing of a word facilitates subsequent (repeat) processing of both that word and words highly associated with it [Meyer and Schvaneveldt, 1971]. Since then, evidence has been accumulating to support an inference that priming is an obligatory, fast-acting and autonomous process—in short, an automatic one (e.g., [Neely, 1977; Seidenberg et al., 1984]).

In one of the clearest demonstrations that priming is automatic and obligatory, Neely [1977] told subjects to expect words from a particular category (building parts) following words from another category (body parts). Subjects were presented word pairs sequentially, and asked to make lexical decisions on the second word. On most trials, when a body part was the prime, the target was in fact a building part (e.g., "heart - door"); however, on some trials, unexpected but associatively related targets were presented (e.g., "heart - chest"). Neely examined these expectancy and associative relationships at several different SOAs, and found that facilitation was determined by associative relationships alone at short SOAs (less than 450 msec), but that as SOA increased, facilitation was increasingly determined by expectancy relationships. By 2000 msec, *only* expectancy relationships were facilitated, and unexpected targets, even when associatively related, were *inhibited*.

These data suggested that lexical access involves an initial process of activation of the node associated with a particular phonological code and its associates, independently of "top down," intentional processes. Only subsequently are intention or context observed to have an effect on lexical processes. Results from a number of studies are consistent with this inference

that "early" priming effects (those existing immediately after presentation of the prime) are automatic and are based only on associative relationships. In a careful study of different types of "association" between words, Seidenberg, Waters, Sanders and Langer [1984] found that certain types of association did, but others did *not*, facilitate performance on a naming task. With an SOA (Stimulus Onset Asynchrony) of 500 or 600 msec, and an ISI (Interstimulus Interval, that is, time between offset of first and onset of second stimulus) of 0 msec, syntactic priming (e.g., "men-swear," "whose-planet"), proportion of related pairs (inducing a relatedness set), and backward associations (e.g., "paper-fly," "toy-tinker") were all found to be non-facilitators. By contrast, forward associative primes (e.g., "fly-paper," "tinker-toy," "dog-bark") *did* facilitate performance on the naming task within the range that Neely found to be "pre-intentional." Consistent with the results of Seidenberg et al., den Heyer, Briand, and Dannenbring [1983] report that manipulating the proportion of related pairs (a manipulation reasonably argued to affect intentional processing) affected performance on a lexical decision task at a 1000 msec SOA but not at a 75 msec SOA.

In all, these and other data suggest an important distinction between associative and intentional effects of a prime on a target with respect to time course of processing. (See [Posner and Snyder, 1975; Shiffrin and Schneider, 1977], for discussion of the more general basis for this distinction.) In the remainder of this paper, we will distinguish between these two types of facilitation effects by referring to them as "associative priming" and "intentional context effects" respectively.

Initiation of associative priming effects seems to occur relatively quickly; a number of studies have demonstrated associative priming with as short as a 40 to 50 msec SOA (e.g., [Fischler and Goodman, 1978; Warren, 1972, 1977]). Less well established is duration of associative priming, as most research has focussed on determining whether and how soon priming effects obtain. In one of the few studies directed towards examining the time course of decay of automatic priming effects, Warren [1972] suggested that "activation" persists for about 40 sec post-stimulus presentation (see also [Loess and Waugh, 1967]). However, while facilitation effects do (or can) seem to persist for a relatively long period of time, it is not clear if those effects are due to associative priming or rather to intentional maintenance of that initially automatic effect. As noted above, Neely [1977] found that associative priming occurs at least as early as 250 msec, while intentional context effects initiate considerably later (at about 750 msec, in his study). It is at least possible, then, that priming effects measured after about 750 msec reflect intentional (though not necessarily conscious) rather than automatic processing. Note, however, that the temporal demarcation dividing automatic associative priming from intentional effects is more likely a function of the *type* of processing going on than an *absolute* temporal parameter.

2.2 Paradigm-based Issues In the Examination of Lexical Processing

Most priming studies have used a paradigm in which subjects are presented with isolated pairs or triplets of words, some subset of which are related associatively. It seems quite likely that, whether or not they are told to look for relationships between words in a pair, subjects would adopt such an approach as a strategy on the assumption that there must be some reason the experimenter isolated these two words together (see [de Groot, Thomassen, and Hudson, 1982; Swinney, 1981], for similar arguments). Also, since stimulus inputs stop after the target item in these paradigms, there is no external stimulus to disrupt or inhibit processing of the target word as there is in normal conversational speech. Thus, the pair or triplet paradigms provide a perfect setting for controlled or intentional effects to take place. Thus, the "pair" paradigm does not seem particularly well suited to studying the time course of automatic associative priming effects in normal, fluent language; while it may accurately reflect initiation of priming, it cannot guarantee separation of intentional from automatic effects.

In a recent study designed to examine this issue, Prather, Pasquotto, Seimen, and Lawson [1986] used a "continuous list" rather than a "pair" paradigm to study the time course of automatic priming. A long list (216) of words was presented one at a time on a video monitor; SOA between words varied from 500 to 2250 msec, and the ratio of associatively related to unrelated words occurring sequentially in the list was 14 to 166 (8.4%). The subjects' task was to read each word as quickly as possible (a "naming" task). As soon as the subject began to say the word, it was removed from the screen and the SOA initiated. The assumption in using this paradigm was that reading lists rather than isolated pairs of words would minimize using a strategy of looking for relationships between words, particularly when very few of the words were related. Also, because words were presented continuously, subjects could not choose (consciously or unconsciously) to delay their response to a target. In all, it is a task that may give a better look at automatic processing effects than does the pair paradigm.

Using this paradigm, Prather et al. found a small, (non-significant) facilitation for forward-associated items at 500 msec. That effect increased and was largest (and significant) at 700 msec, then gradually decreased to the point where there was no significant effect by 2250 msec (see Figure 11.1). These data are taken as evidence that automatic effects diminish passively over time and, at least in this continuous processing task, are gone within approximately 2 seconds. It is important to note that Prather et al. did not find immediate "early" priming using the continuous list paradigm. Given that immediate priming is virtually always found when using the pair paradigm on these materials, it seems reasonable to assume that lack of an immediate effect in this study has to do with the paradigm used. One possible explanation of the differ-

ence is that continuous processing may encompass both the automatic facilitation seen in the pair paradigm *and* an inhibition effect. Processing of a lexical code may yield activation of associates and inhibition of non-associates of words in the list. Together, these might well produce a narrowing of the "window" (change in onset of both the beginning and end points) during which automatic associative priming takes place.

With respect to the general temporal course of processing, however, the basic conclusion is similar whether using continuous or pair presentations: there seems to be a short period of time during which lexical priming is automatic, in the sense of being passive and unintentional, followed rapidly by availability of lexical processes of intentional or controlled processes. When processing is rapid and continuous (a situation that discourages specific intentional strategies between words), automatic associative priming effects are largest roughly 700 msec after occurrence of the word and are virtually gone within 2 sec. When intentional processes are allowed or encouraged, both the initial rise time and the decay of associate priming appears to be somewhat earlier (see, e.g., [Neely, 1977; Simpson, 1984]).

To summarize the evidence from studies of temporal and paradigmatic issues in examinations of lexical processing, it can be seen that data from lexical priming studies suggest that there exists a temporal "window" of autonomous lexical processing during which the data base connected to a phonological/orthographic code and its close associates are activated. This temporal window seems to be relatively narrow, where the specific beginning and end points of the window vary with specific task and materials conditions. However, the important point is that processing of information in that window appears to be immune from intentional effects.

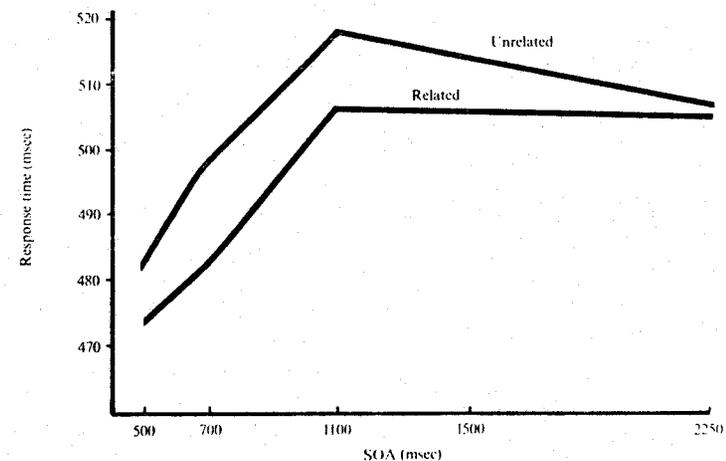


Figure 11.1 Mean reaction time for primed words (Related) and their unprimed controls (Unrelated) when presented in continuous list format at variable SOAs.

3 Accessing Ambiguous Words in the Lexicon: Ambiguity Resolution

Having discussed some aspects of the temporal and paradigmatic boundaries on lexical processing, we now turn to empirical studies examining the question of whether or not prior-contexts (particularly semantic and syntactic contexts) can *restrict* lexical access for an ambiguous word. In particular, we consider whether or not the early "automatic" period in lexical processing is truly autonomous. It is important in this regard to briefly consider the issue of precisely what *types* of effects can be taken as support for and against a modularist position. The lexical modularist position holds that *non-lexical* processes will not directly influence *lexical* access. It is important to note that associative priming (priming between associated lexical items) is actually a within-lexical effect. Consequently, a "context" that is associatively related and (temporally near) an ambiguous word could act on processing of that ambiguous word *within* the lexical module. This type of effect cannot be taken as evidence for interaction *between* modules (see [Fodor, 1983; Seidenberg et al., 1982; Tanenhaus, Carlson and Seidenberg, 1984]). Consequently, to claim strong evidence against a lexical modularist position and in favor of some alternative type of interactive model, either a distant context greater than 200 or so msec) or a non-associative context would have to be shown to restrict access. With that in mind, we review studies that examine the effects of context on ambiguity processing.

3.1 Context Effects and Ambiguous Words

A number of different paradigms have been used to examine the effects of context on ambiguous words, notably: phoneme monitoring, word monitoring, sentence decision, Rapid Serial Visual Presentation (RSVP), modified Stroop (interference) tasks, and various priming paradigms, both within modality (e.g., [Schvaneveldt et al., 1976]) and across modalities (e.g., [Swinney, 1979; Tanenhaus, Leiman, and Seidenberg, 1979]). Only certain of these tasks, however, are of relevance in examining the issue of restricted access. Studies demonstrating that ambiguous words place an increased load on processing based on tasks such as phoneme or word monitoring (e.g., [Foss and Jenkins, 1973]) offer only indirect evidence for multiple access of meanings. If context restricts access, the prediction is not just that there will be a decrease in processing load, but more specifically, that the contextually inappropriate meaning will not be present. It could, for example, be the case that context will reduce processing load in such studies by allowing faster access to one or both meanings, but will *not* restrict access to just one meaning (see, e.g., [Onifer and Swinney, 1981]). In short, processing load offers suggestive evidence, but only tasks sensitive to which meanings have been accessed offer convincing evidence for or

against *restricted* access. Paradigms using priming and interference effects *do* allow for examination of which meanings have been accessed. Consequently, in the following review, we will restrict consideration to studies that have used priming or interference paradigms.

Studies are also differentiated with respect to whether the ambiguous target word is presented in the context of a pair or triplet of isolated words or rather in the context of a sentence. This difference is an important one. As noted earlier, a paradigm in which words are presented as isolated pairs allows and even encourages looking for relationships between words, which seems likely to increase the effects of a context over those effects normally available in discourse processing. Sentence contexts, on the other hand, are more likely to disallow such intentional processing. While all paradigms allow examination of some type of context effects, they cannot be compared directly, nor will results using one type of context necessarily predict effects observed when using the other. Consequently, data from pair and sentence studies is considered separately, even though this paper is intended to focus primarily on ambiguity resolution in normal language comprehension.

3.2 Word Studies

As elaborated earlier, there have been a large number of studies that use word pairs to examine priming effects. However, only a few have systematically examined the effects of a prior associative prime on access of *ambiguous* words. When evaluated in terms of the time course of lexical processing, evidence from those studies is consistent with claims that a prior priming context does not restrict access. This evidence therefore supports claims for autonomy of the lexical access routine.

Simpson [1981] presented ambiguous *words* in pairs with an ambiguous word as the prime and one of its two meanings as the target (e.g., "bank-money," "bank-river"); the control condition was an unrelated ambiguity-target pair (e.g., "calf-money"). While there was *no* semantic context in this experiment, Simpson elected non-equibias ambiguous words, that is, words that had an inherent bias towards one meaning (for example, the "money" meaning of bank is much more frequent than the "river" meaning). The prime word was presented until the subject made a lexical decision, and then the target immediately followed (0 msec ISI). The results showed significant facilitation *on* a lexical decision task for the target word related to the dominant meaning of the ambiguity ("bank-money"), but no facilitation for the subordinate meaning ("bank-river"), relative to the relevant control pair. Simpson argued from these data that meanings of ambiguous words are retrieved in order, and that the dominant meaning is retrieved first and exclusively if there is no contradictory context. In a more recent study, however, Simpson and Burgess [1985] varied the delay between prime and target, and found that while the dominant mean-

ing is facilitated at shorter SOAs (16 msec), both meanings are facilitated at a slightly longer SOA (100 msec); subsequently, activation of the subordinate meaning diminished while that of the dominant meaning increased. It seems, then, that the rise time for the dominant meaning is more rapid and that that meaning *is* selected for continued attention, while the subordinate meaning is accessed but somewhat more slowly, at least when there is no prior semantic bias.

With respect to the effects of a semantic bias, Schvaneveldt, Meyer, and Becker [1976] offer a very comprehensive examination of the effects of context on ambiguity resolution. In their study, Schvaneveldt et al. presented subjects with sequential word triplets. The subjects made a lexical decision on each word, and as soon as the decision was made for one, the next word in the triplet was presented. Experimental triplets included an ambiguous word as the first or second word in the triplet; contexts included a concordant bias, that is, a word that biased the ambiguity towards the target word (e.g., "save-bank-money"), a bias discordant with that word (e.g., "river-bank-money") and a "null" bias (e.g., "river-date-money"). Basically, the results showed that when there was a concordant context, lexical decision on the target word (here, "money") was facilitated relative to the null-context condition; when the context was discordant, the target word was not facilitated.

These data are taken as consistent with an inference that when there is a strong prior biasing context, only one meaning of an ambiguous word, the contextually appropriate meaning, is accessed. However, given Simpson's results with longer SOAs, the conclusion in favor of restricted access must be viewed with some caution. It could as well be argued that the amount of time between words in each triplet was of sufficient duration to allow intentional processing, and therefore selection from rather than restriction of accessed meanings. While the procedure in Schvaneveldt et al. included presenting each word *as* soon as a decision had been made on the prior word, average decision time was 500-600 msec. SOA, then, would be in that range, which clearly is long enough to allow intentional suppression of the non-biased meaning of the ambiguous word (as demonstrated in [Neely, 1977; Simpson and Burgess, 1985]) and therefore to argue against a claim that the biasing context acted to restrict access.

In contrast to the Schvaneveldt et al. results, Oden and Spira [1983] offer evidence consistent with an inference of full elaboration of ambiguous words presented in triplets, regardless of context condition-but again at an ISI sufficient to allow intentional processing. In one condition of their study, subjects heard two biasing words and then an ambiguous word, each presented at "about" a rate of one every 3/4 second; there was a subsequent 500 msec ISI, and the visual presentation of a colored word related to one meaning of the ambiguity. The subjects' task, a variation of the Stroop interference task, was to name the color of the word.

Using this paradigm, Oden and Spira found that *both* meanings of the ambiguous word showed an interference effect, suggesting that both are accessed but that the meaning consonant with the contextual bias consistently showed a larger effect than the other meaning. In short, context affected *degree* of activation but did not restrict activation to just one meaning. As with Schvaneveldt et al., the length of the delay between context and ambiguous words, and between ambiguous words and the colored word to be read, raises problems with respect to inferences about automatic access processes. It is at least possible (if not plausible) that within that 500 msec ISI, subjects intentionally initiated a full elaboration process; subjects might, for example, have noticed that the colored word was sometimes related to the "unexpected" meaning of a word, and subsequently started to elaborate both meanings. It should also be noted that direct comparisons, particularly with respect to time course comparisons, cannot be made between priming tasks and interference tasks. These tasks are quite different, their time courses have not been directly compared, and it cannot be assumed that they would be similar either in rate or direction of changes over time.

Overall, while data from studies using isolated words are consistent with a claim that all meanings of ambiguous words are elaborated at least briefly, they are far from definitive. A particular problem with these studies is failure to examine a broad range of inter-word delays in order to characterize the time course of lexical access in the presence of a context (as Simpson has done for ambiguous words in the absence of context).

3.3 Sentence Studies

The majority of studies using priming to examine the effects of sentence context on the processing of ambiguous words are, as with the word studies, consistent with an exhaustive access position. These include the study of Tanenhaus et al. [1979] using syntactic contexts, and Swinney's studies using semantic contexts with equibias and non-equibias ambiguities ([Onifer and Swinney, 1981; Swinney, 1979; Swinney, Onifer, Prather, and Hirshkowitz, 1979]). While these studies differed in types of ambiguity (noun-verb vs. noun-noun), dominance of meanings, and task (naming vs. lexical decision, during-sentence vs. end-of-sentence), they were consistent in finding that regardless of direction of prior biasing context, priming was demonstrated for words related to *both* meanings of the target ambiguity. As with the word studies, examination of the time course showed that both meanings were available only for a delimited period. Unlike the word studies (where both meanings seemed to become available only after a short delay), both meanings were facilitated when examined *immediately* following the target word.

Tanenhaus et al. [1979] presented subjects with sentences biased towards one meaning of ambiguities that had independent noun-verb meanings (e.g., "watch"). The ambiguity was always the last word in the sentence, and was followed at a 0, 200, or 600 msec ISI by a visually-presented word related to one meaning of the sentence (e.g., "look"). Regardless of contextual bias, words related to both meanings were facilitated on the naming task at the 0 msec delay; by 200 msec, however, only the word related to the contextually appropriate meaning was facilitated. Swinney [1979] similarly presented subjects with sentences biased towards one meaning of an ambiguous word. In this study, however, the ambiguous words were selected to have 2 noun meanings of equal strength, and the task involved lexical decision on a visually-presented word that was presented while the auditory sentence was ongoing. Again, the visual word was presented with either a 0 msec ISI or 3 syllable (approximately 800 msec) ISI following the ambiguous word. And again, words related to both meanings were facilitated downstream.

While supportive of a claim that all meanings of an ambiguous word are accessed regardless of prior bias, one criticism raised about these studies has been that the apparent facilitation of words related to both meanings may be an artifact of using equibias ambiguities (see [Hogaboam and Perfetti, 1975]). It may be that about half the subjects access one meaning first, and half the other meaning first, resulting in both meanings being facilitated *on average* but not by *all* subjects. To test for that possibility, Onifer and Swinney [1981] essentially replicated Swinney [1979] using ambiguous words that showed a strong bias towards one meaning. The original results replicated; words related to both the dominant and subordinate meaning were facilitated, even in the case where the context was biased toward the dominant meaning. Recently, this study itself was replicated by Gildea [1984].

More recently, Oden and Spira [1983] have questioned the claim that context has no effect on the lexical access process. They argue that, while there is convincing evidence that context does *not* restrict access, their research suggests that context does affect *degree* of activation of an ambiguity's multiple meanings. In that sense, context can affect the access process. In the study they offer as support for this position, subjects listened to sentences that included a bias towards one meaning of equibias ambiguities, with the ambiguous word as the last one in the sentence; 500 msec after the end of that sentence, a Stroop, color naming task was performed on a word that was related to one meaning of the sentence. Interference effects were obtained for both meanings, but were significantly greater for the contextually appropriate meaning, i.e., context affected the *degree* of activation of the multiple "word senses" of an ambiguity. Oden and Spira acknowledge, however, a clear problem with their interpretation: because they test 500 msec after the ambiguity, they may be examining a post-lexical *decision* process rather than lexical access. In short, their data are consistent with the preceding studies in suggesting that context has an effect on

relative activation of meanings of an ambiguity shortly after access. However, their data do not speak to the issue of whether or not that effect is taking place within the "window" of automatic lexical processing.

To summarize briefly, those studies which have examined the time course of lexical processing within what we have called the "window" of automaticity support an inference that all meanings of an ambiguous word are made available in this access process, and that one meaning is selected relatively quickly thereafter. The selected meaning is the one consistent with prior bias or, if none is present, with the preference/frequency bias for that word. These data are taken as support for a claim that there is at least a brief period of autonomy, that is, of modularity, during lexical processing.

There is, however, one sentence study in which a prior context *does* seem to restrict access to just one meaning of an ambiguous word. However, for reasons cited earlier, the results of this study also do not contradict claims for autonomy of lexical processing. Rather, they are consistent with the prediction that the one condition under which a context could restrict access would be when the "context" is a lexical associate that, through automatic priming effects, facilitates one meaning of the ambiguity automatically. The effect was obtained by Seidenberg et al. [1982]. Across four studies, they presented subjects with an auditory phrase or sentence that ended with an ambiguous word; then at either a 0 or 200 msec ISI, subjects were presented with a word to read aloud. In one study, a general semantic context was selected (disambiguation with no words or phrases associatively related to the target, e.g., "You should have played the *spade*"). As with other studies reported above, words related to both meanings of ambiguous words were facilitated at 0 msec.

In two other Seidenberg et al. studies, however, the context was selected to be lexical, that is, the context was a strong lexical associate presented within the same phrase as (and therefore temporally near) the target ambiguity (e.g., "Although the *farmer* bought the *straw*"). In that condition, only the contextually-appropriate meaning was facilitated for the noun-noun ambiguities. This result is consistent with the prediction that a temporally close associative context could restrict access to just one meaning as an automatic (within-lexicon) process (see also discussion in [Seidenberg et al., 1982]). This is one of a few studies to use naming rather than lexical decision, and it presents the target word at the end of the sentence which may allow some strategic processing (similar to the argument for "pair" vs. "list" processing presented earlier). Still, the finding is an important one, it is consistent with theoretical predictions, and it warrants further empirical attention.

To summarize, across a number of studies using a variety of paradigms, the data have consistently supported two inferences. (1) Both meanings of an ambiguous word are accessed even in the presence of a biasing context *except* when the context is lexical and temporally near to the target ambiguity. (2) The automatic access routine is temporally delimited to about 200 msec post-am-

biguity. In short, the evidence favors a model in which access is brief but exhaustive except (perhaps) in the case where the context acts within the lexical module. This conclusion is based on the weight of the evidence rather than on any definitive study; as Simpson [1984] points out, the results of most studies are open to question based on methodological issues, and there has been little consistency in materials, tasks, or temporal range explored across studies that would allow the systematic building of a model of the lexicon and its internal processes.

The one remaining serious challenge to this general conclusion comes from a recent set of studies reported by Glucksberg, Kreuz and Rho [1986]. They argue that evidence supporting contextually independent lexical access derived from priming studies (again, arguably the strongest evidence supporting the lexical modularity hypothesis) is confounded due to the possibility of backward priming, and they provide a study that uses a non-word interference paradigm that, so it is claimed, demonstrates contextually directed lexical access. Because this claim is so critical to the issues we have raised above, we will now consider the Glucksberg et al. arguments in some detail.

4 On Backward Priming: Why It Isn't

In what follows, we will first briefly describe the claims made about backward priming and evaluate these claims on logical and empirical grounds. Following that, we will briefly evaluate the Glucksberg et al. evidence and provide some new evidence that bears on their study.

The possibility of something called backward priming has been raised by Glucksberg et al. as an alternative explanation to that typically given to findings of exhaustive access using a cross-modal priming task (CMLP). Glucksberg et al. argue that only the contextually constrained interpretations of an ambiguity are ever accessed when a lexical ambiguity is heard, but that the acoustic form of the ambiguous word remains in "echoic" memory for a period following access of the appropriate interpretation. According to this view, what happens in a CMLP task is that if (when) the contextually inappropriate probe (visual target word) is presented at the end of the ambiguity, that word, too, enters some form of memory. It is here that "backward priming" is argued to take place; the relationship of the (contextually inappropriate) probe word in memory to the *unchosen* meaning of the acoustic form of the ambiguity in echoic memory is somehow noticed and serves to activate this (previously unaccessed) interpretation. This, in turn, serves to make lexical decisions to the contextually inappropriate probe word faster than those to an unrelated control word. In short, the presentation of the inappropriate probe word as a lexical decision target has an effect on the acoustic form of an ambiguous word that

occurred temporally earlier, causing reprocessing of that ambiguity, which results in "access" of the contextually inappropriate meaning of that ambiguity which, in turn, "primes" the decision to the target/probe word itself.

While conceivably true, this seems to be a tortuous explanation of how priming obtains for contextually inappropriate meanings of words in the CMLP task. Consider, for example, that if the contextually inappropriate probe word has to be accessed to begin with for the relationship between itself and the *un-*accessed meaning of the ambiguous word to be noticed, why is there facilitation in making a lexical decision to the probe at all? Surely a lexical decision could have been made in response to the *initial* access of the probe? All of this later processing seems superfluous with regard to actually making a lexical decision to the target. Despite this and other such questions, however, the claim does represent a scenario worth considering and evaluating carefully. In what follows, we deal with a few of what we take to be the major points in the Glucksberg argument, beginning with the concept of backward priming itself.

Backward priming is a concept that has derived from sources independent of the Glucksberg claims, thus strengthening the potential plausibility of their arguments. Most notably, Kiger and Glass [1983] (see also [Koriat, 1981]) demonstrated a relative savings for lexical decisions made in response to visually presented words (stimulus duration of 30 msec) when they were followed immediately (35 msec later) by an associatively related word, as compared to when they were followed by a non-word. The basic question is, is this priming at all?

There is every reason to believe that rather than priming, these researchers have demonstrated a diminution-of-interference effect. After all, asking subjects to make a lexical decision in response to a very briefly reported word, and requiring them to immediately process another word following *that* word (all within 65 msec) is most likely to cause a great deal of interference in lexical decision to that initial word. In fact, the average reaction time to make lexical decisions in experiments 2 and 3 of Kiger and Glass (the only one of their experiments in which a significant experimental-control difference was found) was on the order of 700 msec. This is considerably longer than normal isolated (or "pair") lexical decision RTs (which typically average in the 400-500 msec range). Such long RTs certainly suggest a great deal of interference might be taking place in making the initial decision. The savings discovered for priming associates in this task suggests that the interference lessened somewhat when the following word was associatively related.

The single way in which one might determine if these effects were essentially driven by priming or by inhibition would be to employ a "neutral" prime. Unfortunately, in the one experiment in which a form of "neutral" prime was used by Kiger and Glass (Experiment 1), no significant effects of "backward priming" were obtained. It is probably not accidental that reaction times in this study (their Experiment 1) averaged about 475 msec. Here the "backward"

context followed the target word by a far longer delay than in the other studies, and provided a condition where one might hope to obtain somewhat interference-free reaction times. Thus, it seems to be far from clear that there is any evidence for something that can reasonably be called *backward* priming. It seems much more likely that what we have here is unrelated to priming, backward or forward. What, then, is the basis for believing that backward priming can in fact occur?

In an attempt to determine whether or not we could obtain evidence for backward priming between auditorially and visually presented items, we recently conducted a study involving 48 word-pair associates. In this experiment, the first member of the word pair (or a non-word filler) was heard auditorially and the subject was required to make a lexical decision to that item. The second member of the associate pair or a control word was presented at an ISI of 0, 300, or 600 msec. This was run, with appropriate counterbalancing, on 36 subjects. While null results are normally problematic things to report, here they have interpretive value. As can be seen in Table 1, no priming or interference was found for this cross-modal variation of the Kiger and Glass experiments at any of the SOAs.

Regardless of whether or not one believes that the visual domain effects that Kiger and Glass report represent "backward" priming or lessening of interference effects, it is necessary for the Glucksberg et al. argument that these effects obtain in the cross-modal paradigm. In this attempt to obtain such results (a good faith attempt that we believe was a strong one), we failed to find any evidence that cross-modal processing produces either effect. It might be noted that the lack of an effect in the cross-modal paradigm reinforces the notion that the within-mode paradigm of Kiger and Glass produced interference effects. In the CMLP, visual information cannot physically (peripherally) interfere with the initially presented, auditory item.

Table 1. Lexical decisions in response to the first item of an associated or "control" pair, presented in cross-modal format with variable SOAs.

| Pair Type | SOA (msec) | | |
|---------------|------------|-----|-----|
| | 0 | 300 | 600 |
| Associated | 598 | 609 | 607 |
| Nonassociated | 602 | 604 | 601 |

Hence, at one level, it is not clear that the argument for backward priming made by Glucksberg et al. is an independently motivated one; there appears to be only very suspect evidence to support the notion that backward priming obtains at all, although backward interference effects may indeed occur within a modality.

However, even if there is no independent evidence for backward priming, one must also consider the Glucksberg et al. evidence on its own merits. In their paper, these researchers present clever and interesting experimental work that they take to support the non-modularist position. However, as we attempt to show below, all of their work hangs on slender threads of assumptions that we believe to be incorrect, a belief for which we present some direct empirical support.

Briefly, Glucksberg et al. make the argument that if backward priming exists in the CMLP paradigm, and if they can come up with a form of this paradigm which *prevents* backward priming itself but uses the same general cross-modal approach to studying access of meanings for lexical ambiguities in restricted contexts, then they could determine if the Onifer and Swinney [1981] results were caused by this hypothesized backward priming effect. What they devised was an experiment using the Onifer and Swinney auditory materials (lexical ambiguities in biasing sentential contexts), but with non-word probes rather than word probes. These non-word probes were misspellings of *words* that, were they not misspelled, would have formed good associate probes for the various interpretations of the lexical ambiguities. For example, CONDRY and CONBISHUN were non-word probes for the ambiguity STATE; here, CONDRY is similar to COUNTRY and CONBISHUN is similar to CONDITION, the two interpretations of the ambiguous word STATE.

It was argued that because these were non-words, they could not be accessed lexically and therefore could not provide lexical priming—either forward or backward. Thus, it was argued that if the same type of effect was found for CONDRY but not for CONBISHUN when they were presented immediately following occurrence of the ambiguity STATE in a sentence that was contextually biased toward the "country" interpretation, then this would be evidence that *only* the contextually relevant interpretation of a lexical ambiguity is accessed, that lexical access is neither modular nor contextually independent, and that the effects obtained by Onifer and Swinney were due to backward priming (while these new effects *couldn't* be due to priming). They argue that their data support such an interpretation. However, it appears to us that this is a premature conclusion, and to substantiate our claim, we will briefly present some data that bear directly on the major underlying assumption of Glucksberg's research.

The entire premise of the Glucksberg et al. experiment hinges on the notion that non-words such as CONBISHUN cannot or will not prime (facilitate) the processing of related words. If they do, then whatever results Glucksberg et

al. obtain with their task, they *cannot* be results that are impervious to "backward priming" effects (if such exist at all). And, of course, that was the point of their experiment to run something that couldn't be subject to priming. In their paper, they mention that a pilot study was run examining the effects of the non-word probes on lexical decision to the lexical ambiguities used in the sentence study, and no facilitory effects for the related items were found, relative to the control non-words. This is an absolutely critical finding to the entire premise of the experiment, and it was important to run it as they did. However, we have been examining the effects of phonetic and orthographic priming and interference in our lab for some time as have many others (see, e.g., [Seidenberg and Tanenhaus, 1979]), and have had pilot words which indicated to us that non-words were at times sufficiently close to "real" words, either phonetically or orthographically, to prime decisions to other "real" words.

It certainly seems quite possible that, for example, anyone reading CON-BISHUN might well make a word of it (i.e., might access possible words it could be), regardless of the fact that it was, in a formal sense, a non-word. In fact that is, by and large, the task that we all have as listeners and readers. We make sense of something even if it is misspelled or mispronounced-and such processing appears to be very automatized in adults. Hence, we were exceptionally curious about the lack of an effect of related non-words on the ambiguities. Consequently we ran an experiment using the Glucksberg et al. material in a straightforward priming task.

In this experiment, each of Glucksberg's experimental non-words and their respective control non-words were presented visually immediately prior to and also at an ISI of 300 cosecs prior to the ambiguous word, which also appeared visually. The material set comprised exactly the material set used in the Glucksberg experiment, along with an equal number of non-word target items (48) also preceded by non-words. The design constituted a 4 (between) x 2 (within) subjects design, and was run on 46 subjects. As can be seen in Table 2, significant effects of "related" *non-word* facilitation were obtained, compared to the appropriate Glucksberg "control" non-words. In short, these non-words do prime. Given that non-words prime in this condition, the results of Glucksberg et al. are not explainable in terms of their non-words preventing backwards priming, and thus their experiment neither demonstrates that the Onifer and Swinney results were likely due to backward priming nor that only a single contextually relevant meaning of the ambiguity was accessed in their study.

What, then, do the Glucksberg et al. results indicate? While it is difficult to know precisely without further work, it seems reasonable to focus on the fact that we know little about the nature and time course of interference effects that Glucksberg relies on (RTs to the Glucksberg non-words Mere interference RTs). It may well be that the processing of the Glucksberg non-words was sufficiently slow that the interference effects were a result of processing that

Table 2. Lexical decision reaction times in response to ambiguous words following presentation of the Glucksberg et al. non-word stimuli.

| Non-word Type | ISI(msec) | |
|-------------------|-----------|-----|
| | 0 | 300 |
| Related non-words | 568 | 592 |
| Control non-words | 548 | 558 |

occurred entirely outside of the "window" of lexical access. Or perhaps they reflect nothing of the ambiguity at all, and only reflect processing of the prior biasing context. Whatever the answer, their data do not appear to be problematic for claims of autonomous lexical access.

5 The Lexicon: An Interactive Box with an Autonomous Component Process

In our introductory remarks, we suggested that the test of whether or not lexical priming is modular involves a determination of whether or not non-lexical information affects processing *within* the lexicon. If so, then the process is not encapsulated, that is, is not independent and therefore not an autonomous module. Available evidence suggests that there is a narrow temporal window during lexical processing within which prior semantic and syntactic contexts cannot affect the access routine. Based on studies of lexical ambiguities, the only prior context that has been shown to affect access of an ambiguity within the first few hundred milliseconds after its auditory or visual presentation has been a prior lexical associate presented temporally close to the ambiguity; all other contexts seem to affect selection of the appropriate meaning only after context-independent processing.

Is the lexicon an autonomous module? It certainly appears to be, although it may be that the definitive work in the area is yet to be done. Certainly, as we better understand the nature of automatic and interactive processing routines, we will be better able to address the basic question. It seems sufficient at present to say that, given what we know now, aspects of lexical processing are autonomous, and lexical ambiguity resolution involves a process of exhaustive retrieval of information stored in the data base for a word, followed by post-access, context-determined processing which normally results in a single conscious meaning for an ambiguous word in the normal course of sentence comprehension.

Acknowledgments

Preparation of this manuscript was supported in part by an NIMH grant and by a fellowship from the FIDIA Corporation to the first author, and by NIH Grant NS21806 to the second author. We would like to thank an anonymous reviewer for comments helpful to our revision of this manuscript, particularly with respect to an interpretation of work by Schvaneveldt et al. consistent with our model of the lexicon.

References

- Bowles, N., and L. Poon. 1985. Effects of priming in word retrieval. *Journal of Experimental Psychology: Learning, Memory and Cognition* 11:272-285.
- den Heyer, K., K. Briand, and G. Dannenbring. 1983. Strategic factors in a lexical-decision task: Evidence for automatic and attention-driven processes. *Memory and Cognition* 11:374-381.
- Fodor, J. 1983. *Modularity of Mind*. Cambridge: Bradford-MIT Press.
- Fischler, I. and G. Goodman. 1978. Latency of associative activation in memory. *Journal of Experimental Psychology: Human Perception and Performance* 4:455-470.
- Foss, D. and C. Jenkins. 1973. Some effects of context on the comprehension of ambiguous sentences. *Journal of Verbal Learning and Verbal Behavior* 12:577-589.
- Gildea, P. 1984. Unpublished dissertation, Princeton University.
- Glucksberg, S., R. Kreuz and S. Rho. 1986. Context can restrain lexical access: Implications for models of language comprehension. *Journal of Experimental Psychology Learning, Memory and Cognition* 12:323-335.
- Hogaboam, T. and C. Perfetti. 1975. Lexical ambiguity and sentence comprehension. *Journal of Verbal Learning and Verbal Behavior* 14:265-274.
- Kiger, J. and A. Glass. 1983. The facilitation of lexical decisions by a prime occurring after the target. *Memory and Cognition* 11:356-365.
- Koriat, A. 1981. Semantic facilitation in lexical decisions as a function of prime-target association. *Memory and Cognition* 9:587-598.
- Meyer, D. and R. Schvaneveldt. 1971. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology* 90:227-234.
- Morton, J. 1982. Disintegrating the lexicon: An information processing approach. In J. Mehler, E. Walker and M. Garrett (eds.), *Perspectives on Mental Representation*. Hillsdale: Erlbaum
- Neely, J. 1977. Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited capacity attention. *Journal of Experimental Psychology: General* 106:226-254.
- Onifer, W., and D. Swinney. 1981. Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. *Memory and Cognition* 9:225-236.
- Oden, G. and J. Spira. 1983. Influence of context on the activation and selection of ambiguous word senses. *Quarterly Journal of Experimental Psychology* 35a:51-64.
- Prather, P., J. Pasquotto, A. Seimen and D. Lawson. 1986. *The Time Course of Lexical De-activation*. Manuscript, University of Illinois at Chicago.
- Posner, M. and C. Snyder. 1975. Facilitation and inhibition in the processing of signals. In P. Rabbit and E. Dornic (eds.), *Attention and Performance V*. London: Academic Press.
- Schvaneveldt, R., D. Meyer and C. Becker. 1976. Lexical ambiguity, semantic context, and visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance* 2:243-250.
- Shiffrin, R., and W. Schneider. 1977. Controlled and automatic human information processing II: Perceptual learning, automatic attending, and a general theory. *Psychological Review* 84:127-180.
- Simpson, G. 1981. Meaning, dominance, and semantic context in the processing of lexical ambiguity. *Journal of Verbal Learning and Verbal Behavior* 20:120-136.
- Simpson, G. and C. Burgess. 1985. Activation and selection processes in the recognition of ambiguous words. *Journal of Experimental Psychology: Human Perception and Performance* 11(1):28-39.
- Seidenberg, M., M. Tanenhaus, J. Leiman, and M. Bienkowski. 1982. Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology* 14:489-537.
- Seidenberg, M., G. Waters, M. Sanders, and P. Langer. 1984. Pre- and post lexical loci of contextual effects on word recognition. *Memory and Cognition* 12:315-328.
- Seidenberg, M. and M. Tanenhaus. 1979. Orthographic effects in rhyme and monitoring. *Journal of Experimental Psychology: Human Learning and Memory* 5:546-554.
- Swinney, D. 1979. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior* 18:645-660.
- Swinney, D., W. Onifer, P. Prather and M. Hirshkowitz. 1979. Semantic facilitation across sensory modalities in the processing of individual words and sentences. *Memory and Cognition* 7:159-165.
- Swinney, D. 1981. Lexical processing during sentence comprehension: Effects of higher order constraints and implications for representation. In Meyers,

- Laver, and Anderson (eds.), *The Cognitive Representation of Speech*. Amsterdam: North-Holland.
- Tanenhaus, M., G. Carlson and M. Seidenberg. 1984. Do listeners compute linguistic representations? In Zwicky, Karttunen, and Dauty (eds.), *Natural Language Parsing: Psycholinguistic, Theoretical, and Computational Perspectives*. Cambridge: Cambridge University Press.
- Tanenhaus, M., J. Leiman, and M. Seidenberg. 1979. Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior* 18:427-441.
- Warren, R. 1972. Stimulus encoding and memory. *Journal of Experimental Psychology* 94:90-100.
- Warren, R. 1977. Time and the spread of activation in memory. *Journal of Experimental Psychology: Human Learning and Memory* 3:458-466.