

# Meaning resolution processes for words: A parallel independent model

LESLIE C. TWILLEY and PETER DIXON  
*University of Alberta, Edmonton, Alberta, Canada*

Lexical ambiguity research over the last two decades is reviewed, with a focus on how that literature applies to understanding the resolution of meaning for words. Early models of ambiguity processing dealt almost exclusively with the time course of the effects of context on lexical access, in order to address the issue of modularity of lexical access. Newer models of ambiguity processing accommodate recent findings of early context effects that are contingent on both strength of context and meaning frequency. The most important contribution of these newer models of ambiguity processing is not to the modularity debate, but to investigation of the range of parameters affecting the entire meaning resolution process, including meaning access as well as the integration of meanings into context. As an example of this approach, we describe a simple quantitative model of meaning resolution that subsumes many other models as parametric variations.

One of the most significant challenges to perception and comprehension is indeterminacy. We are continually required to disambiguate stimuli with uncertain identities, using whatever environmental and experiential context is available. Stimulus context has obvious and profound effects on information processing, from object recognition to language understanding to social interaction, and the influence of context on processing is a classic topic of psychological research. Of critical concern are questions about the mechanisms and time course of contextual influence and the insight these questions and their answers provide into representation and process.

In the case of language, ambiguity occurs at all levels of linguistic communication, from perception of individual letters or sounds to discourse processing. The present article is concerned with contextual influence on lexical, or word-level, processing. There is abundant evidence that context affects lexical processing in general. For example, Stanovich and West (1983) found that lexical decision and naming times for sentence-ending words that were congruent with preceding words (e.g., *The puppy chewed on the bone*) were faster than lexical decision and naming times for words ending neutral sentences (e.g., *They said it was the bone*). Similarly, Meyer and Schvaneveldt (1971) showed that lexical decisions to words (e.g., *nurse*) were faster when preceded by a related word (e.g., *doctor*) than when preceded by an unrelated word (e.g., *bread*). The

contextual influence in both of these cases was facilitatory, with congruent sentences or related words making identification of the target word easier. Conversely, incongruent contexts can have inhibitory effects. Stanovich and West, for example, showed that lexical decisions to words ending incongruous sentences (e.g., *He smiled and sat down at the bone*) were slower than lexical decisions to words ending neutral sentences. These (and hundreds more) studies demonstrate that contexts, in forms such as associatively related words or biasing sentences, do affect the processing of subsequent words. The crucial questions concern the nature of the mechanism that generates such effects and when and how the mechanism operates.

Lexical ambiguity enjoys a considerable history as an arena for the study of such questions. Lexical ambiguity occurs when an utterance has more than one possible meaning because it contains a homograph (a word with more than one meaning). A sentence such as *The men stood beside the pitcher* is lexically ambiguous, because the word *pitcher* is a homograph and both of its meanings result in plausible interpretations of the sentence. In the case of structural ambiguity, an utterance has two or more possible interpretations owing to ambiguity in assignment of grammatical roles. So, for example, the sentence *Visiting relatives can be a nuisance* has two interpretations, depending on the role given to *visiting* (relatives that someone is visiting or relatives that are visiting someone). Although the issues and mechanisms we discuss here are relevant to both types of word-level ambiguity, we are concerned primarily with lexical ambiguity.

Homographs are well suited to the study of the role of context in lexical processing, because the different interpretations of a homograph are perfectly controlled with respect to variables such as word frequency, length, and number of syllables: Orthographic (and almost always phonological) representations are identical for both interpretations of a homograph; only the meaning varies. Any

---

This research was supported by a Natural Sciences and Engineering Council of Canada (NSERC) doctoral scholarship and a Killam doctoral scholarship to L.T. and by an NSERC research grant to P.D. The authors thank Gay Bisanz, Norman Brown, Bruce Derwing, Jeff Bisanz, Mike Masson, Greg Simpson, Stephen Paul, and anonymous reviewers for comments and helpful discussions on earlier versions of the paper. Correspondence concerning this article should be sent to L. C. Twilley or P. Dixon, Department of Psychology, University of Alberta, Edmonton, AB, T6G 2E9 Canada (e-mail: ltwilley@telusplanet.net or pdixon@psych.ualberta.ca).

variation in the meaning assigned to a homograph must be ascribed purely to the context in which the homograph appears. Although not all words have more than one distinct meaning, almost all English words do, at least, have different senses (see Anderson & Nagy, 1991; Britton, 1978), and it is our contention that understanding the processes involved in homograph meaning resolution will further our understanding of the resolution of word meaning and language comprehension in general.

The extensive literature on context effects in lexical ambiguity can seem convoluted and contradictory; however, we contend that the application of a simple quantitative model, the independent activation model, provides a framework for a great deal of the extant data on ambiguity resolution and incorporates several common approaches to ambiguity resolution as special cases. It is our hope that this type of model will prove useful in understanding the more general issue of the influence of context on comprehension in many domains.

The paper is organized as follows. First, we describe the lexical modularity debate, a theoretical issue that has dominated research on lexical ambiguity. Our analysis is that theoretical discussion has often confused the time course of meaning resolution with issues related to cognitive modularity, resulting in a lack of consideration of an important class of models dealing with context effects. Second, we review several common views of ambiguity resolution and some of the evidence in their favor. Third, we present a new approach, the independent activation model, that would seem to account for a wide range of different results. Finally, we discuss the implications of this approach for understanding the time course of meaning resolution, the fate of unselected meanings, the development of meaning resolution processes, and individual differences in meaning access and integration.

### THE LEXICAL MODULARITY DEBATE

Research on lexical ambiguity began to appear with regularity in the psychological literature in the 1970s and has continued to be an important topic in psycholinguistics (for reviews, see Bubka & Gorfein, 1989; Rayner, Pacht, & Duffy, 1994; Simpson, 1984, 1994). One of the central questions has been, What is the time course of the influence of contextual information on low-level perceptual processing during completion of cognitive tasks? There is little disagreement that contextual influences on the final outcome of processing are abundant in many tasks; when such influences occur is the prevalent empirical question. Early and rapid effects of context have been taken to imply that low-level lexical processing, such as lexical access, occurs in parallel with higher level cognitive operations, such as the integration of accessed information with the surrounding context. The lack of such effects has been taken as evidence for the modularity of lexical processing.

In our view, a clear discussion of this issue requires a distinction between two types of processes: *access* pro-

cesses and *integration* processes (Marslen-Wilson, 1989). For written and spoken communications, access refers to the process of mapping incoming perceptual information onto lexical representations; integration entails combining lexical representations with relevant internal information. Such distinctions are implicitly or explicitly assumed in most models of lexical processing (e.g., Forster, 1981; Seidenberg, 1985; Stanovich, 1991). Access is generally conceived of as those processes that result in lexical representations being activated when a word is presented either visually or auditorily. Operationally, access has often been defined as the point at which a word can be pronounced or distinguished from a nonword (hence the use of the naming and lexical decision tasks, respectively). However, we think it is important to distinguish access processes from those mechanisms that operate on or distinguish among word meanings; we use the term *access* to refer solely to those processes that relate the perceptual input to the appropriate lexical representations. In contrast, integration processes are those that affect the activation of word meanings on the basis of other relevant information, such as surrounding words, prior knowledge, and so on. When multiple interpretations are possible, as is the case for homographs, integration procedures may be involved in selecting an appropriate meaning of the homograph by using contextual information. The value of distinguishing access from integration is that it allows us to be clear as to how different sources of information are hypothesized to be used in the course of word recognition in a context. In particular, access refers to the use of perceptual information from the environment, whereas integration refers to the use of internal information arising from mental processing of the context. We use the term *meaning resolution* to refer to the entire constellation of processes that are used to activate and integrate word meanings.

Modular and nonmodular views of meaning resolution differ primarily in terms of the presumed relationship between meaning access and meaning integration. Modular theorists (e.g., Fodor, 1983; Forster, 1981; Seidenberg, 1985) argue that many perceptual and cognitive processes are carried out by separate and distinct structures. Each of these structures, or modules, has limited interactions with other modules. In particular, each module is assumed to be cognitively impenetrable, in the sense that it is uninformed as to the details of computations within other modules. In word recognition, it is often argued that a mental structure containing information about words, the *lexicon*, constitutes a module. Because modules are cognitively impenetrable, access proceeds solely on the basis of perceptual input and information internal to the lexicon, without input from syntactic, semantic, or pragmatic levels. Thus, modular models of meaning resolution have at their core the assumption that access processes are independent of integration processes.

On the other hand, nonmodular theorists (e.g., Becker, 1980; McClelland, 1987; Neely, 1976, 1977; Posner & Snyder, 1975) contend that context can affect lexical ac-

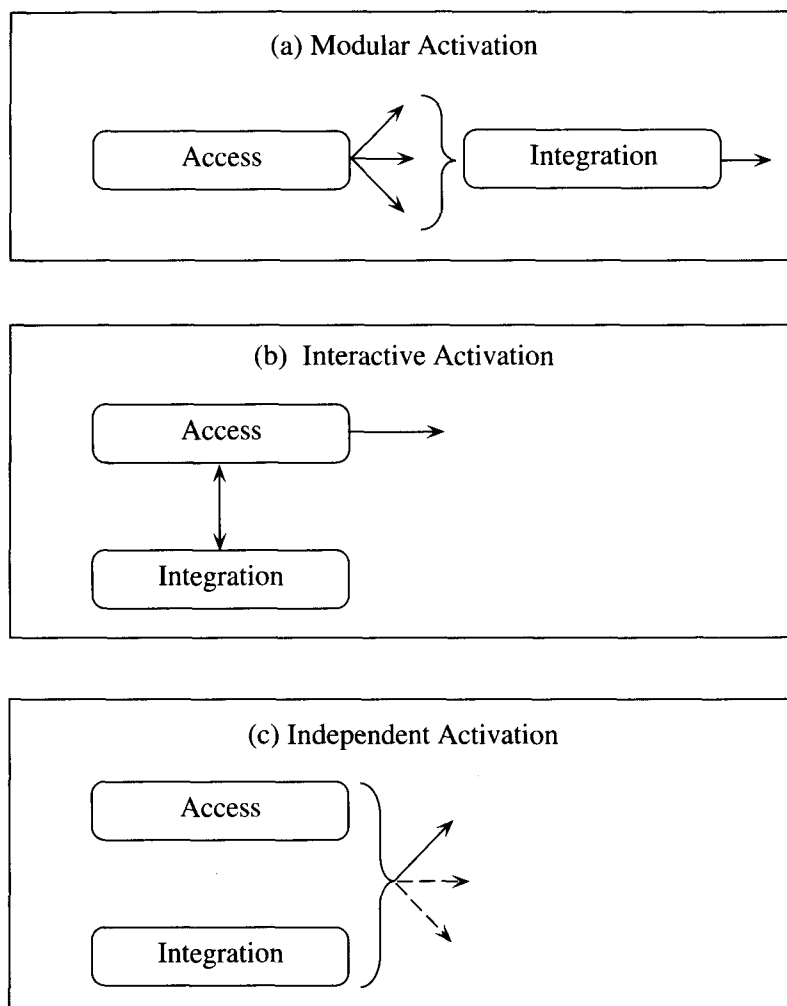
cess; access and integration interact and are not independent. In the prototypical nonmodular approach (e.g., McClelland's [1987] interactive activation model), information flow is multidirectional, with both feedforward and feedback relations between the access and the integration levels of processing. In nonmodular theories, it is generally assumed that this kind of interaction occurs throughout most of the time course of lexical processing and that, in principle, various levels of contextual constraint could affect the activation of word meanings at very early stages of lexical processing.

In this discussion, we are using the terms *interactive* and *independent* to refer specifically to the presence or absence of these immediate feedback loops: In our terms, an interactive model is one in which information flows rapidly back and forth between two or more processes, whereas an independent model is one in which information flows primarily from one process to another in one direction. There is a larger sense in which all models of word recognition are interactive, because contextual information and perceptual information are combined in order to reach some final meaning resolution. We believe that our more restrictive usage is more in keeping with the vocabulary typically adopted in the discussion of parallel distributed processing models. Moreover, it seems evident that all models of word recognition must involve some kind of feedback, since a principal source of contextual information derives from the preceding sentence or discourse, which, in turn, is built up out of word meanings. However, we presume that such interactions are slow (i.e., it is likely that they occur over the course of seconds or more) and are mediated by higher level comprehension processes and representations; the distinction we are making here concerns the direct relationship between access and integration during the initial stages of word recognition.

Evidence for the independence or dependence of access and integration is usually sought in the time course of meaning activation. In modular views, it is assumed that access and integration are independent of one another because they operate serially, one after the other. For example, it is common to refer to the lexical access *stage* and to contrast it with subsequent *post access* processing that performs meaning integration (e.g., Forster, 1981; Seidenberg, 1985; Stanovich, 1991). The assumption of seriality is plausible, because at least some types of context effects develop with time (see, e.g., Neely, 1977; Rayner & Pollatsek, 1989). In nonmodular views of meaning resolution, it is assumed that access and integration processes operate in parallel and interact throughout the time course of meaning activation (e.g., McClelland, 1987; Seidenberg & McClelland, 1989). Thus, evidence that suggests an early and rapid effect of context has been taken to support parallel access and integration processes and, by implication, nonmodular views of meaning resolution. Similarly, the failure to find early effects of context has been taken to implicate serial, modular views.

However, the link between the time course of context effects and the temporal arrangement of access and integration processes is not as straightforward as it might seem. It is necessary to distinguish patterns of meaning activation over time from the operation of the processes that produce those patterns. Activation patterns are inferred from observable performance on cognitive tasks and, presumably, reflect the relative availability of the various homograph meanings. However, the processes that produce those activations may change over time with little discernible effect on the availability of various meanings. Alternatively, the same constellation of activation processes may produce qualitatively different activation patterns over time, given different stimuli, contexts, or tasks. In particular, activation processes that operate in parallel can produce results that suggest seriality, whereas a sequence of activation processes may mimic parallelism. This point is distinct from the well-known mimicry between serial and parallel processes with respect to response time (e.g., Townsend, 1976): We are arguing that processes that operate in parallel (e.g., access and integration processes) can produce sequential mental events (e.g., one meaning may become active after another), rather than noting that serial and parallel processes can produce similar effects on response time. Thus, a pattern of activation over time reflects the potential contribution of all the processes that have operated to that point, and it is necessary to remain wary of the temptation to interpret a given activation pattern as being necessarily the result of a particular processing architecture.

Moreover, the question of whether access and integration are serial or parallel is not identical to the question of whether access and integration are independent. This point is illustrated by considering the possible arrangements of access and integration, as shown in Figure 1. Panel a shows a common modular view of meaning resolution: Access and integration operate serially and are independent of one another. For example, access processes based on perceptual information may activate several possible meanings of a word, and subsequent processes select and integrate one contextually appropriate meaning on the basis of information from the surrounding context. Panel b shows a common nonmodular view of meaning resolution. Here, access and integration operate in parallel and interact. In particular, integration processes affect ongoing access processes, so that perceptual information is used to activate only meanings that are contextually appropriate and to minimize activation of inappropriate meanings. This view is nonmodular because the use of perceptual information by the access mechanism can be affected by the current context; in this sense, access is cognitively penetrable and does not constitute a cognitive module. However, our analysis suggests that a third arrangement is also possible: Access and integration may be independent (as in the modular view) and parallel (as in the nonmodular view). This arrangement is shown in panel c. In this case, access and integration operate con-



**Figure 1. Possible arrangements of access and integration. In panel a, access and integration operate serially and are independent of one another. In panel b, access and integration operate in parallel and interact. In panel c, access and integration are parallel but independent.**

currently and jointly activate one or more meanings. However, in this case, the access and integration processes themselves do not interact; for example, the use of perceptual information by the access process would remain unchanged, regardless of whether or not a meaning is contextually appropriate. The independent activation model, described below, is an instance of this approach.

In sum, our view is that the lexical modularity debate has been hampered by a false dichotomy between serial, independent models, on the one hand, and parallel, interactive models, on the other. Although a few authors have argued for parallel and independent models (e.g., Oden & Spira, 1983), this approach seems to have been largely overlooked. We have developed a simple quantitative model of ambiguity resolution, the independent activation model, that can account for the major findings in the literature. This model is a hybrid of modular and nonmodular approaches because, although access and

integration are independent in the model (as in modular models), they also operate in parallel (as in interactive models). We anticipate that our model will serve to demonstrate that concentration on single issues, such as modularity or the lack thereof, does not serve the interest of broadening our knowledge of meaning resolution and language comprehension in general. Rather, our model demonstrates that, by moving beyond the modularity debate, we can clarify the role of ambiguity research in our understanding of language comprehension in general.

#### MODELS OF MEANING RESOLUTION

Below, we summarize influential models of lexical ambiguity resolution and describe some of the characteristic evidence that has been marshaled for each. This collection of data forms the foundation for the development of our independent activation model.

### Independent Serial Models

**Exhaustive access models.** The prototypical modular models of meaning access are exhaustive access models of ambiguity processing. In these models, all known meanings of a homograph are automatically accessed when that homograph is presented, regardless of context; context effects on meaning resolution are attributed to post-access processing, with an assumption that integration follows access (e.g., Conrad, 1974; Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Tanenhaus, Leiman, & Seidenberg, 1979; Till, Mross, & Kintsch, 1988). In extreme versions of exhaustive access, even meaning frequency does not affect homograph processing (Lucas, 1987; Onifer & Swinney, 1981; Swinney, 1979); frequent and infrequent meanings are not only accessed automatically, in parallel, but to equal degrees.

The primary source of evidence for exhaustive access models is the finding that, under many circumstances, no context effects are found early in processing. Swinney's (1979) cross-modal lexical decision methodology provides an example of this kind of finding. In Swinney's task, sentences (some of which contained homographs) were read to subjects over headphones. A sentence was either strongly predictive of one meaning of the homograph or was not biased toward either meaning of the homograph. During some of the sentences, a letter string was presented visually, and subjects were asked to decide whether the letter string was a word. The letter string was related to the contextually biased meaning of the homograph, was related to the contextually inappropriate meaning of the homograph, or was unrelated to either meaning. Swinney found that lexical decision times for both the contextually appropriate and the contextually inappropriate test words were faster than those for unrelated test words, regardless of the preceding context. The interpretation of this result is that both meanings of the homograph were accessed, regardless of contextual appropriateness. In a subsequent experiment, letter strings were not presented until three syllables after the homograph. In this case, facilitation occurred only for contextually appropriate test words. Swinney claimed that the lexical decision times in the delayed condition reflected postaccess integration processes that selected the contextually appropriate meaning within several hundred milliseconds.

**Ordered access model.** Another type of serial, independent model is the ordered access model. This model is similar to the exhaustive access model in that it assumes that contextual influence occurs postaccess. However, in this case, meaning access is related to meaning frequency (see, e.g., Forster & Bednall, 1976; Hogaboam & Perfetti, 1975; Holmes, 1979). The serial, self-terminating process of the access and integration of meanings is ordered by meaning frequency. The most common meaning is always accessed first. If the most common meaning is not appropriate to the context, the second most common meaning is accessed and integrated with the context; otherwise, meaning resolution termi-

nates. Thus, in ordered access models, whether multiple meanings are involved in meaning resolution depends on the suitability of the most common meaning to the context. As in exhaustive access models, meaning integration occurs only after access.

The principal data in favor of models such as the ordered access model are effects of meaning frequency. For example, in Hogaboam and Perfetti (1975), subjects attempted to detect, as quickly as possible, whether sentence-final words were ambiguous. This task is appropriate for studying the order of meaning access, because it requires access of more than one meaning for successful completion. Responses were faster for sentences in which the less frequent meaning of the homographs was appropriate than they were for sentences in which the more common meaning was appropriate. Hogaboam and Perfetti claimed that ambiguity judgments were easier when the secondary meaning was suggested by the preceding context because both primary and secondary meanings would have already been accessed. When the primary meaning was consistent with the preceding context, the secondary meaning would not normally be accessed; thus, the ambiguity judgment was delayed by the need to search for a second meaning. Hogaboam and Perfetti concluded that lexical access for homographs is ordered by relative meaning frequency. Their findings are consistent with modularity, as are exhaustive access models, because context does not influence access.

### Parallel Interactive Models

**Context-dependent models.** Context-dependent models are parallel and interactive. In this approach, the context preceding the homograph affects meaning access, and access is completed for only those meanings that are consistent with the context (see, e.g., Glucksberg, Kreuz, & Rho, 1986; Schvaneveldt, Meyer, & Becker, 1976; Simpson, 1981; Van Petten & Kutas, 1988). The crucial findings for such models consist of effects of context that occur early in lexical processing. Schvaneveldt et al. provide an example of this type of evidence. In their research, three letter strings were presented to subjects on each trial, one at a time, and subjects made a lexical decision to each of the three strings in turn. In the trials of interest, the second letter string was a homograph, and the first and third letter strings were words that were related to the meanings of the homograph. In the *concordant* condition, the first and third words were related to the same meaning of the homograph. In the *discordant* condition, the first and third words were related to different meanings of the homograph. The dependent measure was lexical decision time to the third word. Relative to unrelated controls, lexical decision time to the third word was facilitated in the *concordant* condition, but not in the *discordant* condition. Schvaneveldt et al. claimed that appropriate context narrowed access of word meanings for a word subsequent to a homograph. When the word following the homograph was related to the same

meaning (i.e., in the *concordant* condition), processing of that word was facilitated owing to the prior access of the meaning. When the word following the homograph was related to another meaning of the homograph (i.e., in the *discordant* condition), processing of that word was not aided, because the meaning related to that word was never accessed. Thus, the data suggest that access and integration are parallel, because early availability of meanings was a function of context.

Schvaneveldt et al.'s (1976) study, as well as some others supporting temporal concurrence of access and integration, has been criticized on the grounds that the context effects observed could be caused by lexical priming. Lexical priming is due to associative relationships between words in the lexicon and is presumably mediated by spreading activation from one word to related words. For example, when *nurse* is presented, all of the words to which *nurse* is associated are activated in the lexicon. If *doctor* is presented shortly afterwards, it will be easier to identify than it would be if no word, or an unrelated word, preceded it, because its activation level is still elevated owing to spreading activation from *nurse*. In a modular model, such activation is a by-product of access and does not involve integration processes. Thus, fast lexical-priming effects are consistent with the notion of a lexical module, because the effects are due entirely to access processes within the module.

More compelling evidence for context-dependent models requires demonstration of context effects even when the prior context contains no words that are associatively related to the homograph. It is commonly agreed that the effects of sentence and discourse context are the result of integration with syntactically or pragmatically related concepts and are not simply by-products of access (Seidenberg et al., 1982). For example, the word *shoe* is facilitated in the sentence *The puppy chewed on the shoe*, even though none of the words preceding *shoe* is strongly associated to *shoe* (Stanovich & West, 1983). Some research using materials of this type has demonstrated effects on meaning activation only after a considerable delay (e.g., Seidenberg et al., 1982); thus, these results provide support for the modular view of lexical access, rather than for context-dependent models. However, as described below, more recent research has found early effects of context in the absence of lexical associates, supporting parallel interactive models.

**Context-sensitive models.** Context-sensitive models of ambiguity processing are motivated by the finding that certain experimental situations may indeed provide evidence of early activation of contextually inappropriate meanings but these meanings will generally be activated to a lesser degree or more slowly than contextually appropriate meanings. The contextually inappropriate meaning of a homograph has been found to show *no* facilitation or *less* facilitation than the appropriate meaning, with delays of 100 msec or less in lexical decision (Simpson, 1981, Experiment 2; Swinney, 1979), naming (Seidenberg et al., 1982; Simpson & Krueger, 1991; Tanen-

haus et al., 1979), and color-word naming (Oden & Spira, 1983) tasks. (See Simpson, 1984, for a discussion of this across-task trend.) The event-related potential data of Van Petten and Kutas (1987, 1988) also show that, even though both appropriate and inappropriate meanings are facilitated early on in a naming task, the onset of this facilitation is delayed slightly for inappropriate meanings. According to context-sensitive models, whether a difference in activation between appropriate and inappropriate meanings is observed will depend on the sensitivity of the experimental situation. Variability across studies in context types and strength, the range of meaning frequencies used, tasks, timing, and dependent measures contributes to experimental sensitivity differences. The main tenet of context-sensitive models is that meaning activation is a function of relative meaning frequency and context strength. Thus, context-sensitive models are parallel and interactive. They differ from context-dependent models in that they do not suggest that *only* contextually appropriate meanings are accessed; rather, context and meaning frequency jointly determine meaning activation levels along a continuum.

### Resolving the Evidence

Experimental support has been found for independent serial models and for parallel interactive models of access under a wide variety of conditions, and there does not seem to be any simple way to resolve the conflicting results (for reviews, see Bubka & Gorfein, 1989; Rayner et al., 1994; Simpson, 1984, 1994). In particular, Simpson (1994) provided considerable support for the conclusion that type of task (naming, lexical decision, or Stroop task), position of homograph within sentence context (medial or sentence-final), and type of context presentation (cross-modal or unimodal) were not predictive of results in lexical ambiguity studies; support for both modular and interactive models has been reported under each of these conditions. Instead, there is now abundant evidence that the patterns of results that are obtained in various studies are related to the nature of the *stimuli* that are used. Below, we review some of the evidence for the role of two stimulus factors in meaning resolution: context strength and meaning frequency.

**Context strength.** It seems clear that a crucial difference between experiments that demonstrate early effects of context on meaning resolution and those that do not concerns the nature of the context. In particular, contexts that produce early and robust effects of context may be stronger or more potent than those that produce minimal effects. For example, Simpson (1981) found a pattern of results reminiscent of exhaustive access when the context was weakly biased toward a meaning, but results suggesting context-dependent access with strongly biased contexts. Similar results have been obtained by Simpson and Krueger (1991). However, a clear understanding of the nature of context strength has been elusive. In addition to the inclusion of lexical associates, other suggestions for definitions of context strength have included

the use of pragmatic biases toward homograph meanings or of stereotypic phrases using the homograph (Carpenter & Daneman, 1981), the repetition of homographs or their meanings (Rayner et al., 1994; Simpson & Kang, 1994), the priming of characteristic features of homograph meanings (e.g., Tabossi, 1988, 1989, 1991; Tabossi, Colombo, & Job, 1987), and an overlap in meaning features between context and target words in Stroop tasks (Paul, Kellas, Martin, & Clark, 1992). Our view is that strong contexts for homographs are those that (by whatever means) successfully narrow alternatives down to a single plausible meaning, not only by suggesting one meaning, but also by excluding alternative meanings (see Twilley, 1996). This is not inconsistent with approaches taken by the researchers referred to above and is similar to claims made by others (e.g., Dopkins, Morris, & Rayner, 1992; Holmes, 1979; Reder, 1983).

**Meaning frequency.** The effects of meaning frequency are related to the ubiquitous effects of word frequency that are found in virtually all language processing tasks. However, word frequency counts (e.g., Carroll, Davies, & Richman, 1971; Kučera & Francis, 1967) do not have separate entries for each of the multiple interpretations of a homograph and, consequently, cannot be used to estimate frequency of occurrence of the meanings in the language. Instead, estimates of *relative* meaning frequency are obtained by requiring subjects to provide semantic associates to homographs and calculating the proportion of responses made to each meaning (e.g., Twilley, Dixon, Taylor, & Clark, 1994). Effects of relative meaning frequency have been found in numerous tasks with homograph stimuli, and these effects have been used to model aspects of meaning resolution and language processing. Some of these insights are reported below.

One source of evidence on the role of meaning frequency comes from ambiguity detection tasks, often used to support ordered access models (e.g., Forster & Bednall, 1976; Hogaboam & Perfetti, 1975; Holmes, 1979; Neill, Hilliard, & Cooper, 1988). As was discussed above, when the sentence context is biased toward the high-frequency or dominant interpretation of the homograph, ambiguity detection is relatively slow, presumably because the low-frequency interpretation is difficult to find. Detection of ambiguity is faster when the low-frequency or subordinate interpretation is supported by context. Decisions concerning syntactic role have also been used to study the relative effects of meaning frequency on ambiguity processing (e.g., Forster & Bednall, 1976; Kinoshita, 1985). In these tasks, subjects are required to judge the grammaticality of sentences. Grammaticality judgments of sentences with homographs are faster when high-frequency meanings are used than when low-frequency meanings are used; this result suggests that the high-frequency meaning is available first. Hogaboam and Perfetti noted that the types of effects mentioned above demonstrate that meaning access for homographs is, indeed, influenced by the structure of the

lexicon (as indexed by relative meaning frequency) and that response times can be used to make inferences about this structure.

Eye movement monitoring during reading comprehension has shown that gaze durations are shorter for homographs with one very high frequency meaning than for homographs with two approximately equally frequent meanings (e.g., Duffy, Morris, & Rayner, 1988; Rayner & Duffy, 1986). Presumably, the gaze duration reflects the time to find an interpretation of the homograph. Because the high-frequency meaning is easier to find than either of two meanings of moderate frequency, the gaze duration is shorter. In this case, relative meaning frequency results supported the assumption that processing difficulty is reflected in word-by-word reading times and that such effects are immediate (see Just & Carpenter, 1980).

Priming tasks also show effects of meaning frequency. Naming of homographs is faster when the prime is related to the high-frequency meaning than when the prime is related to the low-frequency meaning (Simpson & Foster, 1986; Simpson & Krueger, 1991). Similar results are obtained in lexical decision tasks (Burgess & Simpson, 1988; Frost & Bentin, 1992; Lucas, 1987; Simpson, 1981; Simpson & Burgess, 1985; Tabossi et al., 1987). Each of these priming studies was designed to reveal whether access and integration were serial or parallel. As a group, they are equivocal on this point, with some studies supporting seriality and other studies supporting parallelism. However, these studies, in combination with others, have led to a crucial insight: Meaning frequency and context strength have early interactive effects in homograph meaning resolution.

**Combined effects of meaning frequency and context strength.** Context effects vary with both context strength and relative meaning frequency (Carpenter & Daneman, 1981; Duffy et al., 1988; Holmes, 1979; Kellas, Paul, Martin, & Simpson, 1991; Kinoshita, 1985; Rayner & Frazier, 1989; Simpson, 1981; Simpson & Krueger, 1991; Tabossi et al., 1987). For example, Simpson (1981) found exhaustive access, early effects of context, and early effects of meaning frequency, depending on the particular combination of meaning frequency and context strength used. When a homograph was preceded by a sentence context that was weakly biased toward the low-frequency meaning, both interpretations of the homograph were primed equally. When the homograph was preceded by a strongly biasing sentence, only the contextually appropriate interpretation was primed, regardless of its meaning frequency. When there was no context (homograph only), only the high-frequency interpretation was primed. Simpson's (1981) study has been criticized because an interstimulus interval (ISI) of up to 120 msec occurred between sentence and target, allowing the possibility that access and integration were serial and that context effects occurred after access. However, Simpson and Krueger (1991) replicated the Simpson (1981) data without this delay, demonstrating facilitation only for

contextually appropriate sentences in strong contexts and different facilitation patterns for dominant and subordinate meanings in unbiassing contexts.

A different, but related, interaction of context and meaning frequency was observed by Sereno and Rayner (1993; see, also, Sereno, 1995). In their study, context sentences were followed by target sentences. The context sentences were either consistent with the meaning of the target sentences or inconsistent. The fast-priming paradigm was used, in which a homograph was presented very briefly (for 35 msec) just as readers fixated on a critical word in the target sentence. The critical word was related to either the dominant or the subordinate meaning of the homograph, and homographs were classified as biased (having a dominant meaning with a relative frequency of more than 75%) or balanced (having a dominant meaning with a relative frequency of less than 75%). Priming effects were measured by comparing reading times for critical words preceded by related homographs or by unrelated words. There was an interaction between context type (consistent vs. inconsistent), meaning frequency (dominant vs. subordinate), and bias (biased vs. balanced): Sereno and Rayner found priming effects only for the dominant meanings of biased homographs in consistent contexts.

The pattern of results observed in these studies clearly indicates that meaning activation is determined jointly by meaning frequency and context strength; early context effects are not solely a function of associative relationships. In particular, no effects of meaning frequency are expected in strong contexts; in weak or absent contexts, meaning frequency is expected to be the dominant contributor to meaning resolution. Models of meaning resolution must be able to produce both exhaustive access and context-dependent access patterns of activation, depending on experimental conditions (i.e., context strength and meaning frequency). Thus, a resolution of the evidence seems to require some form of context-sensitive model of meaning resolution. However, we argue that context sensitivity does not necessarily require an interactive model of meaning resolution; rather, such effects can be generated by access and integration processes that operate in a parallel and independent fashion. As a demonstration of this point, we describe the independent activation model below.

### THE INDEPENDENT ACTIVATION MODEL

In the present section, we describe a quantitative model of meaning resolution, the independent activation model, which simulates the joint effects of context strength and meaning frequency, as well as mechanisms related to individual differences and development. As will become clear in the following discussion, the model can easily incorporate the kinds of effects that have been observed in the meaning resolution literature with a simple structure and underlying logic. The general structure of the model is first described, along with its nature of opera-

tion; subsequently, we describe the way in which the model accounts for a variety of results concerning meaning resolution. For illustrative purposes, the model is described in neural network terms. However, we wish to emphasize that this is purely heuristic and that the model is also consistent with a variety of other types of cognitive mechanisms and formalisms.

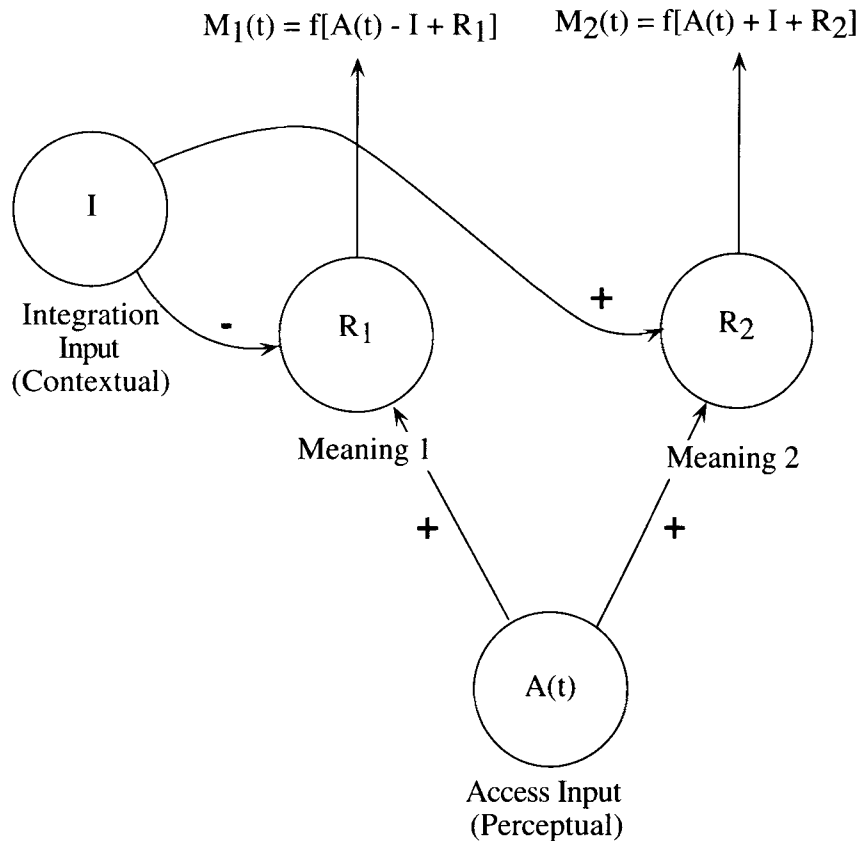
### Structure of the Model

In the independent activation model, meaning resolution is achieved by access and integration processes operating in parallel but independently, as in Figure 1c. For illustration, we assume that each meaning of a homograph is represented by a node and that access and integration processes provide independent input to those nodes. This structure is depicted in Figure 2.

One of the most common instantiations of semantic knowledge is that of a network of nodes (often called a semantic network); our model uses this general approach. Examples of such an approach can be found in Balota (1994), Collins and Loftus (1975), Kawamoto (1988, 1993), MacDonald, Pearlmutter, and Seidenberg (1994), Posner and Snyder (1975), Seidenberg and McClelland (1989), and Waltz and Pollack (1985), but see McKoon and Ratcliff (1992) for an alternative. Each meaning node has a resting activation level and receives input from access and integration processes. The meaning nodes for a homograph increase in activation above their resting level when the homograph is perceived. So, for example, if the word *board* is perceived, the concepts *board as a piece of wood* and *board as a committee* (as well as other meanings, such as *boarding a ship* and *room and board*) will both become active, owing to input from access processes, with the *wood* meaning being considerably more active than the *committee* meaning (relative meaning frequencies, .70 and .07, respectively; Twilley et al., 1994).

We assume that the lexical system depicted in Figure 2 is embedded in a larger comprehension system that includes perceptual processes, working memory, and long-term memory. As each word is encountered, perceptual operations identify the features and constituent elements of the word and provide the access input that activates word meanings. As the meanings are activated, they become available to working memory and may be incorporated in the development of sentence and discourse level representations. We assume that the word meanings are not copied directly into working memory but, rather, are recoded in the service of ongoing comprehension processes that are driven, in part, by other long-term memory information, the demands of the task, and the goals of the comprehender. However, the precise nature of these working memory representations and the operation of the comprehension processes is beyond the scope of this article. As working memory representations are built, they, in turn, provide input to the lexical system and can affect the activation of subsequent word meanings as they are encountered. This provides the major contribution to the integration input indicated in Figure 2. The input can





**Figure 2. Structure of the independent activation model. Each node represents a homograph meaning, with an associated resting level determined by meaning frequency. Access input to each node is generated by the perceptual system and varies over time; integration input is determined by prior context and is either positive or negative. The activation level of a node at any point in time is a monotonic function of the sum of integration, access, and resting-level inputs.**

be either positive (supportive of a meaning) or negative (contraindicative of a meaning). For example, if the word *board* appears in the sentence *The carpenter sawed the board*, the integration input will be strongly positive for the *wood* meaning and strongly negative for the other meanings. Input to the meaning nodes from access processes occurs parallel to but independent of integration input.

A second component of the integration input is provided by any closely related word nodes that are simultaneously active. In modeling the activation of homograph meanings, we adopt the simplifying assumption that this additional lexical input is zero. This assumption is justified in simulating many results because the experimental materials are carefully designed to eliminate lexical associates of the homograph in the immediately preceding context. However, intralexical inputs play an essential role in the priming measures often used to measure the activation of homograph meanings. In particular, if an associate of one meaning of a homograph is presented shortly after the homograph, the activation of the (unambiguous) target meaning would be affected by input coming from the associated homograph meaning, if it is

active. This mechanism allows the model to relate the activation of homograph meanings to priming effects that might be measured experimentally.

The core of the model is the meaning nodes. There are some proponents of the idea that word meanings are unitary (e.g., Armstrong, Gleitman, & Gleitman, 1983; Smith, Shoben, & Rips, 1974), whereas in feature theories of meaning (e.g., Medin & Schaffer, 1978; Rosch, 1975), word meanings consist of a collection of semantic features. Many supporters of feature theories have argued that meanings are more than mere feature lists; for example, there have been arguments that theories of the world are used to tie features together (Murphy, 1991; Murphy & Medin, 1985) or that features are abstracted from stored episodic traces of all exposures to words (Anderson & Nagy, 1991; see, also, Balota, Ferraro, & Connor, 1991, and Masson & Freedman, 1990, for similar arguments about meaning representation). In the independent activation model, word meanings are discussed as unitary entities; that is, there is local representation of meaning (cf. Cottrell, 1988; Waltz & Pollack, 1985). In many respects, it is reasonable to think of the nodes as logogens

for the separate meanings (cf. Morton, 1969, 1979). However, in the independent activation model, the inputs vary over time: At any point in time, the activation of each homograph meaning is simply a monotonic function of the sum of the inputs to the associated node.

In distributed representations, a word's meaning is distributed across many semantic features (see, e.g., Kawamoto, 1988, 1993; McClelland & Kawamoto, 1986); we demonstrate below that a distributed version of our model generates predictions that are identical to those of the local version. We purposefully do not advocate one scheme as superior to the other: Such speculation is beyond the scope of this model. Rather, we note that, because either a local or a distributed representation scheme can successfully underlie simulations of the major results in the field, we can learn a great deal about homograph meaning resolution without specifying details of representation. Of course, a complete model of meaning resolution will eventually have to specify underlying representation, and we expect that such specification will be implemented when it becomes clear which type of scheme best represents neural structure.

There are four senses in which the model represents independent activation of word meanings. First, there is no interaction between access and integration processes: Both contribute to the activation of word meanings in the same manner, regardless of the nature of the other input. This would not be the case, for example, if access input were only directed to contextually appropriate meanings or if the perceptual system were actively primed for features and letters that were expected on the basis of the context. Second, there is no feedback from the word meanings to the access input processes, and the perceptual processes operate in the same manner, regardless of the level of activation of the meanings. This structure can be contrasted, for example, with the operation of models such as the interactive activation model of word recognition proposed by McClelland and Rumelhart (1981). In that model, letter detectors that are consistent with a word are reinforced with additional input from the word detector whenever it is activated. Third, there is no interaction among the word meanings themselves. For example, the meanings representing different possible interpretations of a word are not mutually inhibitory, nor does the structure of the model preclude the possibility that several meanings may be active simultaneously. Finally, there are no internal dynamics involved in the process of activating word meanings. This means that the activation of a node is determined simply by the current level of access and integration inputs and is independent of the node's immediate history of activation levels or inputs.

There are, however, three aspects of the model to which the term *interactive* might be applied. First, the access and integration inputs combine to determine the level of activation of word meanings, and, as will be detailed below, we assume that this combination is nonlin-

ear. The effect of this nonlinear combination rule is that the change in meaning activation associated with a given change in access input will vary with the level of integration input. In other words, the pattern of meaning activation displays an interaction effect, as opposed to purely additive effects. Although the term *interactive* is appropriate as a description of this activation pattern, it should not be confused with the notion of process interaction, discussed earlier: We assume that integration and access are independent (noninteractive) processes, even though their outputs are combined in a nonlinear fashion at the level of word meanings.

Second, there is an implicit feedback loop between the level of word meanings and the integration input: When a word meaning is activated, the information may be incorporated into a working memory representation of the sentence and discourse, which may, in turn, affect the subsequent activation of word meanings. Interaction at this level is logically required in any model of meaning resolution. In our approach to understanding homograph meaning resolution, we assume that the processing of lexical input in working memory is relatively slow (on the order of half a second or more) and that, consequently, the potential for feedback from the activation of a word meaning to itself via the integration input is minimal. This allows us to model the integration input as effectively constant over the course of immediate word recognition.

Finally, there is also the potential for interactive feedback to occur among related word meanings. As was mentioned previously, we assume that an active word meaning provides integration input to related word meanings and that the related meanings, in turn, could provide input to the first meaning. However, because we hypothesize that the integration input coming from other meaning units is relatively weak and the integration input by itself would be insufficient to activate the related meaning, this form of feedback has minimal effect on the behavior of the model. Weak lexical input of this sort would be sufficient to produce lexical priming. For example, if one meaning is fully active, integration input coming from that unit could speed the activation of a related meaning that is presented immediately afterward. However, in this case, feedback from the second meaning would have little effect on the first, since it is already fully active.

There are four substantive assumptions in the model. First, in order to determine activation of each meaning, it is assumed that the inputs to each meaning are summed and then scaled into the range 0–1. This means that, if both access and integration inputs are moderately strong, the activation will be near the peak activation of 1 and further increases in input strength will have little effect on the activation level. This assumption is at the core of many of the simulations described below. For mathematical tractability, we use the logistic function to accomplish this scaling. In other words, if the net input to a unit

is  $N = A + I$ , where  $A$  is the input from access processes and  $I$  is the input from integration processes, the activation of a meaning is

$$M = \frac{1}{1 + e^{-2(N+R)}} = \frac{1}{1 + e^{-2(A+I+R)}}, \quad (1)$$

where  $R$  determines the resting level. This assumption determines how the access and integration input are combined at the level of word meanings and determines much of the interesting behavior of the model.

Second, it is assumed that integration input is determined by prior context and is relatively constant throughout the course of meaning resolution. Prior context generally consists of the activation of meaning units for words appearing prior to the homograph, as well as pertinent world knowledge, and it is likely that it is stored in a working memory system, as was described above. Although integration input is amassed over the period of time preceding the appearance of the homograph, it is assumed that this process is slow and that the input does not change appreciably during the period of time it takes to process the homograph. This assumption reflects the conditions examined in most experiments on ambiguity resolution: Contextual information is presented several hundred milliseconds prior to the presentation of a homograph, and meaning activation is measured over the course of the first second or two after presentation of the homograph. Clearly, the assumption of constant integration input would have to be changed under other circumstances. For example, it seems unlikely that the nature of the context will remain constant over the course of several seconds or more, and it is reasonable to suppose that contextual information provided immediately prior to a homograph will have limited immediate effect on meaning resolution. However, our working hypothesis is that, when the context is presented substantially prior to the homograph (i.e., several hundred milliseconds earlier, as is commonly done), there may be little meaningful variation in the contextual input to each meaning during the first second after homograph presentation.

The third assumption of the model is that, in situations in which prior context disambiguates the meaning of a homograph, integration input enhances the contextually appropriate meaning, while suppressing the contextually inappropriate meaning. Moreover, we assume that these effects are symmetric. In other words,  $I$  is *added* for interpretations that are consistent with the prior context, and  $I$  is *subtracted* for interpretations that are inconsistent with the context. (In situations in which the prior context is unbiased or neutral, we assume that  $I$  is 0.) The use of symmetric enhancement and suppression is mathematically simpler than alternative formulations (i.e., it requires fewer parameters). Moreover, the assumption is consistent with the only study of homograph meaning resolution we know of that looked at the issue of enhancement and suppression symmetry (Simpson & Kang, 1994). Given the dearth of evidence to the contrary, we

have assumed symmetry in our model. However, whether prior context generates symmetric enhancement and suppression of word meanings remains an important empirical question.

The final assumption of the model is that the input from access processes is determined by perceptual encoding mechanisms and varies over time. Many of the interesting predictions of the model depend on the nature of the time-varying input; we hypothesize that the input starts at zero when the word is first encountered, rises rapidly to a peak within several hundred milliseconds, and then decays relatively slowly back to zero. In other words, we assume that the perceptual input is transitory, even though a stimulus may remain in view. This would be appropriate, for example, if the lexical system is responsible only for signaling the content of new lexical input. After the meaning of a word is identified, other mechanisms (in working memory, e.g.) would be responsible for incorporating and maintaining that word meaning. The decay in access input consists of a form of rapid habituation to sustained lexical information and is analogous to operation of transient visual pathways, albeit at a quite different level and time scale.

One class of encoding processes that could produce input of this form consists of a combination of sampling and decay. For example, assume that the perceptual system has available a quantity  $f$  of evidence in favor of a particular lexical item; conceivably,  $f$  might represent the number of visual features consistent with that word. Access processes would sample the available evidence randomly at a rate  $s$ . This means that the probability of information remaining unsampled will decay exponentially over time—that is,  $U(t) = e^{-st}$ . However, we assume that the information also decays at a rate  $d$ . Consequently, the probability of information remaining active will also decay exponentially:  $V(t) = e^{-dt}$ . If the decay and sampling are independent, the probability of information being available to the access process is simply the product of information sampled and remaining active, or  $[1 - U(t)]V(t)$ . This means that the total information provided by the access process can be represented by a difference of exponentials:

$$A(t) = f(e^{-dt} - e^{-(d+s)t}). \quad (2)$$

This curve is illustrated for the values  $d = .001$  and  $s = .01$  in Figure 3; these values are used for all of the simulations reported below (with the exception of those for the performance of young children, for whom we assume these values are inappropriate; see below). However, our approach to meaning resolution would be unchanged for a range of values for  $d$  and  $s$ ; all that is required is that the input wave form has the general qualitative character shown in the figure. In particular, our choice of parameter values and even the form of Equation 2 are unimportant, as long as the input reaches a maximum in roughly 200–300 msec and decays rather slowly subsequent to that peak.




Figure 3. Access input as a function of time. This is a difference of exponentials (Equation 2) with  $s = .01$  and  $d = .001$ . The curve has been normalized to have a peak value of 1.

In sum, Equations 1 and 2 imply that, when the meaning is consistent with the context, its activation at time  $t$  is given by

$$M(t) = \frac{1}{1 + \exp\left\{-2\left[f\left(e^{-dt} - e^{-(d+s)t}\right) + I + R\right]\right\}}. \quad (3)$$

Similarly, when the meaning is inconsistent with the context, its activation is given by

$$M(t) = \frac{1}{1 + \exp\left\{-2\left[f\left(e^{-dt} - e^{-(d+s)t}\right) - I + R\right]\right\}}. \quad (4)$$

In order to generate a simulation of the model, parameter values are selected for  $f$ ,  $I$ , and  $R$ , and the equations are evaluated at each point in time. The simulations reported below are based on default values of  $f = 20$ ,  $I = 5$ , and  $R = -9$  (along with the parameters that determine the access wave form,  $d = .001$  and  $s = .01$ ). Logical considerations impose a variety of constraints on the values that these parameters must have. For example, the combination of resting level and integration strength has to be low enough that meanings are not spontaneously "recognized" without any access input from the perceptual system. Similarly, the access input has to be strong enough that meanings achieve a substantial amount of activation even in the absence of contextual support. Within these constraints, the parameter values trade with one another, so that qualitatively similar curves are generated with a range

of parameter values and there is no clear basis for preferring one set of parameter values to another. Although a small amount of exploration was required to find parameter values that produced intuitively reasonable simulations, curves were not "fit" in the sense that a systematic search of the parameter space was required in order to generate particular predictions.

### Model Simulations

Below, we describe a series of simulations designed to illustrate how the model accounts for the varying patterns of results in studies of homograph processing. In most of those studies, the dependent variable was a measure of priming, and it is typically assumed that a particular meaning of a homograph is active if it can affect the recognition or activation of related words. The independent activation model predicts the level of activation of the homograph meanings, and we regard the construction of a response model capable of predicting priming effects as a somewhat separate issue. In a subsequent section, we will describe a possible account of priming that is based on the effects of intralexical integration inputs. For the moment, however, we adopt the heuristic that high levels of activation (e.g., those greater than .9) will produce strong priming of related words, whereas low to moderate levels of activation (e.g., those less than .4) will not.

**Exhaustive access.** As was discussed above, a common finding in research on meaning resolution is exhaustive access: Both meanings are active shortly after presentation of a homograph, but only the contextually appropriate meaning is active after a longer delay. This pattern of activation is to be expected when context is weak or absent, and the model produces this pattern when a relatively low value is used for the strength of the integration process (parameter values for this and all subsequent simulations are listed in Table 1). Figure 4 shows the simulation of the model when the value of  $I$  is reduced to 4 (from the default value of 5). As can be seen, both meanings are highly active shortly after the presentation of the homograph, although the contextually appropriate meaning is activated more quickly. Indeed, even when exhaustive access patterns are observed empirically, there is a tendency (rarely statistically significant within a given study, but consistent across studies; see Lucas, 1999, for a meta-analysis) for the appropriate meaning to show slightly larger priming effects than the inappropriate. The difference in rate of activation between appropriate and inappropriate meanings in our model nicely simulates this subtlety in the empirical data. Larger effects of the differential contextual support are found only after several hundred milliseconds. Thus, this pattern of activation corresponds to the exhaustive access pattern of priming effects observed in experiments such as those of Seidenberg et al. (1982): Words related to both meanings are primed at short delays, but only the contextually appropriate meaning is primed at longer delays. In the model, the reason this pattern of results occurs is that the input from the perceptual system is strong enough initially to

**Table 1**  
**Simulation Parameters**

Simulation	<i>d</i>	<i>s</i>	<i>f</i>	<i>I</i>	<i>R</i>
Exhaustive access (Figure 4)	.001	.01	20	4	-9
Context-dependent access (Figure 5)	.001	.01	20	6	-9
Context $\times$ frequency interaction (Figure 6)					
Subordinate meaning	.001	.01	20	5	-10
Dominant meaning	.001	.01	20	5	-8
Fast priming (Figure 7)					
Subordinate meaning	.001	.01	6	5	-10
Dominant meaning	.001	.01	6	5	-8
Balance (Figure 8)					
Subordinate meaning neutral context	.001	.01	20	0	-12
Dominant meaning neutral context	.001	.01	20	0	-6
Balanced meaning neutral context	.001	.01	20	0	-9
Subordinate meaning biased context	.001	.01	20	4	-12
Dominant meaning biased context	.001	.01	20	4	-6
Balanced meaning biased context	.001	.01	20	4	-9
Priming (Figure 9)					
Target	.001	.01	20	4	-9
Homograph, weak context	.001	.01	20	4	-9
Homograph, strong context	.001	.01	20	6	-9
Development (Figure 12)					
Subordinate meaning, older children	.001	.01	20	0	-10
Dominant meaning, older children	.001	.01	20	0	-8
Subordinate meaning, younger children	.0005	.005	20	0	-10
Dominant meaning, younger children	.0005	.005	20	0	-8
Individual differences (Figure 13)					
High-capacity readers	.001	.01	20	4	-9
Low-capacity readers	.001	.01	20	2	-9


overwhelm any distinctions arising from context and to activate both meanings, even if only one is favored by the context. However, as the strength of the perceptual input decays, the contextual bias toward one meaning will keep that meaning active, while suppressing the alternative meaning.

**Context-dependent access.** Some research (described earlier) shows a context-dependent pattern of access in strong contexts. The model predicts this access pattern if the strength of the integration input is increased. In particular, if the value of *I* is increased to 6, the model produces results that look like the predictions of the context-dependent model. This is shown in Figure 5. The access input succeeds in activating the contextually appropriate meaning in conjunction with the input from integration processes. However, the contextual bias against the inappropriate meaning of the homograph is now strong enough to keep this alternative meaning from becoming more than minimally active. Our interpretation of the results shown in Figures 4 and 5 is that the essential difference between experiments showing the exhaustive access pattern and those showing context dependence is in the nature of the stimuli used to provide contextual support. In particular, with strong contexts, the integration process is sufficiently strong to select only that meaning, but with weaker contexts, integration is not strong enough to keep the contextually inappropriate meaning from becoming momentarily active (cf. Simpson, 1981).

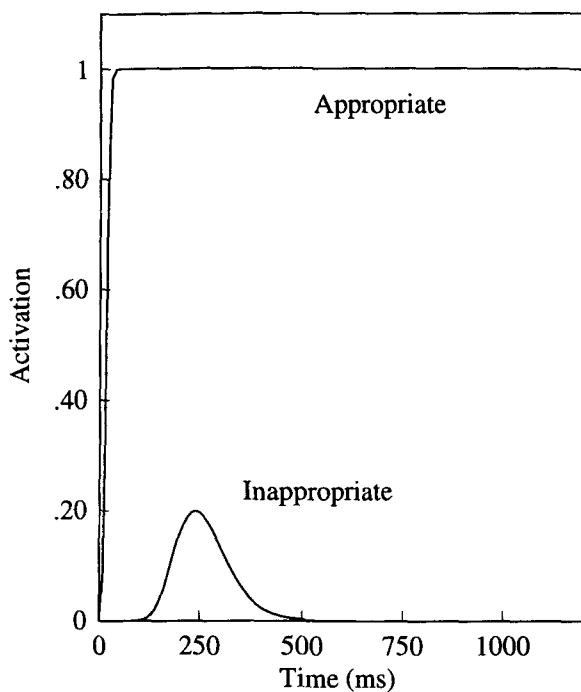
In the model, we use varying values of the parameter *I* to capture the differences in contexts provided in these

studies. This approach embodies the assumption that contexts differ primarily in terms of a single characteristic, strength of constraint. We cannot defend this assumption with any compelling arguments and know of no simple way in which the strength of a context can be assessed a priori. In fact, it is quite possible that contexts of different types operate differently and that the model would have to be elaborated to incorporate a variety of different kinds of integration processes. For example, in a subsequent section, we will consider the possibility that enhancement and suppression of homograph meanings might be different. At present, however, the differing patterns of context effects can be explained simply by assuming that contexts, whatever their nature, vary in the strength of their input to the lexical system.

**Context  $\times$  frequency interactions.** According to the model, differences in meaning frequency can modulate the effects of context, such as those modeled above. For example, a subordinate (less frequent) meaning may not become active at all if it is not supported by context (i.e., the pattern of activation will suggest context-dependent access), whereas the dominant (more frequent meaning) may become active regardless of the context (i.e., the pattern of activation will suggest exhaustive access). This interaction is illustrated in Figure 6. In order to produce this simulation, we made the plausible assumption that variations in meaning frequency among the stimuli are associated with variations in resting level. In this case, the resting level (*R*) was reduced from -9 (the default value)



**Figure 4.** Simulation of exhaustive access. Predicted activation levels are shown for the contextually appropriate and contextually inappropriate meanings with strong access input and moderate integration input.



**Figure 5. Simulation of context-dependent access. Predicted activation levels are shown for the contextually appropriate and contextually inappropriate meanings with strong access input and strong integration input.**

to  $-10$  to represent a subordinate meaning and was increased to  $-8$  to represent a dominant meaning. We used the default value of context strength ( $I = 5$ ). The left panel depicts the pattern of activation for the subordinate meaning. As is shown, it becomes fully active when it is contextually appropriate but shows little activation when it is contextually inappropriate. The right panel depicts the patterns of activation for the dominant meaning. In this case, the meaning becomes completely active both when it is contextually appropriate and when it is contextually inappropriate, although the meaning remains active only when it is contextually appropriate. This interaction of meaning frequency and context is similar to that reported by Simpson and Krueger (1991; see also Simpson, 1981). Notably, the time course of activation in the model is solely a function of the varying strength of the perceptual input, together with appropriate values of the resting levels to represent differing meaning frequencies.

The patterns of activation shown in Figure 6 may also underlie the effects observed in the ambiguity judgment task used by Hogaboam and Perfetti (1975). They found that the time needed to detect the ambiguity of a sentence-final homograph was faster when the sentence was consistent with a subordinate meaning than when it was consistent with the dominant meaning. When the sentence is consistent with the subordinate meaning, both the appropriate (subordinate) meaning and the inappropriate (dominant) meaning are fully active. In this case, it should

be relatively easy for subjects to detect the potential ambiguity, because both meanings are available. When the sentence is consistent with the dominant meaning, the appropriate (dominant) meaning is still fully active, but the inappropriate (subordinate) meaning reaches only a low activation level. Here, one would expect it to be more difficult to retrieve the subordinate meaning of the homograph, and consequently, ambiguity detection time should be greater. Hogaboam and Perfetti interpreted their results in terms of an ordered access model in which the dominant meaning is accessed first, followed by the subordinate meaning. However, the present simulation demonstrates that a serial mechanism is not required to produce their results; they can be explained in a straightforward manner by the parallel activation mechanism in the independent activation model.

Sereno and Rayner (1993; Sereno, 1995) reported a somewhat different interaction between meaning frequency and context, using the fast-priming procedure. In their results, priming results were observed only for the dominant meaning, when it was contextually appropriate; little priming was found for the context or for the subordinate meaning under any circumstances. However, in their task, the homograph was visible for only about 35 msec. Thus, it would be reasonable to model their results with a relatively weak perceptual input. Figure 7 shows the result of using precisely the same parameters as those used in the simulation shown in Figure 6, with the exception that the strength of the perceptual input ( $f$ ) was reduced from 20 to 6. The results are similar to those obtained by Sereno and Rayner: The only case in which a meaning is highly activated (and in which substantial priming would result) involves a meaning that is both dominant and contextually appropriate. Our interpretation is that the particular form of interaction between meaning frequency and context observed in the fast-priming paradigm is dependent on the brief exposure of the homograph, which, in turn, leads to only weak perceptual input to the homograph meanings. Notably, the *re-ordered access* model put forth by Sereno and Rayner is not necessary to explain this pattern; rather, their results can be simulated with a simple and plausible parametric variation in the independent activation model.

**Effects of homograph balance.** With additional assumptions, the model can predict some aspects of the effects of relative meaning frequency observed by Rayner and colleagues on fixation duration during comprehension (Duffy et al., 1988; see, also, Rayner et al., 1994). They observed that when prior context was consistent with both meanings of a homograph, fixation duration on balanced homographs (i.e., those with two equally frequent meanings) was longer than that on control words. However, fixation duration on unbalanced homographs (i.e., those with one meaning that is substantially more frequent) was not elevated, relative to controls. A different pattern emerged when the prior context was consistent with the less frequent meaning of the homograph. In this

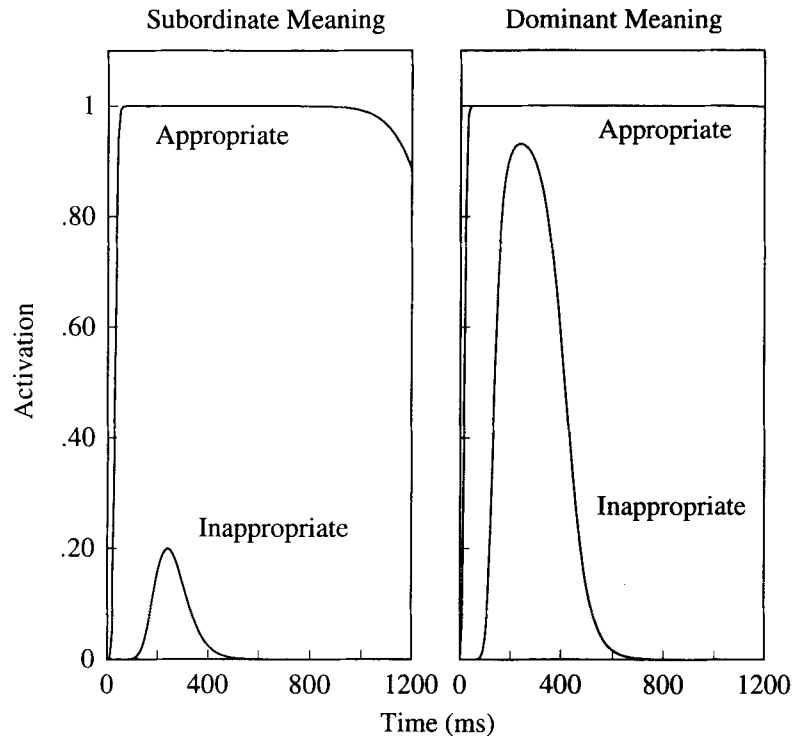


Figure 6. Simulation of interaction of context and meaning frequency. The left panel shows the predicted activation of the subordinate meaning when it is either contextually appropriate or contextually inappropriate; the right panel shows the predicted activation of the dominant meaning when it is either contextually appropriate or contextually inappropriate.

case, fixation duration was longer on unbalanced homographs (i.e., homographs for which the contextually inappropriate meaning was dominant), and little effect of ambiguity was observed for balanced homographs. These effects were attributed to competition among possible meanings for integration into the text representation (Dopkins et al., 1992; Duffy et al., 1988; Rayner et al., 1994). For example, when there is no biasing context and the two meanings are of comparable frequency, the two meanings are equally available, and the ensuing competition slows processing. In contrast, when one meaning is dominant, that meaning is accessed first, and there is little competition. However, slower fixations result if the prior context favors the subordinate meaning, because the contextual support speeds the access of that meaning, so that it is available at the same time as the more frequent meaning.

Our account of these phenomena is similar in that we also assume that fixation duration is slowed when the two meanings of a homograph are active simultaneously. However, the relationship between the duration of eye fixations and the predicted levels of activation is complicated by the fact that fixation durations are typically much shorter than the total duration of meaning resolution hypothesized in the model. For example, fixations are often less than 300 msec, whereas, according to the model, it may take as much as 500 msec to suppress a

contextually inappropriate meaning (e.g., as shown in Figure 4). Consequently, we assume that fixation duration is primarily sensitive to the initial stages of meaning activation, and we model the effects observed by Rayner and colleagues by examining the predicted activation of the two meanings at a 100-msec delay. In particular, we assume that fixation duration is slowed by the presence of competing interpretations of a homograph when both are active at 100 msec, and there should be no effect of ambiguity when only one meaning is available at this initial stage.

With this assumption, the model predicts the patterns of fixation duration observed by Rayner and colleagues when the unbalanced homograph meaning frequencies are sufficiently disparate. The first panel in Figure 8 shows the pattern of activation at 100 msec with a neutral context. When the two meanings are of equal frequency (i.e.,  $R = -9$  for both meanings) and there is no biasing context (i.e.,  $I = 0$ ), both meanings are active, and, as a consequence, fixation duration should be relatively slow. When one meaning is substantially more frequent (i.e.,  $R = -6$  and  $-12$  for the dominant and subordinate meanings, respectively), only the more frequent meaning is active at 100 msec, and there should be no effect of ambiguity on fixation duration. In the second panel, contextual support for the subordinate meaning was added




Figure 7. Simulation of interaction of context and meaning frequency with weak access input. The left panel shows the predicted activation of the subordinate meaning when it is either contextually appropriate or contextually inappropriate; the right panel shows the predicted activation of the dominant meaning when it is either contextually appropriate or contextually inappropriate.

( $I = 4$  for the subordinate meaning and  $I = -4$  for the dominant meaning). With this biasing context, the patterns are reversed: Only the subordinate meaning of a balanced homograph is fully active, whereas both meanings are active for strongly unbalanced homographs. As Duffy et al. (1988) found, fixation duration on the unbalanced homograph should be slower, because of the presence of a competing meaning. The independent activation model produces these patterns without any sequences of operations or serial access stages; rather, all the meanings are activated independently and in parallel.

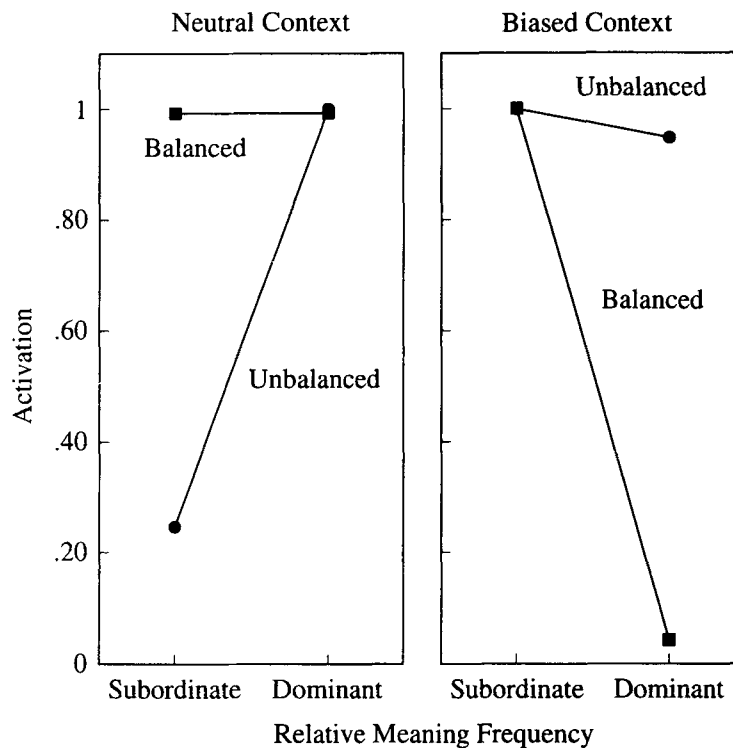
### Priming Effects

With a few additional assumptions, the model produces quantitative estimates of priming effects. Although we regard this component of the model as speculative and somewhat tangential to our main argument, it is useful to demonstrate the nature of the development that is required to produce such predictions. In a priming paradigm, a lexical associate of one of the homograph meanings is presented shortly after the homograph, and recognition time for the target word is measured (e.g., in a lexical decision task). It is usually assumed that, when the related meaning of the homograph is active, the target recognition will be faster, as compared with a control condition in which the target is unrelated to either meaning of the homograph. In order to simulate priming ef-

fects in the model, we assume that full activation of a meaning is not necessary for it to be recognized but that meanings are recognized whenever their activation is substantially above zero. The magnitude of priming effects can be calculated, if we assume a specific value for the recognition threshold. In the following development, we assumed that the target meaning can form the basis of a word recognition response as soon as the activation reaches a value of .2. Using the default parameters from the previous simulations and an integration input of 0, the activation level reaches this recognition threshold in 58 msec. However, the precise value of the threshold is unimportant, and other values of the threshold produce comparable estimates.

Lexical priming is produced when a related homograph is presented prior to the target. If the homograph-target stimulus onset asynchrony (SOA) is short enough, the homograph meanings may still be active, and the related meaning would contribute to the integration input for the target. We assumed that this input would be relatively weak even when the homograph meaning is fully active. As a consequence, we used parameter values for the integration input that are appropriate for a weak context manipulation (such as that described above in the Exhaustive Access section). Thus, for the related target meaning,  $I(t) = 4 M(t)$ , where  $M$  represents the activation of the homograph meaning at time  $t$ . As can be seen in Fig-





**Figure 8.** Simulation of the effect of balance and contextual support on initial levels of activation. The left panel shows the pattern of activation without contextual support 100 msec after presentation for balanced homographs (two equally frequent meanings) and unbalanced homographs (one subordinate and one dominant meaning). In the right panel, contextual support in favor of the subordinate meaning was added. (One of the equally frequent meanings of the balanced homograph was arbitrarily designated as the subordinate meaning.)

ures 4–6, the contextually appropriate meaning of the homograph is fully active almost immediately after presentation, so for targets related to the appropriate meaning, the integration input would be essentially 4 for the initial stages of their meaning activation. With this additional input, the time to reach the threshold of .2 decreases to 25 msec. In other words, the priming effect for contextually appropriate meanings should be  $58 - 25 = 33$  msec and should be constant over a large range of homograph–target SOAs.

The pattern of activation of contextually inappropriate meanings is more complex and varies with the strength of the integration input to the homograph meaning. Figure 9 shows two possible patterns. The curve labeled “Inappropriate–Weak Context” indicates the predicted priming when the target is related to a homograph meaning receiving an integration input of  $-4$  (which produces the pattern of activation shown for the inappropriate meaning in Figure 3). Although substantial priming is found when the target is presented immediately after the homograph, very little priming is obtained when there is a small delay. The curve labeled “Inappropriate–Strong Context” indicates the predicted priming when the integration input to the homograph meaning is  $-6$  (which

produces the pattern of activation shown for the inappropriate meaning in Figure 9). In this case, very little priming is observed at any SOA. In general, the patterns of priming effects mirror the patterns of homograph meaning activation, albeit with a different scale and with less sensitivity to low levels of activation.

Clearly, this development cannot be regarded as a theory of word recognition or priming. There is a great deal of processing that must be involved in making a lexical decision response; for example, and the actual magnitude of the predicted priming effects cannot be regarded as anything more than illustrative. The point of this exercise is merely to indicate one way in which the theoretical quantities of meaning activation can be connected in a concrete way to observable priming effects.

### Distributed Lexical Representations

Thus far, our description of the independent activation model has been based on the assumption that the meanings of homographs are represented as single units. However, this assumption has been purely for the sake of expositional convenience. In particular, the relationships and predictions of the model can also be cast in terms of distributed lexical representations. Given recent

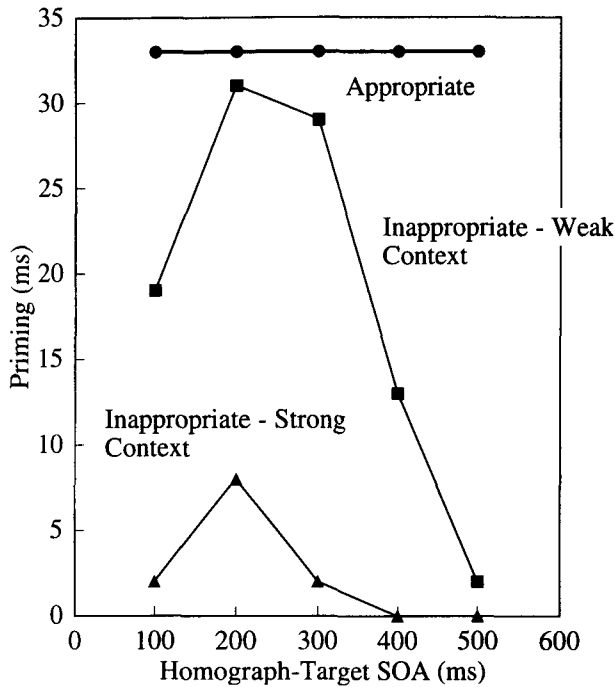


Figure 9. Simulation of priming effects as a function of time and context. The simulated priming effect is the decrease in the time for a meaning related to one of the homograph meanings to reach a recognition threshold.

controversy about distributed representations of meanings of ambiguous words (Besner & Joordens, 1995; Joordens & Besner, 1994; Masson & Borowsky, 1995; Rueckl, 1995), we felt that consideration of distributed representations was in order.

In the following section, we will develop a parallel representational model of the lexicon in which inputs

and outputs are distributed and show that precisely the same equations and predictions apply. The model is illustrated in Figure 10. Each interpretation of a homograph is represented by a pattern of activation in the output units at the top of the figure; the output units can be thought of as meaning features that together code the meaning of words. Each of these meaning feature units receives input from access and integration units at the bottom. As in the localist representation described previously, the access units code orthographic information from the perceptual system, whereas the integration units code information about the context and are based on representations in working memory. As before, we assume that the activation of each of the output units is simply the sum of its input, logarithmically scaled to the range of 0–1. That is, the output of each unit is

$$M_j = \frac{1}{1 + e^{-2(N_j + R_j)}}, \quad (5)$$

where  $N_j$  is the net input and  $R_j$  determines the resting level. The input is derived from a set of input units, divided between access and integration, as shown in Figure 11. The activation levels of the total set of input units can be described as a vector and decomposed, using vector addition, as  $A + I$ , where  $A$  codes the activations of the access units (with 0 for the values of the integration units) and  $I$  codes the activations of the integration units (with 0 for the values of the access units). The net input to the meaning units,  $N$ , is modulated by a matrix of connection weights,  $W$ :

$$N = W(A + I), \quad (6)$$

where  $A$  and  $I$  are vectors representing the pattern of inputs over the access and integration units, respectively, and the  $i$ - $j$ th element of  $W$  represents the strength of the connection between the  $i$ th input unit and the  $j$ th mean-

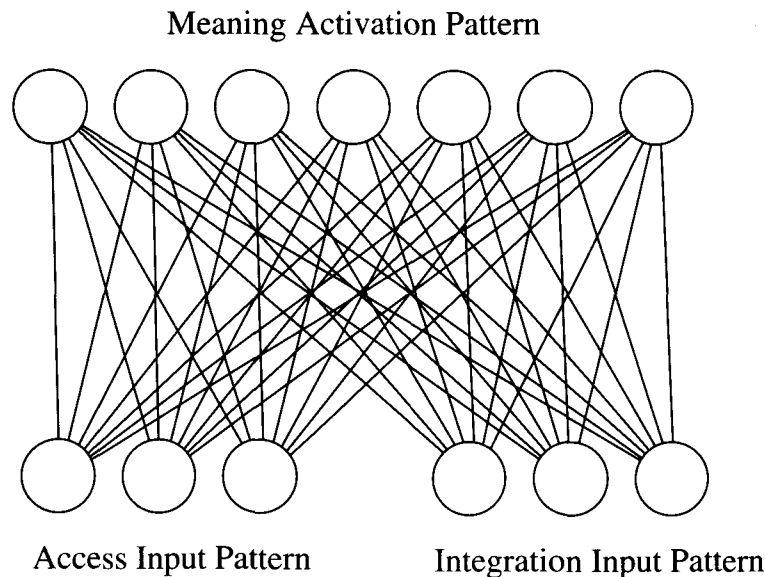


Figure 10. Structure of the independent activation model with distributed representations.




Figure 11. Hypothesized access input as a function of time for younger and older children. The curve for older children uses the same parameters as those used for adults,  $s = .01$  and  $d = .001$ ; the curve for younger children uses half those values,  $s = .005$  and  $d = .0005$ . The curves have been normalized to have a peak value of 1.

ing unit. The present formulation is similar to a linear associator system, such as those studied by Kohonen (1987) and others. However, in a linear associator, matrix multiplication (as in Equation 6) determines the nature of the output, not just the net input. In contrast, our system is effectively a linear associator where the output values are constrained to be in the range of 0–1.

Connection weights in this kind of system can be learned by using the generalized delta rule (Rumelhart, Hinton, & Williams, 1986). In applying this procedure, a parallel distributed system is presented with an input pattern, and the output pattern is compared with a target. The connection weights are altered incrementally in proportion to the strength of the input and the magnitude of the discrepancy between target and output. This learning procedure can be applied to the lexical system depicted in Figure 10 by assuming that, for a given homograph, the system would learn to produce one meaning when presented with the appropriate orthographic input and supporting context and to produce a different meaning when the context supports the alternative meaning. In other words, suppose that the vector  $A$  represents the pattern of access inputs corresponding to a homograph and that  $I_1$  and  $I_2$  represent patterns of integration inputs for two different interpretations of the homograph. Then, the output target should be  $T_1$  when the system is presented with

$A + I_1$  and  $T_2$  when presented with  $A + I_2$ , where  $T_1$  and  $T_2$  represent the patterns of meaning features in the two homograph meanings. In addition, when the system is presented with the orthographic input without any supporting context or with a neutral context, the output should be a mixture of the two meanings; that is,  $A$  alone should produce

$$\frac{T_1 + T_2}{2}.$$

With these assumptions, the exact form of the connection matrix,  $W$ , can be specified, and the pattern of activation can be predicted for any given input (see the Appendix). In order to assess the extent to which a meaning is active, we compute the dot product of the target vector with the obtained pattern of activation, normalized by the length of the target vector:

$$M_1^* = \frac{T_1 \cdot M}{\|T_1\|^2}. \quad (7)$$

This is essentially a measure of the extent to which features that are present in the target pattern of activation are also active in the obtained pattern. In particular, it is shown in the Appendix that, when presented with access input of strength  $\alpha$  and integration input of strength  $\beta$  in favor of meaning 1, the activation of that meaning is

$$M_1^* = \frac{1}{1 + e^{-2(A^* + I^* + R^*)}}, \quad (8)$$

where  $A^*$  is a function of  $\alpha$ ,  $I^*$  is a function of  $\beta$ , and  $R^*$  determines the resting level. This, of course, is precisely the same formulation described in Equation 1 for localist representations.

We assume that the role of the activation vectors in comprehension processes is also comparable with that which would occur with a localist representation. In particular, after a meaning has been activated, the vector of features would be available to higher level comprehension processes and would likely be incorporated into working memory representations of the sentence and discourse. We assume that the lexical system is responsible for responding to sequences of words encountered one at a time; the present system is sufficient in this regard. How a sequence of word meanings are used in comprehension and maintained in working memory is a complex problem beyond the scope of the present article. However, one aspect of the interface between patterns of lexical activation and working memory deserves comment. The simulations described above make it clear that, in many circumstances, both meanings of a homograph would be available for a brief period after the word is encountered. With a distributed representation, this creates ambiguity about which features should be used in comprehension. However, this ambiguity is exactly analogous to the ambiguity that occurs when multiple meanings are momentarily active in a localist representation. In either case, the sim-

ulations indicate that the ambiguity is resolved over time and that the system generally reaches and maintains a state where only a single meaning is active. Thus, one possibility is that working memory integrates lexical activations over time and that transitory activations of contextually inappropriate meanings have relatively little impact.

The distributed representation described here provides a parsimonious account of the selection of word senses. Many relatively unambiguous words have several related meanings, only one of which is appropriate in a given context. For example, *booster*, meaning an electronic amplifier, shares meaning components with *booster*, meaning a fan or enthusiast, even though the interpretations are mutually exclusive. In order for the independent activation model to select one meaning by using a localist representation, one would have to hypothesize that a separate node exists for each meaning sense. This seems awkward, since almost all words have a variety of different but overlapping senses. However, this is less of a problem in a distributed representation. For example, it would be natural to assume that two related senses would share some features but not others, so that the meaning vectors for the two senses would be correlated but not identical. Because these two meaning vectors would, nonetheless, be linearly independent (although not orthogonal), developments analogous to those presented above would apply. In particular, features that are active in one meaning sense but not in the other would be sensitive to contextual influences, just as the features in unrelated meanings are. In contrast, features that are common to both senses would be fully active regardless of the nature of the context. Thus, the independent activation model, together with the distributed representational scheme described above, would be capable of sense disambiguation without requiring a separate meaning node in the lexical system for each meaning sense.

The conclusion that we draw from this development is that the mathematics of the independent activation model hold regardless of at least some of the details of lexical representation. Importantly, either local or distributed models of lexical representation can be assumed to underlie meaning access and integration without affecting the validity of our model. The crucial aspect of the model is that the system be capable of activating more than a single lexical representation at a time and that the access and integration inputs make independent contributions to the activation of different interpretations. As long as these constraints obtain, it does not matter whether lexical items are conceived of as single nodes (with associated lexical knowledge) or as distributed patterns over a set of interpretive semantic units; the same predictions and mathematical formulation apply to both.

### Relations to Other Models of Lexical Ambiguity

**Context-sensitive models.** The independent activation model is a context-sensitive model in the sense that access and integration go on concurrently and in parallel. However, unlike most context-sensitive models, ac-

cess and integration in the independent activation model are independent and do not interact; the access processes provide equal support for all possible meanings, regardless of context. Moreover, there are no serial stages in the model (e.g., an access stage is not followed by an integration stage), and there are no complex process dynamics (e.g., there is no interactive feedback from the meaning nodes to access input). Rather, we have found that simply adding the resting level, perceptual inputs, and contextual inputs (thus granting each independent status) can provide an accurate description of the findings on meaning resolution.

**Classical models of perception.** The independent activation model can be regarded as a development of the logogen model developed by Morton (e.g., 1969, 1979) in which each homograph meaning is represented by a different logogen. In particular, both models share the assumption that perceptual processes and context contribute independently to the activation of a meaning. McClelland (1991) termed models with this property *classical models of perception* and contrasted them with interactive activation models, in which perceptual and contextual information interact to produce activation levels. Other classical models include the choice model of Luce (1963) and the fuzzy logical model of Oden and Massaro (1978). The formal relation between the logogen model and the independent activation model can be expressed as follows. In the logogen model, each unit has an activation level that is simply the sum of perceptual input and contextual input. In principle, these inputs are counts of the number of features that are consistent with the word represented by the logogen, but they can be approximated by real values (Morton, 1969). In other words, the activation level of a logogen is  $L = P + C$ , where  $P$  is the perceptual input and  $C$  is the contextual input. Furthermore, the activation level is subject to random noise, approximated by Morton (1969) as a logistic distribution. This entails that the probability of a logogen being greater than its threshold  $v$  is

$$\Pr(\text{Exceeding threshold}) = \frac{1}{1 + e^{-2(L-v)}} = \frac{1}{1 + e^{-2(P+C-v)}}. \quad (9)$$

This is formally identical to the activation level in the independent activation model (expressed in Equation 1), if  $v = -R$ ,  $P = A$ , and  $C = I$ .

However, despite this formal relationship between the two models, there are fundamental conceptual differences. To begin with, activation levels in the independent activation model do not correspond to the probability that the homograph is recognized. For example, skilled readers can accurately recognize words even when they are unexpected or incongruous in context, as long as the words are clearly perceptible; thus, the recognition probability should be close to 1. However, simulations such as that shown in Figure 4 suggest that the activation level for a meaning that is incongruous in context should be about

.2 in the independent activation model. This difference between meanings that are minimally active (e.g., an activation of .2) and those that are fully active (e.g., with an activation of .9 or higher) is essential to the predictions of the model, because only the latter case is associated with strong priming of related meanings. However, both meanings should be easily recognized, as long as the words are substantially above threshold. Thus, activation levels cannot be interpreted directly as recognition probabilities, despite the formal relation to the logogen model. Rather, in order to estimate recognition rates at low levels of perceptibility, one would have to make assumptions about noise in the system, add a threshold, and then compute the probability that the activation is greater than the threshold (as Morton does in his description of the logogen model).

Another crucial difference between the two conceptualizations is that activation levels in the independent activation model are logistically scaled to the range of 0–1. In contrast, in the logogen model, the logistic distribution only arises because of noise in the system, and the actual activation level is simply the sum of the inputs. The implications of this difference could be seen most clearly if one were required to choose between the two possible meanings of a homograph. Such an operation is similar to the forced-choice tasks modeled by Morton (1969, 1979), and his predictions can be applied directly. Using the mathematical development of Luce (1963), Morton demonstrated that the probability of selecting one of two logogens is

$$\text{Pr(Choose Meaning 1)} = \frac{C'_1 P'_1}{C'_1 P'_1 + C'_2 P'_2}, \quad (10)$$

where  $C'$  and  $P'$  are strength measures that depend only on the contextual and perceptual inputs, respectively. If the logogens correspond to the two meanings of a homograph, this reduces to  $C'_1/(C'_1 + C'_2)$ , because the two meanings have identical perceptual inputs. In other words, the probability of selecting one meaning over the other is unaffected of the strength of the perceptual input. These predictions do not apply to the independent activation model, because activation levels are logistically scaled, rather than simply being the sum of perceptual and contextual information. For example, one could easily imagine that the meanings of a homograph are equally likely to be selected if both are fully active (indeed, this is comparable with the *exhaustive access* priming prediction). Thus, the conceptually different treatments of meaning activation in the two models lead to quite different predictions.

The most important unique element in the independent activation model is the idea that perceptual input varies over the course of meaning activation, rising rapidly to a peak within several hundred milliseconds and then decreasing more slowly. This characteristic of the model is essential for predicting the temporal patterns described in the simulations above. In particular, the time-varying perceptual input underlies the ability of the model

to settle on one meaning of a homograph even though both meanings may be active initially (e.g., in Figure 4). Although Morton (1969) suggests that the activation of logogens must decay, this property was intended to allow the system to respond to successive words during normal language processing and is not involved in the dynamics of recognizing individual words.

In sum, the independent activation model shares with the logogen model (as well as with other classical models) the idea that context and perception have independent effects. Moreover, there are formal relations between the models (e.g., Equations 1 and A7). However, there are important differences: Despite the formal similarity, activation levels in the interactive activation model are conceptually different than recognition probabilities in the logogen model; the activation of each meaning is a logistic, rather than a linear, function of the inputs; and the perceptual input varies over time in the independent activation model, rather than being static. These elements are essential to the behavior of the independent activation model.

**Interactive activation models.** Kawamoto (1993; see, also, Kawamoto, 1988) proposed an interactive activation model of homograph meaning resolution. In his model, each lexical entry is represented as a distributed representation that encodes spelling, pronunciation, part of speech, and meaning. Recognizing a homograph in context is simulated by activating units corresponding to the word's spelling and units corresponding to the meaning context. The system subsequently settles into a state corresponding to one particular meaning of the homograph. Thus, the structure of his model is similar in some respects to Figure 10: Input to the system is provided in the form of activation levels in perceptual and contextual units, and interpretation of a homograph is represented by a pattern of activation over a set of semantic units.

There are many differences between our approach and Kawamoto's, but we think two are fundamental. The first is that the semantic units in Kawamoto's model receive direct input not only from perceptual and contextual units, but also from other semantic units; there is no comparable immediate interaction among the meaning units in the independent activation model. The second difference is in the nature of the activation function of each of the units. In the independent activation model, the activation of each output unit at any point in time is determined simply by the instantaneous sum of the inputs. In Kawamoto's model, however, activation levels grow incrementally over time. This means that the activation at any point in time is a joint function of the previous state of the system, as well as of the perceptual and contextual inputs. The system as a whole gradually settles into a state that minimizes the difference between each unit's activation level and its net input; this asymptotic state represents the selected meaning. The implication of this approach is that meaning resolution involves a complex interaction over time between access and integration. Although perception and context provide independent in-

puts to the system, the model cannot be described as performing independent activation in the same sense as that for the model developed here.

As a consequence of this difference in meaning resolution, the two models provide distinct accounts of the temporal dynamics of meaning resolution. In the independent activation model, changes in meaning activation over time are assumed to be a product of the transitory nature of the perceptual input. Although we have suggested a possible mechanism that could generate the assumed time course of the access input, this mechanism is, in principle, external to the meaning resolution process. In contrast, in Kawamoto's (1993) interactive activation approach, temporal dynamics are inherently a property of the lexical system. This theoretical distinction could be evaluated empirically. For example, the approach taken in the independent activation model implies that temporal dynamics are related to perceptual mechanisms and, consequently, may be quite different when lexical items are presented in different perceptual modalities: The resolution of auditorily presented homophones may follow a different time course than the resolution of visually presented homographs, because two different perceptual systems are involved. In contrast, Kawamoto's model provides no basis for expecting the temporal dynamics to vary with perceptual modality. The approach taken in the independent activation model also suggests that there may be some similarity between meaning resolution with visual input and the temporal dynamics of other visual tasks. For example, Busey and Loftus (1994) estimated information extraction rates on the order of  $0.02 \text{ msec}^{-1}$  in a task requiring subjects to report briefly presented digits. This value is similar to the information sampling rate assumed here for the access process ( $s = .01 \text{ msec}^{-1}$ ). In our model, it is natural to find similar time constants in tachistoscopic letter or digit recognition, because both are determined by the same perceptual mechanisms. However, according to the interactive activation approach, temporal dynamics in meaning resolution are determined by the lexical system, and one would have to conclude that such a correspondence is coincidental. Our impression is that the evidence to date on temporal dynamics is not sufficiently detailed to allow a clear choice with respect to such predictions.

Another essential difference between the present approach and that of Kawamoto (1993) is in the nature of meaning representation. His semantic representation consists of a set of features that are either present or absent in a particular interpretation, just as in the distributed lexical representation described above. However, in his case, presence is indicated by the value +1, whereas absence is indicated by -1. As a consequence, all of the target patterns in Kawamoto's model have the same maximal length, and fully activating one pattern precludes the complete activation of any other pattern. In contrast, the distributed representation developed above assumes that meanings are coded by only the presence of features and that two meanings could be simultaneously active, if

the constituent features of both meanings are active. Interestingly, the inability of Kawamoto's model to fully activate two meanings simultaneously does not prevent it from predicting something like the exhaustive access pattern found in many cross-modal priming experiments. As Kawamoto demonstrated, the model initially activates both interpretations to a certain extent by activating those units that the two interpretations have in common. Only later, as the system settles into a single interpretation, does the activation of alternative meanings decline.

The interactive activation approach generates an important qualitative prediction that is not shared by the independent activation model: In the interactive model, the peak activation for the subordinate meaning should be earlier than that for the dominant meaning, and there should be some period of time in which the activation of one meaning is increasing while the activation of the alternative meaning is declining (e.g., Figure 6 in Kawamoto, 1993). This follows from the fact that there is some distance between the energy basins corresponding to the two interpretations of a homograph and that the system must move away from one basin as it settles into the other. The independent activation model makes a strong prediction to the contrary: The sole source of temporal changes in the model is the time-varying perceptual input, and this input is identical for the two meanings of a homograph. This implies that one meaning cannot be increasing in activation while the other is decreasing. For example, in the right-hand panel of Figure 6, the dominant meaning (in its appropriate context) reaches its maximum prior to the peak of the subordinate meaning in the same (now inappropriate) context and remains fully active until after the peak of the subordinate meaning. Although some data are consistent with the present predictions (e.g., Simpson & Burgess, 1985; Simpson & Krueger, 1991), the effects are small, and more evidence is required in order to answer this question.

Another type of interactive activation model was used by Cottrell (1988) to simulate ambiguity resolution. Cottrell used a localist representation in which different homograph interpretations were represented by different nodes; this is similar to our initial exposition of the independent activation model. As a consequence, it shares with the present model the property that different interpretations of a homograph can be fully active at the same time. Each of these nodes received both access and integration input in a fashion analogous to our approach. However, Cottrell's model makes use of interactive activation in order to settle on a particular meaning: Access input filters through a meaning discrimination net, and pathways representing distinct meanings are mutually inhibitory. (A related approach has been explored by Waltz & Pollack, 1985.) In this type of framework, all the meanings receive some input initially; the inhibition will tend to narrow the access input to a single interpretation. However, a crucial property of Cottrell's model is that the interaction of access and integration is limited. Input from the context is only received by the meaning

nodes, and there is no direct connection between context and the access pathways. Although contextual influences may eventually feed back indirectly through access pathways, this can only occur after meaning nodes have reached threshold. As a consequence, the system generally produces exhaustive access, and effects of context are only found after the multiple interpretations have been activated.

Conceptually, one may divide the operation of the Cottrell (1988) model into two phases. During the first phase, access and integration are independent, and the activation of meaning interpretations is comparable with that assumed in the independent activation model. The access input in Cottrell's model is itself the product of a complex interaction among the different meanings of a homograph that depends on their frequency and semantic distance. Just as in the independent activation model, however, this process is unaffected by surrounding context, and the initial activation of the meaning nodes is simply a function of the sum of the access and the integration inputs. On the other hand, Cottrell's model is unlike the present approach, in that context does not directly inhibit inappropriate meanings. This means that both meanings of a homograph will be active initially and that the model does not capture the early effects of context found by Tabossi (1988, 1989, 1991) and others.

The second phase occurs after the interpretations reach threshold and start to produce top-down feedback to the access pathways. During this phase, the system is more fully interactive, and effects of context intermingle with the access input, to settle on a single interpretation. During this latter phase, the dynamics of meaning activation are dissimilar to those of the independent activation model. However, there is a sense in which the use of thresholds in Cottrell's (1988) model is related to the mechanisms used here. The logistic scaling assumed to represent the activation of meanings in the independent activation model minimizes the effect of contextual inputs when the access input is low; contextual effects are only observed when the access and integration inputs combine to produce moderate levels of net input. This mechanism is comparable with the use of thresholds in Cottrell's model; in this case, context only produces feedback when the combined integration and access inputs activate a meaning unit above threshold.

In general, we feel that the evidence on homograph meaning resolution does not permit a definitive choice between the present approach and simulation models, such as those of Kawamoto (1993) and Cottrell (1988). Moreover, each model could be recommended on the basis of other criteria. The model of Kawamoto is consistent with the growing tradition of interactive activation models of cognition (e.g., Masson, 1991; McClelland, 1987; McClelland & Elman, 1986; Seidenberg & McClelland, 1989), which have been argued to be an advance over mathematical or symbolic models because they describe computation by using biologically plausi-

ble mechanisms. On the other hand, the independent activation model builds on the success of the logogen model of word recognition and other classical models of perception. Moreover, we have shown above how the model is consistent with distributed lexical representations and connection weights that are learned with experience.

Our preference for the independent activation model is based, in part, on its mathematical tractability: It admits of a simple formulation that readily generates qualitative and quantitative predictions under a wide range of circumstances. Furthermore, some predictions of interactive activation models, such as those of Kawamoto (1993) and Cottrell (1988), may depend on detailed computer simulations. As an example, consider the following prediction. The equations of the independent activation model imply that the activation levels of two meanings of a homograph always have the same rank order: The meaning the system eventually selects will be at least a little more active at all times, and the contextually inappropriate meaning is never more active than the appropriate meaning, even if it is more frequent. This is consistent with behavioral data (e.g., Simpson & Burgess, 1985; Simpson & Krueger, 1991). The same qualitative pattern is evident in Kawamoto's simulations. However, it is not clear to us whether this pattern is an intrinsic characteristic of models such as Kawamoto's or whether it depends on particular aspects of the training regime, stimulus representations, and initial conditions.

## IMPLICATIONS AND EXTENSIONS

Although modularity (or the lack thereof) of meaning access is an important question, our understanding of meaning resolution must not be limited to this issue. We need to understand both meaning access and meaning integration, if we are to be able to answer larger questions about lexical representation and process as they apply to reading for comprehension. A great proportion of previous empirical work with homographs was concerned with the question of whether one or more meanings were initially accessed. Processes leading to selection and integration of meanings were largely ignored, because the main interest was the modularity debate, rather than meaning resolution. However, the evidence indicates that extreme positions on modularity are not tenable; rather, the evidence suggests that the focus of ambiguity research must be broadened to include the entire meaning resolution process, not simply the first several hundred milliseconds of processing (cf. Carpenter & Daneman, 1981). Below, we discuss several issues that are underspecified in the literature on ambiguity resolution: meaning suppression, the development of meaning resolution processes, and the locus of individual differences. We see the independent activation model as providing a focus and vocabulary for framing these issues, although it is likely that the model will have to be reworked as further evidence accumulates.

### Meaning Suppression

A crucial problem, for which there is only a modest amount of data, concerns the fate of inappropriate meanings. Data relevant to this question have been obtained in studies in which context–target ISI is manipulated. Cross-modal lexical decision and naming data suggest that, when an inappropriate meaning is activated initially, its activation decreases to minimal or nonexistent activation with increasing ISIs (Onifer & Swinney, 1981; Swinney, 1979; Tanenhaus et al., 1979; Van Petten & Kutas, 1987). Unfortunately, the studies described above used a restricted range of ISIs and often incorporated an unspecified level of meaning frequencies and context types. As a result, the evidence they provide on the time course of meaning activation is imprecise.

More detailed results were obtained in two studies by Simpson and his colleagues (Simpson & Burgess, 1985; Simpson & Krueger, 1991). Simpson and Burgess presented homographs as primes for targets in a lexical decision task. They found that facilitation of responses to subordinate meanings of homographs (relative to unrelated words) followed an inverted-U pattern, with minimal facilitation at an SOA of 16 msec, peak facilitation at SOAs of 100 and 300 msec, moderate facilitation at a 500-msec SOA, and minimal facilitation at a 750-msec SOA. Dominant meanings had maximal activation by 100 msec with only a small decline in facilitation at longer SOAs. Simpson and Krueger's study is the only one in the literature in which the critical factors of context strength and meaning frequency are evaluated over a range of time intervals. They found an inverted-U pattern of facilitation of responses to subordinate meanings (relative to unrelated words) in nonbiasing contexts across ISIs of 0, 300, and 700 msec (dominant meanings maintained constant facilitation across ISIs, equal to the peak level of facilitation for the subordinate meanings at a 300-msec ISI), replicating the results of Simpson and Burgess. When homograph meanings were primed by contexts rated as strong but not containing lexical associates, responses to inappropriate meanings were not facilitated at any ISI (dominant meanings maintained a constant level of activation).

Several researchers have argued that one of the mechanisms behind these activation patterns must be active suppression, rather than passive decay alone (e.g., Gernsbacher & Faust, 1991a, 1991b; Paul et al., 1992; Simpson, 1984; Simpson & Kang, 1994; Twilley, 1996). Suppression processes are often implicated whenever inappropriate meanings are less active than appropriate meanings at short SOAs, because it is unlikely that such immediate effects would be due to decay processes (Onifer & Swinney, 1981; Tanenhaus et al., 1979). In our terms, decay occurs when a meaning gradually becomes less active, because of a failure to provide further positive input. Suppression occurs when the activation level of a meaning is reduced by processes that provide negative, inhibitory input. Suppression might be exhibited in two ways. We use the term *context-relative* suppression to

refer to situations in which the presence of an inappropriate context suppresses the activation of a meaning to a level below that which would be obtained with an unbiased (baseline) context. Baseline contexts to be compared with inappropriate contexts must be carefully chosen, so as to minimize bias as to which meaning is intended. For example, context-relative suppression could be implicated in the observation of a lower level of activation of *directors* following an inappropriate context, such as *He sawed the board*, as compared with unbiased contexts, such as *He chose the board* or *The word is board*. We use the term *meaning-relative* suppression to refer to the suppression of a meaning to a level of activation below its resting state; this can be assessed by comparing activation levels of words related to meanings with the activation level of words unrelated to the homograph. For example, meaning-relative suppression would have occurred if *directors* was less active than *detectors* following *He sawed the board* (note that unrelated words should be matched to related words on variables such as word frequency, number of syllables, etc.).

Evidence for context-relative suppression very early in processing is mixed. Some researchers found that meanings are less active in inappropriate contexts than in unbiased contexts with short SOAs (Schvaneveldt et al., 1976; Simpson, 1981; Simpson & Krueger, 1991). Paul et al. (1992) used a Stroop task to measure homograph meaning activation levels. Their data show that, at 0 ISI, homograph meanings are less active in inappropriate contexts (e.g., *He had to wear an old cast* [dominant meaning] and *They were all part of the cast* [subordinate meaning]) than in neutral contexts (e.g., *You realize that the word is cast*) when evaluated with targets related to high-salient features of the homograph meaning (e.g., *actors* is inappropriate to the dominant meaning, and *plaster* is inappropriate to the subordinate meaning). With low-salient targets, subordinate meanings (evaluated with Stroop responses to *actors* following *He had to wear an old cast*) are suppressed, relative to the neutral condition at 0 ISI, but dominant meanings are equally active in neutral and inappropriate contexts at 0 ISI. However, target meanings following appropriate contexts were also less active than those following neutral contexts, causing interpretation difficulties for Paul et al.'s results. Using a lexical decision task with a 0-msec ISI, Kellas et al. (1991) found identical results: suppression of dominant and subordinate inappropriate meanings with high-salient targets and suppression of subordinate meanings, but neither facilitation nor suppression of dominant meanings with low-salient targets. Appropriate meanings were facilitated in all cases except dominant meanings with low-salient targets. In Swinney's (1979) studies, there was no difference between inappropriate and unbiased conditions with short ISIs.

There is some evidence consistent with context-relative suppression at longer SOAs (e.g., Simpson & Krueger, 1991, for suppression of subordinate meanings only). Paul et al.'s (1992) data show suppression of dominant and sub-



ordinate inappropriate meanings at a 600-msec ISI. Subordinate appropriate meanings showed suppression at this ISI with high-salient targets (see above), whereas dominant appropriate meanings evaluated with high-salient targets showed no suppression or facilitation. With low-salient targets, neither dominant nor subordinate meanings showed any appreciable facilitation at a 600-msec ISI. Other data show that meanings are equally active in inappropriate contexts and unbiased contexts (Swinney, 1979). In all but the Kellas et al. (1991) results, cited above, suppression effects were not a focus of the experiments; the presence of effects was determined by our post hoc inspection of the conditions. In summary, although many researchers have argued for the presence of suppression effects, data to support this contention are suggestive but not overwhelming.

The evidence for meaning-relative suppression is very weak. At both short and long SOAs, most data provide no evidence of meaning-relative suppression (Onifer & Swinney, 1981; Simpson, 1981; Simpson & Krueger, 1991; Tabossi & Zardon, 1993; Tanenhaus et al., 1979; Van Petten & Kutas, 1987). Simpson and Kang (1994, discussed below; see, also, Simpson & Kellas, 1989) provided the only evidence we know of for meaning-relative suppression.

The independent activation model includes both decay and context-relative suppression mechanisms. The decline in activation of a contextually inappropriate meaning over time is due to a reduction in positive input from perceptual processes. According to our assumptions concerning the nature of the access processes, this decline is caused by the loss of information about visual features in the stimulus and constitutes a decay mechanism. In addition, the activation levels of inappropriate meanings are reduced, because of the negative input based on prior context; this produces context-relative suppression. The model does not produce meaning-relative suppression, because the perceptual input is always strong enough to keep a meaning above its resting level, even when that meaning is inhibited by an inappropriate context. It is further assumed in the model that contextual suppression and enhancement are symmetric (see the earlier discussion); that is, context serves to enhance the appropriate meaning and to suppress the inappropriate meaning to an equal degree. However, this assumption is only a matter of convenience, and the model could easily incorporate the assumption that enhancement and suppression have different strengths.

One possibility might be to assume in the model that there is no suppression and that context acts only in a positive manner to enhance consistent meanings only. Because of the structure of the model, integration inputs can trade with the resting level input to produce the same level of activation. Thus, it is technically possible to produce some of the same patterns of activation without suppression, if one lowers the resting level input for all meanings. This would mean, for example, that an inappropriate meaning would be minimally activated with-

out contextual support. However, it would also imply that, when words are encountered in isolation, in the absence of supporting context, none of the meanings would be fully activated. This seems implausible; consequently, a suppression mechanism seems to be required for a parsimonious interpretation of the model.

These decay and suppression mechanisms in the independent activation model are sufficient to account for much of the time course data reported above. According to the model, variations in activation over time are due to changes in the perceptual input to all the meanings. The existing evidence indicates that this input must reach a peak around 300 msec after presentation of the homograph and that it decays rather slowly thereafter (cf. Figure 3). Furthermore, changes in activation over time should be independent of context and meaning frequency. Although strong context or large differences in meaning frequency may mask temporal variations in perceptual input, whenever any variation is observed over time, it should always follow precisely the same pattern. In particular, a crossover activation pattern, in which one interpretation is more active than another at one point in time but less active later in time, would be inconsistent with the basic tenets of the model. What evidence there is on this issue (e.g., Simpson & Burgess, 1985; Simpson & Kang, 1994) conforms to the prediction of the model and indicates that activation is at a maximum at around 300 msec for any meaning, irrespective of meaning frequency or contextual input.

Evidence for meaning-relative suppression at longer intervals is more difficult to incorporate into the model. A clear example of this form of suppression is found in the results of Simpson and Kang (1994; see, also, Simpson & Kellas, 1989). In this research, a repetition-priming procedure was used in which homographs were presented twice; the measure of interest was level of activation on the second presentation. Homographs preceded target words by 200 msec; the task was to name the target word. The homograph meaning that was related to the target word was either the same or different across presentations. For example, if *bank* was followed by *save* in one trial and by *money* on a subsequent trial, the same meaning of *bank* had to be accessed in both trials. If, however, *bank* was followed by *save* in one trial and by *river* on a subsequent trial, the financial meaning of *bank* had to be suppressed to enable responding to the alternate meaning. Simpson and Kang found that, when a response was required to a homograph meaning that had previously been suppressed, the response to the previously suppressed meaning was inhibited to the point where responding was slower than that for unrelated stimuli. Simpson and Kang found this inhibition to occur only on the second presentation of the homograph, to be specific to competing meanings, to be long lasting (across as many as 12 priming trials), and to be unaffected by meaning frequency (in contexts that appear to be highly constraining). Meaning-relative suppression with such characteristics cannot be produced by the independent activation

model as it stands, and other kinds of meaning resolution processes may need to be postulated in order to explain these results. One possibility is relatively long-lasting reductions in resting levels, analogous to the mechanisms in negative priming (see, e.g., Tipper, 1985; cf. Neill, 1989), although this idea is speculative at this point. Our view is that the data are too sparse to permit confidence as to the appropriate avenue for modeling such effects, and we have not implemented them in the independent activation model.

### Development of Meaning Resolution Processes

A valuable method for discovering the parameters affecting any process is to probe development of the process. Because the critical variables early in development may differ from those that are critical later in development, developmental studies can yield important insights into underlying structures or processes. With respect to meaning resolution processes, such investigations provide a unique opportunity to disentangle the relative contribution of meaning access and meaning integration and, in particular, to evaluate the assumption of the independent activation model that access and integration are functionally independent. Evidence that these two processes follow different developmental pathways would suggest that access and integration are not intrinsically dependent on one another. We consider first some of the available developmental evidence on meaning resolution with oral language input and then turn to written language input. In both cases, the independent activation model provides a straightforward account of developmental change.

Swinney and Prather (1989) investigated lexical access in very young children in a task requiring meaning resolution. Children heard a sentence over headphones while looking at a picture. The experimental sentences were ambiguous, owing to the presence of homophonic homographs, and the pictures were related to one meaning of the homograph. The children simply had to decide whether the picture was of an edible object or not; the dependent measure was time to make the edibility response. Very young children (4.0–4.7 years) made faster responses following ambiguous sentences than following unambiguous control sentences (i.e., they showed facilitation) only when the picture was related to the dominant meaning of the homograph. Somewhat older children (4.8–5.6 years) showed facilitation when either meaning of the homograph was suggested by the picture. Swinney and Prather interpreted their data as showing that all children exhaustively access homograph meanings but that “rise-time” (Swinney & Prather, 1989, p. 237) was slower for subordinate meanings for very young children.

Simpson and Foster (1986) had children name words that had been primed by homographs. Students in Grades 2 and 4 showed facilitation for responses to words related to a meaning of the homograph prime (relative to unrelated words), irrespective of the dominance of the meanings (i.e., both dominant and subordinate meanings were

primed). For sixth graders, facilitation occurred only for responses to words related to the dominant meaning. A second experiment manipulated prime–target SOA in order to examine the hypothesis that the age effects might simply be explained by differential speeds of processing. Although Simpson and Foster did find differential speeds of processing for different ages (in support of Swinney & Prather’s [1989] suggestion that “rise-time” varies with age), these differential rates were not responsible for the critical finding that meaning frequency effects develop with age. Simpson and Foster found that, with a 300-msec prime–target SOA, children of all three grades (2, 4, and 6) showed facilitation of responses to related targets over unrelated targets, with no effects of dominance. With an SOA of 750 msec, however, the sixth graders showed facilitation for dominant associates and inhibition for subordinate associates, whereas second and fourth graders continued to exhibit facilitation of responses to both subordinate and dominant meanings. Simpson and Foster concluded that sixth graders were able to use meaning frequency to guide meaning selection but that second and fourth graders were incapable of doing so.

The findings of Swinney and Prather (1989) and Simpson and Foster (1986) form a coherent picture: At very young ages, children demonstrate access of multiple meanings for words, but the processes that allow selection on the basis of meaning frequency (and, it is likely, context; see Simpson & Lorsch, 1983) remain underdeveloped until about the age of 12 years (see Simpson, Krueger, Kang, & Eloffson, 1994, for a similar argument). Very young children activate word meanings somewhat more slowly than do older children and adults. Because dominant meanings are activated quickly and easily from a very young age, the difference in speed of activation has little effect on dominant meaning activation. However, it results in a shift in the pattern of activation for subordinate meanings. At short SOAs, as in the Swinney and Prather study, in which sentences and pictures were simultaneously present, subordinate meanings are not yet activated for very young children, although they are for slightly older children. At intermediate SOAs (such as the 300-msec SOA condition in Experiment 2 of Simpson & Foster), subordinate meanings are active for both young and older children. Finally, at longer SOAs (such as the 750-msec SOA condition in Simpson & Foster’s study), the activation of subordinate meanings has begun to decrease for older children but remains high for younger children.

The data of Swinney and Prather (1989) and Simpson and Foster (1986) can be modeled in the independent activation model simply by assuming that the perceptual input function is roughly half as fast in younger children as in older children and adults (see Kail, 1991a, 1991b, 1992, for relevant data). Early in development, the perceptual input may increase relatively slowly and remain strong for a long period of time. However, as people learn to decode words more quickly, the access system may

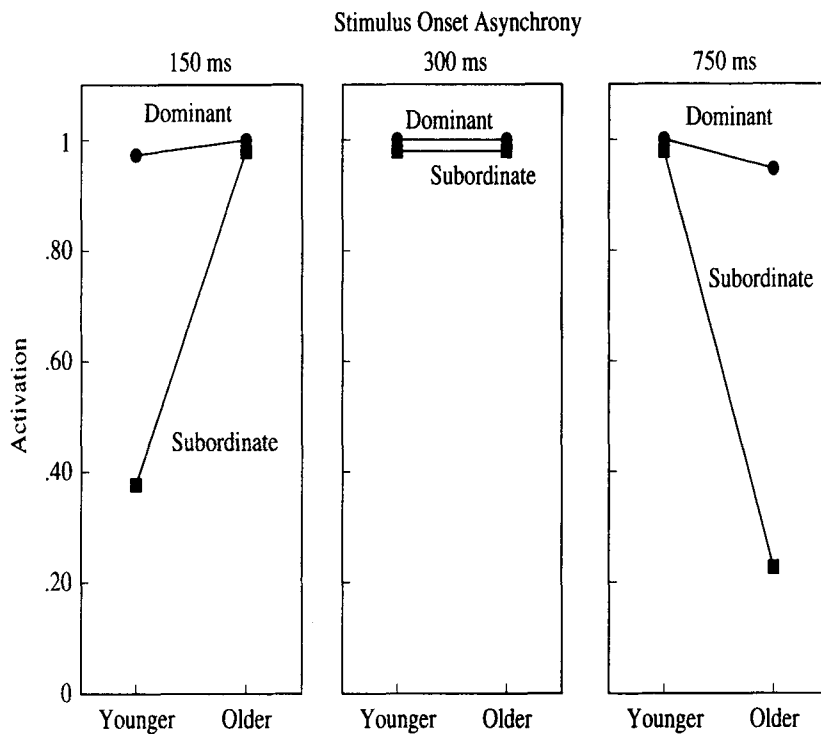


Figure 12. Simulation of development of meaning resolution. The three panels show the predicted activation levels of the dominant and subordinate meanings at stimulus onset asynchronies of 150, 300, and 750 msec. In each case, the activation levels are shown for both the slow access function hypothesized for younger children and the fast access function hypothesized for older children. The activations for subordinate meanings have been reduced by .02 for clarity of presentation.

adapt, to produce a more rapid and short-lasting perceptual input. Presumably, this would allow one to identify words presented in rapid sequence more easily, as in normal conversation or natural reading. These contrasting input functions are shown in Figure 11. The curve for older children is the same as that shown in Figure 3 and was plotted using the default parameters  $s = .01$  and  $d = .001$ ; the curve for younger children was based on rate parameters of half that size ( $s = .005$  and  $d = .0005$ ). Thus, the two curves have identical qualitative characteristics and differ only in the time base. These two input curves were used to simulate the activation of dominant and subordinate meanings for younger and older children. The simulation used the resting levels that were used previously for dominant and subordinate meanings,  $R = -8$  and  $-10$ , and an integration input of zero. The results of this simulation are shown in Figure 12. The left panel depicts the results with a short SOA of 150 msec that is likely to be representative of the method used by Swinney and Prather. The calculations were based on Equation 3, using a fixed value of 150 for  $t$ , but with the different values of  $R$ ,  $d$ , and  $s$  representing dominant and subordinate meanings and younger and older children. The middle and right panels show the simulation of the pattern of data observed in Experiment 2 of Simpson and Foster with SOAs of 300 and 750 msec. The results were

calculated as in the left-hand panel, but with values of 300 and 750, respectively, for  $t$ . The simulation illustrates how simple quantitative changes in processing speed can produce qualitatively different patterns of meaning activation. In keeping with the evidence on the development of meaning resolution, the simulation shows complete activation of the dominant meaning for both older and younger children but slower and longer lasting activation of the subordinate meaning for the younger children.

Despite the important contribution of the research described above, the study of the development of lexical ambiguity processing is impoverished. Little of the research in this area has been methodologically rigorous, and the completed work is far from conclusive. Comprehension of ambiguous stimuli by children younger than 10 or 12 does not involve rapid suppression of inappropriate or subordinate meanings; rather, the inappropriate or subordinate meanings remain activated for at least 750 msec. So, what are the means by which meaning resolution is accomplished in younger children? Perhaps, young children must undertake a laborious context-checking procedure in which all the possible interpretations of a stimulus are checked for plausibility before a decision can be made. Or perhaps integration processes, as well as access processes, are simply delayed in young children.

Any comprehensive account of meaning resolution must involve a systematic investigation of the patterns of access and integration of word meanings over time, across a range of ages, with variations in meaning frequency and context strength. Many central questions are as yet unstudied. Are context strength effects similar in children and adults? Do meaning frequency and context strength effects develop at the same time? An intriguing question is whether access and integration processes develop similarly in the oral lexicon and the print lexicon. The answer to this question would have important implications for arguments about the existence of language processes that are general, rather than modality specific. Furthermore, data designed to answer such questions would provide fertile ground on which to test ambiguity-processing models such as the independent activation model; most salient is the issue of independence of access and integration. If they are indeed independent, their developmental paths would be expected to differ.

### Individual Differences in Meaning Resolution

Another approach to understanding meaning resolution processes is to consider individual differences. Like developmental data, individual difference data can be used to assess the relationship between meaning access and integration. For example, if meaning access covaries with some reader-specific variables, whereas meaning integration covaries with others, it would suggest that access and integration are functionally independent. In addition, an understanding of the variables that are related to meaning resolution processes may provide some insight into the nature of those mechanisms. In the present section, we review evidence that individual differences in processing lexical ambiguity are related to working memory capacity. In particular, we hypothesize that individuals with large working memory capacity use context more effectively. We assume that the representation of prior context is maintained in working memory and that those with a larger capacity may be able to maintain more articulated and elaborated representations of that context. These more extensive representations exert more of an effect on subsequent lexical representations, and, as a consequence, these individuals show stronger effects of context in homograph meaning resolution.

One source of evidence on individual differences in lexical ambiguity processing is a series of experiments by Gernsbacher and her colleagues (Gernsbacher & Faust, 1991a, 1991b; Gernsbacher, Varner, & Faust, 1990). In these studies, the stimuli were sentences containing a lexical ambiguity, followed by test words. The dependent measure was the time to judge the relatedness of the test word to the sentence. The results indicated that readers with high scores on a comprehension test devised by Gernsbacher and Varner (1988) were less affected by contextually inappropriate meanings of the homograph; in other words, high comprehenders were more effective at using context to suppress irrelevant meanings. This conclusion is consistent with the hypothesis that working

memory capacity is related to use of context in meaning resolution, if it is assumed that the high comprehenders identified by Gernsbacher had relatively large working memory capacities.

Another source of evidence for the hypothesis comes from the work of Daneman and Green (1986). They found that high working memory capacity was associated with a superior ability to infer the meanings of unknown words from context. Their study did not involve homographs, but there is a sense in which an unknown word constitutes a form of lexical ambiguity; an unknown word is ambiguous, not because several meanings receive input from access processes, but because *no* meaning receives much input. However, a suitably close synonym of the unknown word could receive a certain amount of contextual support, even though it is inconsistent with the perceptual input. Moreover, if the contextual constraint is sufficiently strong, the reader may discern the intended meaning on the basis of this minimal activation of the synonym. On this view, the extent to which an appropriate meaning can be activated without access input provides an index of the strength of the integration input, since there is no other way for a meaning to become active. Thus, the finding that high-capacity readers can infer an appropriate meaning successfully suggests that, for these readers, a meaning sometimes can be activated purely on the basis of context. Our interpretation is that the results both of Daneman and Green and of Gernsbacher indicate that high-capacity readers show stronger effects of context.

Daneman and Carpenter (1983) and Miyake, Just, and Carpenter (1992, 1994) have also argued that working memory capacity is related to lexical ambiguity resolution. For example, Daneman and Carpenter found that readers with small working memory capacities were less likely to be able to resolve misinterpretations of lexical ambiguities. That is, after reading aloud a sentence containing a homograph whose meanings depended on pronunciation (i.e., the fish meaning of *bass* vs. the guitar meaning), subjects with a small working memory span were less likely to be able to reinterpret the sentence correctly if they had initially mispronounced the homograph. Results such as these have been interpreted in terms of the capacity theory of working memory proposed by Just and Carpenter (1992). According to this view, readers with larger working memory capacities are better able to maintain activation of multiple word meanings and are, therefore, better able to resolve interpretation problems that may arise from ambiguous words. A similar approach to the processing of syntactic ambiguity has been proposed by MacDonald, Just, and Carpenter (1992) and King and Just (1991).

This account of individual differences in ambiguity processing is somewhat different from our conclusions, based on the work of Gernsbacher (Gernsbacher & Faust, 1991a, 1991b; Gernsbacher, Varner, & Faust, 1990) and Daneman and Green (1986). For example, Gernsbacher's results suggest that readers with large working memo-

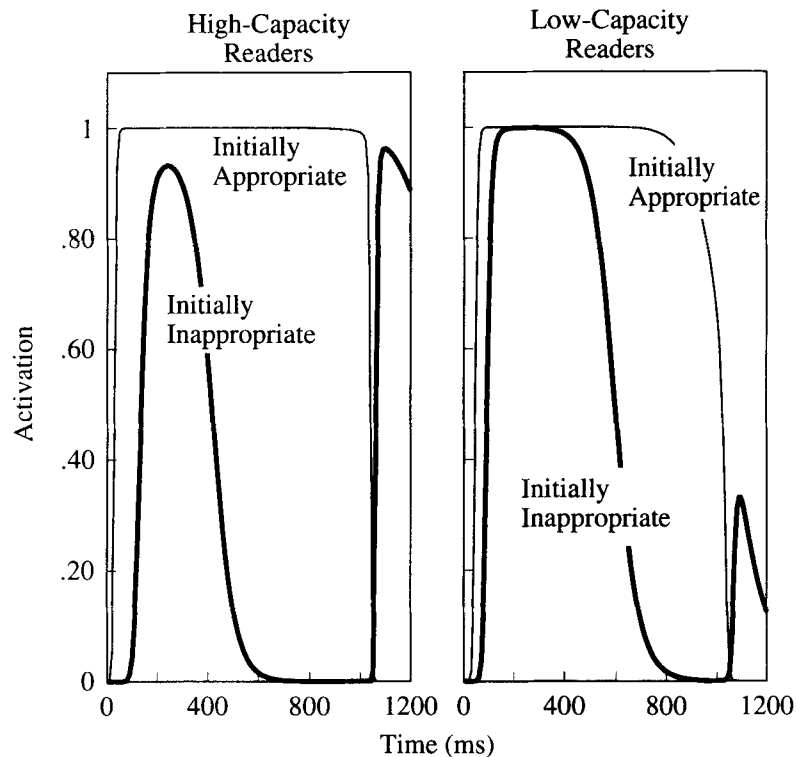


Figure 13. Simulation of individual differences in meaning activation. The left panel depicts the results effect of changing context after 1,000 msec for high-capacity readers with  $I = \pm 4$ . The right panel depicts the predictions for low-capacity readers with weaker integration input ( $I = \pm 2$ ).

ries show greater effects of context (an integration process), whereas the results of Daneman and Carpenter (1983) suggest that readers with large working memories are more successful at activating multiple meanings (an access process). Although it is conceivable that both accounts are correct, we prefer a more parsimonious account: In the Daneman and Carpenter studies, the initial meaning resolution process (based on the perceptual input and prior context) is conflated with recovery processes that intervene when the first interpretation of a homograph proves incompatible with the subsequent context. Thus, both Gernsbacher's evidence on the suppression of inappropriate meanings and Daneman and Carpenter's on recovery suggest that better readers show stronger effects of context. The difference is that the initial suppression of inappropriate meanings is based on prior context, whereas recovery in the Daneman and Carpenter paradigm is based on subsequent context.

This account is illustrated in Figure 13, using the framework of the independent activation model. The curves show the predicted activations of two meanings when prior context supports one meaning of the homograph and subsequent context is consistent with the other meaning. The simulation in the left-hand panel used weak contextual strength in order to produce an exhaustive access pattern of activation. The parameter values were the same as those used in the exhaustive access

simulation shown in Figure 4:  $I = 4$  for the appropriate meaning, and  $I = -4$  for the inappropriate meaning. As in the earlier case, both meanings are activated initially, but the contextually inappropriate meaning is suppressed within a few hundred milliseconds. In order to simulate the effect of subsequent context, the context was gradually changed after 1,000 msec, so that by 1,100 msec, the integration input favored the opposite meaning. In other words, after 1,100 msec,  $I = -4$  for the meaning that was initially appropriate, and  $I = 4$  for the meaning that was initially inappropriate. As is shown in the figure, the change in contextual input successfully suppresses the meaning that was active at the outset and activates the other meaning. This is the pattern of activation that we expect for high-capacity readers on the basis of the results of Daneman and Carpenter (1983).

The curves in the right-hand panel illustrate the patterns of activation we expect for low-capacity readers. Our assumption is that readers with low working memory capacity show weaker effects of context, and the value of  $I$  was reduced from 4 to 2 to reflect this difference; in other respects, the simulation was precisely the same as in the left-hand panel. The figure shows that, even with the reduced effect of context, the exhaustive access pattern is still obtained. Initially, both meanings are activated, and the inappropriate meaning is suppressed within a few hundred milliseconds. However, the

suppression of the inappropriate meaning is somewhat slower than that predicted for high-capacity readers. This aspect of the simulation is similar to the pattern of results obtained by Gernsbacher and colleagues in the relatedness task: Relative to high-capacity readers, low-capacity readers have difficulty suppressing the inappropriate meaning. Crucially, the simulation in the right-hand panel shows a different pattern when the contextual input is changed after 1,000 msec. Unlike high-capacity readers, low-capacity readers fail to strongly reactivate the meaning that is now contextually appropriate. In this case, the access input has declined from its initial peak of around 200–300 msec, and the weaker integration input assumed for these readers is insufficient to compensate for the declining input from the perceptual system. Our interpretation is that these readers would have trouble interpreting the sentence correctly under these circumstances and would be likely to either refixate the homograph or misunderstand the sentence. In contrast, our simulation suggests that high-capacity readers successfully reactivate the correct meaning on the basis of the subsequent context, even without refixating the homograph.

Evidence reported by Miyake et al. (1994) is qualitatively consistent with the account illustrated in Figure 13. In that study, subjects read lexically ambiguous sentences in which homographs were preceded by nonbiasing contexts. Readers with large reading spans spent little extra time reading the subsequent disambiguating region, regardless of whether it was consistent with the dominant or the subordinate meaning of the homograph. In contrast, low-span readers had increased reading times when the subsequent context was consistent with the subordinate meaning of the homograph. Presumably, low-span readers had difficulty reactivating the subordinate meaning and, as a consequence, had difficulty understanding the sentence. Miyake et al. (1994) interpreted this result as suggesting that high-span readers had maintained both possible meanings of the homographs. However, it is also possible that these readers were simply efficient at reactivating subordinate meanings when they were more consistent with the subsequent context, just as in the simulation shown in Figure 13. In other words, the Miyake et al. (1994) results fit with the general conclusion that the advantage of high-capacity readers lies in integration processes, rather than in access processes. It should be pointed out, however, that Miyake et al. (1994) used long delays between the presentation of the homograph and the presentation of the disambiguating context, and this would make it difficult to simulate their results in the independent activation model, using the parametric assumptions adopted here.

The view that readers with large working memory capacities have stronger integration processes also provides a straightforward account of the pattern of results observed by Twilley and Dixon (1993). In a relatedness task similar to that used by Gernsbacher, we had high- and low-capacity readers judge the relatedness of test words to sentences ending in homographs. The test word was related to one meaning of the homograph, and the

sentence was either ambiguous (more than one interpretation was possible) or contextually biased (so that only one interpretation was possible). We found that high-capacity readers were less likely than low-capacity readers to make errors in judging the relatedness of test words to ambiguous sentences, as well as to sentences biased toward a meaning other than that related to the test word. This finding is paradoxical on the view that high-capacity readers simply maintain more alternative interpretations; in that case, one would expect high-capacity readers to make *more* errors than low-capacity readers, when faced with a test word related to an inappropriate meaning, since the high-capacity readers would be more likely to maintain the inappropriate meaning than would the low-capacity readers. Rather, the results are consistent with the view that large working memory capacity is associated with stronger effects of context. This would allow high-capacity readers to activate multiple meanings when they are contextually appropriate and to suppress alternative meanings when they are contextually inappropriate and, hence, to make fewer errors for both ambiguous and inappropriately biased sentences.

## CONCLUSIONS

We have demonstrated that a simple quantitative model, the independent activation model, can account for a large variety of complex patterns of results with only a few parameters and some simple and reasonable assumptions about meaning activation. The model has four substantive assumptions: The effects of prior context are immediate and constant over the bulk of the meaning resolution process; context serves to symmetrically enhance appropriate meanings and suppress inappropriate meanings; perceptual support for both interpretations of a homograph increases to a maximum within about 300 msec and decays relatively slowly thereafter; and the activation of a meaning is based on the sum of the input from access processes and integration processes, together with an input that determines the resting level of activation. A variety of evidence suggests that the parameters of meaning resolution in this model reflect context strength, development, and reading skill.

The success of this model does not lead us to conclude that the model is a sophisticated account of meaning resolution; it is not. In fact, we regard the model as one of the simplest possible accounts that would be capable of generating any quantitative predictions at all. The model simply provides a framework for interpreting results and generating more telling research questions. Instead of an answer to the question of the nature of meaning resolution, it is a starting point from which more precise and extensive investigations of the parameters of meaning resolution can proceed.

As it stands, however, the model provides an interesting resolution to the lexical modularity debate. Historically, it has been argued that early effects of prior context on meaning activation provide support for nonmodular

positions, whereas the failure to find such early effects of context provide evidence for the modular position. By these criteria, recent evidence (that only a single meaning of a homograph is made available relatively quickly in sufficiently constraining contexts) is in favor of the nonmodular position. This has led to the popularity of various context-sensitive models. Although this evidence strongly suggests that meaning access and integration go on in parallel, we argue that it does not necessarily imply that access and integration interact. In particular, in the independent activation model of meaning resolution, it is assumed that meaning access and integration are parallel but independent processes, both providing separate inputs to possible word meanings. Thus, we have demonstrated that interaction between access and integration processes is not necessary for early context effects to be observed. Although our model is nonmodular in the sense that access and integration go on simultaneously, it is a modular model in the sense that access processes do not interact with integration processes.

Our view is that further investigations of lexical ambiguity must go beyond the question of modularity and that research in the area is best served by considering the entire range of processing involved in arriving at an appropriate interpretation of an ambiguous stimulus. Although the question of modularity of access is important, it should not supplant investigation of the broader range of meaning resolution processes. Our review indicates that a number of important questions relevant to the larger question of how meaning resolution is accomplished have already begun to be answered. Meaning activation processes are strongly affected by meaning frequency and context type in a variety of tasks; context strength may be best defined as a combination of effective enhancement of appropriate meanings and effective suppression of inappropriate meanings; integration processes are potentially fast-acting; developmental research has indicated that meaning access processes are automatic even in preaders but that integration processes are not fully developed until late childhood; and variation in access and integration processes appears to be linked to working memory capacity and reading skill. Finding definitive answers to issues such as these provide a broader and more important research agenda than simply deciding whether integration processes can affect meaning activation within the first 300 msec of processing.

#### REFERENCES

- ANDERSON, R. C., & NAGY, W. E. (1991). Word meanings. In P. D. Pearson (Ed.), *Handbook of reading research* (2nd ed., pp. 690-724). New York: Longman.
- ARMSTRONG, S. L., GLEITMAN, L. R., & GLEITMAN, H. (1983). What some concepts might not be. *Cognition*, **13**, 263-308.
- BALOTA, D. A. (1994). Visual word recognition: The journey from features to meaning. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 303-358). New York: Academic Press.
- BALOTA, D. A., FERRARO, F. R., & CONNOR, L. T. (1991). On the early influence of meaning in word recognition: A review of the literature. In P. J. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 187-222). Hillsdale, NJ: Erlbaum.
- BECKER, C. A. (1980). Semantic context effects in visual word recognition: An analysis of semantic strategies. *Memory & Cognition*, **8**, 493-512.
- BESNER, D., & JOORDENS, S. (1995). Wrestling with ambiguity—Further reflections: Reply to Masson and Borowsky (1995) and Rueckl (1995). *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 515-519.
- BRITTON, B. K. (1978). Lexical ambiguity of words used in English text. *Behavior Research Methods & Instrumentation*, **10**, 1-7.
- BUBKA, A., & GORFEIN, D. S. (1989). Resolving semantic ambiguity: An introduction. In D. S. Gorfein (Ed.), *Resolving semantic ambiguity* (pp. 3-12). New York: Springer-Verlag.
- BURGESS, C., & SIMPSON, G. B. (1988). Cerebral hemisphere mechanisms in the retrieval of ambiguous word meanings. *Brain & Language*, **33**, 86-103.
- BUSEY, T. A., & LOFTUS, G. R. (1994). Sensory and cognitive components of visual information acquisition. *Psychological Review*, **101**, 446-469.
- CARPENTER, P. A., & DANEMAN, M. (1981). Lexical retrieval and error recovery in reading: A model based on eye fixations. *Journal of Verbal Learning & Verbal Behavior*, **20**, 137-160.
- CARROLL, J., DAVIES, P., & RICHMAN, B. (1971). *The American Heritage word frequency book*. Boston: Houghton Mifflin.
- COLLINS, A., & LOFTUS, E. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, **82**, 407-428.
- CONRAD, C. (1974). Context effects in sentence comprehension: A study of the subjective lexicon. *Memory & Cognition*, **2**, 130-138.
- COTTRELL, G. W. (1988). A model of lexical access of ambiguous words. In S. I. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Computational, linguistic, and psychological perspectives* (pp. 179-194). San Mateo, CA: Morgan Kaufmann.
- DANEMAN, M., & CARPENTER, P. A. (1983). Individual differences in integrating information between and within sentences. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **9**, 561-583.
- DANEMAN, M., & GREEN, I. (1986). Individual differences in comprehending and producing words in context. *Journal of Memory & Language*, **25**, 1-18.
- DOPKINS, S., MORRIS, R. K., & RAYNER, K. (1992). Lexical ambiguity and eye fixations in reading: A test of competing models of lexical ambiguity resolution. *Journal of Memory & Language*, **31**, 461-476.
- DUFFY, S. A., MORRIS, R. K., & RAYNER, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory & Language*, **27**, 429-446.
- FODOR, J. A. (1983). *Modularity of mind*. Cambridge, MA: MIT Press.
- FORSTER, K. I. (1981). Priming and the effects of sentence and lexical contexts on naming time: Evidence for autonomous lexical processing. *Quarterly Journal of Experimental Psychology*, **33A**, 465-495.
- FORSTER, K. I., & BEDNALL, E. S. (1976). Terminating and exhaustive search in lexical access. *Memory & Cognition*, **4**, 53-61.
- FROST, R., & BENTIN, S. (1992). Processing phonological and semantic ambiguity: Evidence from semantic priming at different SOAs. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 58-68.
- GERNSBACHER, M. A., & FAUST, M. E. (1991a). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 245-262.
- GERNSBACHER, M. A., & FAUST, M. E. (1991b). The role of suppression in sentence comprehension. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 97-128). Amsterdam: Elsevier.
- GERNSBACHER, M. A., & VARNER, K. R. (1988). *The multimedia comprehension battery* (Rep. No. 88-04). Eugene: University of Oregon, Institute of Cognitive and Decision Sciences.
- GERNSBACHER, M. A., VARNER, K. R., & FAUST, M. E. (1990). Investigating differences in general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 430-445.
- GLUCKSBERG, S., KREUZ, F. J., & RHO, S. H. (1986). Context can constrain lexical access: Implications for models of language comprehension. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **12**, 323-335.
- HOGABOAM, T. W., & PERFETTI, C. A. (1975). Lexical ambiguity and sentence comprehension. *Journal of Verbal Learning & Verbal Behavior*, **14**, 265-274.
- HOLMES, V. M. (1979). Accessing ambiguous words during sentence

- comprehension. *Quarterly Journal of Experimental Psychology*, **44A**, 13-35.
- JOORDENS, S., & BESNER, D. (1994). When banking on meaning is not (yet) money in the bank: Explorations in connectionist modeling. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 1051-1062.
- JUST, M. A., & CARPENTER, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review*, **87**, 329-354.
- JUST, M. A., & CARPENTER, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, **99**, 122-149.
- KAIL, R. (1991a). Developmental change in speed of processing during childhood and adolescence. *Psychological Bulletin*, **109**, 490-501.
- KAIL, R. (1991b). Processing time declines exponentially during childhood and adolescence. *Developmental Psychology*, **27**, 259-266.
- KAIL, R. (1992). Processing speed, speech rate, and memory. *Developmental Psychology*, **28**, 899-904.
- KAWAMOTO, A. H. (1988). Interactive processes in the resolution of lexical ambiguity. In S. I. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Computational, linguistic, and psychological perspectives* (pp. 195-228). New York: Morgan Kaufmann.
- KAWAMOTO, A. H. (1993). Nonlinear dynamics in the resolution of lexical ambiguity: A parallel distributed processing account. *Journal of Memory & Language*, **32**, 474-516.
- KELLAS, G., PAUL S. T., MARTIN, M., & SIMPSON, G. B. (1991). Contextual feature activation and meaning access. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 47-71). Amsterdam: Elsevier.
- KING, J., & JUST, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of Memory & Language*, **30**, 580-602.
- KINOSHITA, S. (1985). Sentence context effects on lexically ambiguous words: Evidence for a postaccess inhibition process. *Memory & Cognition*, **13**, 579-595.
- KOHONEN, T. (1987). *Content-addressable memories*. New York: Springer-Verlag.
- KUČERA, J., & FRANCIS, W. M. (1967). *Computational analysis of present day American English*. Providence, RI: Brown University.
- LUCAS, M. M. (1987). Frequency effects on the processing of ambiguous words in sentence contexts. *Language & Speech*, **30**, 25-46.
- LUCAS, M. M. (1999). Context effects in lexical access: A meta-analysis. *Memory & Cognition*, **27**, 385-397.
- LUCE, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1). New York: Wiley.
- MACDONALD, M. C., JUST, M. A., & CARPENTER, P. A. (1992). Working memory constraints on the processing of syntactic ambiguity. *Cognitive Psychology*, **24**, 56-98.
- MACDONALD, M. C., PEARLMUTTER, N. J., & SEIDENBERG, M. S. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review*, **101**, 676-703.
- MARSLÉN-WILSON, W. (1989). Access and integration: Projecting sound onto meaning. In W. Marslen-Wilson (Ed.), *Lexical representation and process* (pp. 3-24). Cambridge, MA: MIT Press.
- MASSON, M. E. J. (1991). A distributed memory model of context effects in word identification. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 233-263). Hillsdale, NJ: Erlbaum.
- MASSON, M. E. J., & BOROWSKY, R. (1995). Unsettling questions about semantic ambiguity in connectionist models: Comment on Joordens and Besner (1994). *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 509-514.
- MASSON, M. E. J., & FREDMAN, L. (1990). Fluent identification of repeated words. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 355-373.
- MCCLELLAND, J. L. (1987). The case for interactionism in language processing. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 3-35). Hillsdale, NJ: Erlbaum.
- MCCLELLAND, J. L. (1991). Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology*, **23**, 1-44.
- MCCLELLAND, J. L., & ELMAN, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, **18**, 1-86.
- MCCLELLAND, J. L., & KAWAMOTO, A. (1986). Mechanisms of sentence processing: Assigning roles to constituents of sentences. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1. Foundations* (pp. 272-325). Cambridge, MA: MIT Press.
- MCCLELLAND, J. L., & RUMELHART, D. E. (1981). An interactive model of context effects in letter perception: Pt 1. An account of basic findings. *Psychological Review*, **88**, 375-407.
- MCKOON, G., & RATCLIFF, R. (1992). Inference during reading. *Psychological Review*, **99**, 440-466.
- MEDIN, D. L., & SCHAFFER, M. M. (1978). A context theory of classification learning. *Psychological Review*, **85**, 207-238.
- MEYER, D. E., & SCHVANEVELDT, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, **90**, 227-234.
- MIYAKE, A., JUST, M. A., & CARPENTER, P. A. (1992, November). *Working memory constraints on the maintenance of multiple interpretations of lexical ambiguities*. Paper presented at the 33rd Annual Meeting of the Psychonomic Society, St. Louis.
- MIYAKE, A., JUST, M. A., & CARPENTER, P. A. (1994). Working memory constraints on the resolution of lexical ambiguity: Maintaining multiple interpretations in neutral contexts. *Journal of Memory & Language*, **33**, 175-202.
- MORTON, J. (1969). The interaction of information in word recognition. *Psychological Review*, **76**, 165-178.
- MORTON, J. (1979). Word recognition. In J. Morton & J. C. Marshall (Eds.), *Psycholinguistics 2: Structures and processes* (pp. 107-156). London: Paul Elek.
- MURPHY, G. L. (1991). Meaning and concepts. In P. J. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 11-35). Hillsdale, NJ: Erlbaum.
- MURPHY, G. L., & MEDIN, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, **92**, 289-316.
- NEELY, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory & Cognition*, **4**, 648-654.
- NEELY, J. H. (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.
- NEILL, W. T. (1989). Lexical ambiguity and context: An activation-suppression model. In D. S. Gorfein (Ed.), *Resolving semantic ambiguity* (pp. 63-83). New York: Springer-Verlag.
- NEILL, W. T., HILLIARD, D. V., & COOPER, E. (1988). The detection of lexical ambiguity: Evidence for context-sensitive parallel access. *Journal of Memory & Language*, **27**, 279-287.
- ODEN, G. C., & MASSARO, D. W. (1978). Integration of featural information in speech perception. *Psychological Review*, **85**, 172-191.
- ODEN, G. C., & SPIRA, J. L. (1983). Influence of context on the activation and selection of ambiguous word senses. *Quarterly Journal of Experimental Psychology*, **35A**, 51-64.
- 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.
- PAUL, S. T., KELLAS, G., MARTIN, M., & CLARK, M. B. (1992). Influence of contextual features on the activation of ambiguous word meanings. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 703-717.
- POSNER, M. I., & SNYDER, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55-85). Hillsdale, NJ: Erlbaum.
- RAYNER, K., & DUFFY, S. A. (1986). Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory & Cognition*, **14**, 191-201.
- RAYNER, K., & FRAZIER, L. (1989). Selection mechanisms in reading lexically ambiguous words. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 779-790.
- RAYNER, K., PACTH, J. M., & DUFFY, S. A. (1994). Effects of prior encounter and global discourse bias on the processing of lexically ambiguous words: Evidence from eye fixations. *Journal of Memory & Language*, **33**, 527-544.



- RAYNER, K., & POLLATSEK, A. (1989). *The psychology of reading*. Englewood Cliffs, NJ: Prentice-Hall.
- REDER, L. M. (1983). What kind of pitcher can a catcher fill? Effects of priming in sentence comprehension. *Journal of Verbal Learning & Verbal Behavior*, **22**, 189-202.
- ROSCH, E. H. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, **104**, 192-233.
- RUECKL, J. G. (1995). Ambiguity and connectionist networks: Still settling into a solution—Comment on Joordens and Besner (1994). *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 501-508.
- RUMELHART, D. E., HINTON, G. E., & WILLIAMS, R. J. (1986). Learning representations by back-propagating errors. *Nature*, **323**, 533-536.
- SCHVANEVELDT, R. W., MEYER, D. E., & BECKER, C. A. (1976). Lexical ambiguity, semantic context, and visual word recognition. *Journal of Experimental Psychology: Human Perception & Performance*, **2**, 243-256.
- SEIDENBERG, M. S. (1985). The time course of information activation and utilization in visual word recognition. In D. Besner, T. Waller, & G. E. MacKinnon (Eds.), *Reading research: Advances in theory and practice* (Vol. 5, pp. 199-252). Orlando, FL: Academic Press.
- SEIDENBERG, M. S., & MCCLELLAND, J. L. (1989). A distributed, developmental model of visual word recognition and naming. *Psychological Review*, **96**, 523-568.
- SEIDENBERG, M. S., TANENHAUS, M. K., LEIMAN, J. M., & BIENKOWSKI, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology*, **14**, 489-537.
- SERENO, S. C. (1995). Resolution of lexical ambiguity: Evidence from an eye movement priming paradigm. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 582-595.
- SERENO, S. C., & RAYNER, K. (1993, November). *Lexical ambiguity resolution: Evidence from an eye movement priming paradigm*. Paper presented at the 34rd Annual Meeting of the Psychonomic Society, Washington, DC.
- SIMPSON, G. B. (1981). Meaning dominance and semantic context in the processing of lexical ambiguity. *Journal of Verbal Learning & Verbal Behavior*, **20**, 130-136.
- SIMPSON, G. B. (1984). Lexical ambiguity and its role in models of word recognition. *Psychological Bulletin*, **96**, 316-340.
- SIMPSON, G. B. (1994). Context and the processing of ambiguous words. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 359-374). New York: Academic Press.
- SIMPSON, G. B., & BURGESS, C. (1985). Activation and selection processes in the recognition of ambiguous words. *Journal of Experimental Psychology: Human Perception & Performance*, **11**, 28-39.
- SIMPSON, G. B., & FOSTER, M. R. (1986). Lexical ambiguity and children's word recognition. *Developmental Psychology*, **22**, 147-154.
- SIMPSON, G. B., & KANG, H. (1994). Inhibitory processes in the recognition of homograph meanings. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory mechanisms in attention, memory, and language* (pp. 359-381). New York: Academic Press.
- SIMPSON, G. B., & KELLAS, G. (1989). Dynamic contextual processes and lexical access. In D. S. Gorfein (Ed.), *Resolving semantic ambiguity* (pp. 40-56). New York: Springer-Verlag.
- SIMPSON, G. B., & KRUEGER, M. A. (1991). Selective access of homograph meanings in sentence context. *Journal of Memory & Language*, **30**, 627-643.
- SIMPSON, G. B., KRUEGER, M. A., KANG, H., & ELOFSON, A. C. (1994). Sentence context and meaning frequency effects in children's processing of ambiguous words. *Journal of Research in Reading*, **17**, 62-72.
- SIMPSON, G. B., & LORSBACH, T. C. (1983). The development of automatic and conscious components of contextual facilitation. *Child Development*, **54**, 760-772.
- SMITH, E. E., SHOBEN, E. J., & RIPS, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, **81**, 214-241.
- STANOVICH, K. E. (1991). Word recognition: Changing perspectives. In R. Barr, M. L. Kamil, P. Mosenthal, & P. D. Pearson (Eds.), *Handbook of reading research* (Vol. 2, pp. 418-451). New York: Longman.
- STANOVICH, K. E., & WEST, R. F. (1983). On priming by a sentence context. *Journal of Experimental Psychology: General*, **112**, 1-36.
- SWINNEY, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning & Verbal Behavior*, **18**, 645-659.
- SWINNEY, D. A., & PRATHER, P. (1989). On the comprehension of lexical ambiguity by young children: Investigations into the development of mental modularity. In D. S. Gorfein (Ed.), *Resolving semantic ambiguity* (pp. 225-238). New York: Springer-Verlag.
- TABOSSI, P. (1988). Accessing lexical ambiguity in different types of sentential context. *Journal of Memory & Language*, **27**, 324-340.
- TABOSSI, P. (1989). What's in a context? In D. S. Gorfein (Ed.), *Resolving semantic ambiguity* (pp. 25-39). New York: Springer-Verlag.
- TABOSSI, P. (1991). Understanding words in context. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 1-22). Amsterdam: Elsevier.
- TABOSSI, P., COLOMBO, L., & JOB, R. (1987). Accessing lexical ambiguity: Effects of context and dominance. *Psychological Research*, **49**, 161-167.
- TABOSSI, P., & ZARDON, F. (1993). Processing ambiguous words in context. *Journal of Memory & Language*, **32**, 359-372.
- TANENHAUS, M. K., LEIMAN, L. M., & SEIDENBERG, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning & Verbal Behavior*, **18**, 427-440.
- TILL, R. E., MROSS, E. F., & KINTSCH, W. (1988). Time course of priming for associate and inference words in a discourse context. *Memory & Cognition*, **16**, 283-298.
- TIPPER, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *Quarterly Journal of Experimental Psychology*, **37A**, 571-590.
- TOWNSEND, J. T. (1976). Serial and within-stage independent parallel model equivalence on the minimum completion time. *Journal of Mathematical Psychology*, **14**, 219-238.
- TWILLEY, L. C. (1996). *The role of competing meanings in homograph meaning resolution*. Unpublished doctoral dissertation, University of Alberta.
- TWILLEY, L. C., & DIXON, P. (1993, July). *Working memory as a predictor of errors in lexical ambiguity processing*. Poster presented at the 3rd Annual Conference of the Canadian Society for Brain, Behaviour, and Cognitive Science, Toronto.
- TWILLEY, L. C., DIXON, P., TAYLOR, D., & CLARK, K. (1994). University of Alberta norms of relative meaning frequency for 566 homographs. *Memory & Cognition*, **22**, 111-126.
- VAN PETTEN, C., & KUTAS, M. (1987). Ambiguous words in context: An event-related potential analysis of the time course of meaning activation. *Journal of Memory & Language*, **26**, 188-208.
- VAN PETTEN, C., & KUTAS, M. (1988). Tracking the time course of meaning activation. In S. I. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Computational, linguistic, and psychological perspectives* (pp. 431-475). San Mateo, CA: Morgan Kaufmann.
- WALTZ, D. L., & POLLACK, J. B. (1985). Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science*, **9**, 51-74.

## APPENDIX

### Derivation of Meaning Activation With Distributed Representations

We assume that the vectors representing target meanings consist of patterns of 0s and 1s, where 1 would mean that a particular semantic feature is present in the interpretation of a homograph, whereas 0 would mean that it is absent. After sufficient training, the output of the system would be arbitrarily close to the target values. However, because of the nature of Equation 5, the system can only approximate these values; producing an output of precisely 0 or 1 would require that the net input to the unit be infinite. Consequently, we assume that the

training targets for the output meaning vectors consist of patterns of  $\epsilon$  and  $1 - \epsilon$  (rather than 0 and 1), where  $\epsilon$  is a suitably small value. The system would be able to match these modified targets with arbitrary precision after sufficient training. Finally, we assume that the resting levels are sufficient to produce approximately 0 output from each unit when there is no input. (In a more complete development,  $R$  would be learned and would reflect the frequency of various possible meanings.) In other words,

$$R_i = \frac{1}{2} \ln \left( \frac{\epsilon}{1 - \epsilon} \right) = R^* \quad (\text{A1})$$

(In this and the following equations, an asterisk superscript is used to denote an input vector that would produce a given pattern of  $\epsilon$  and  $1 - \epsilon$  after logistic scaling.)

The effect of training with the delta rule on the connection weights is well established (e.g., Kohonen, 1987). For example, if a linear associator is trained to produce output  $B$  when presented with input  $C$ , then the matrix of connection weights has the form

$$\frac{1}{\|C\|^2} BC^T \quad (\text{A2})$$

The present model is similar to a linear associator at the level of the net input to meaning level; that is, Equation 6 is formally equivalent to the assumptions made in a linear associator. This relationship allows us to derive the form of the connection weight matrix in a simple fashion. If the training is successful, the net inputs would be sufficient to produce patterns of  $\epsilon$  and  $1 - \epsilon$  that approximate the correct target values. We use an asterisk superscript to indicate such net input values, so that a net input vector of  $T_1^*$  produces approximately  $T_1$  and a net input vector of  $T_2^*$  produces approximately  $T_2$ . These values are given by

$$T_{ij}^* = \frac{1}{2} \ln \left[ T_{ij} \left( \frac{1 - \epsilon}{\epsilon} \right) + (1 - T_{ij}) \left( \frac{\epsilon}{1 - \epsilon} \right) \right] - R_{ij} \quad (\text{A3})$$

If the system produces approximately the correct target values after training, this means that  $W(A + I_1)$  produces  $T_1^*$ ,  $W(A + I_2)$  produces  $T_2^*$ , and  $WA$  produces  $(T_1^* + T_2^*)/2$  as net inputs. Equation A2 can then be used to identify the value of  $W$  that would be learned to produce these values. In particular, if  $A$ ,  $I_1$ , and  $I_2$  are orthogonal (independent), the connection weight matrix would be

$$W = \frac{1}{2} \left[ \frac{1}{\|A\|^2} (T_1^* + T_2^*) A^T + \frac{1}{\|I_1\|^2} (T_1^* + T_2^*) I_1^T + \frac{1}{\|I_2\|^2} (T_1^* + T_2^*) I_2^T \right] \quad (\text{A4})$$

Comparable results apply if the system is simultaneously trained on any number of homographs and unambiguous words, as long as the input vectors are linearly independent.

Having learned these connection weights, the system responds in exactly the same way as the localist representation to access and integration inputs of various strengths. When presented with access input of strength  $\alpha$  and integration input of strength  $\beta$  in favor of meaning 1 (i.e.,  $\alpha A + \beta I_1$ ), the net input would be  $N = W(\alpha A + \beta I_1)$ . Using Equation 10, this becomes

$$N = \frac{\alpha}{2} (T_1^* + T_2^*) + \frac{\beta}{2} (T_1^* + T_2^*) \quad (\text{A5})$$

This expression can be used to derive the net input to units that are active in the target pattern—that is,  $N_j$ , when  $T_{1j} = 1$ . Because  $T_1$  and  $T_2$  are orthogonal, it follows that whenever  $T_{1j} = 1$ ,  $T_{2j} = 0$ , and consequently, from Equations A1 and A3,

$$T_{1j}^* = \ln \frac{1 - \epsilon}{\epsilon}$$

and  $T_{2j}^* = 0$ . Thus, whenever  $T_{1j} = 1$ ,  $N_j$  can be written as  $A^* + I^*$ , where

$$A^* = \frac{\alpha}{2} \ln \frac{1 - \epsilon}{\epsilon}$$

and

$$I^* = \frac{\beta}{2} \ln \frac{1 - \epsilon}{\epsilon}.$$

We assume that the activated features of the distributed representation are available directly to working memory and comprehension processes. However, as an index of the activation of the system, we can measure the extent to which a particular pattern is active with the normalized dot product of the target pattern and the output of the system. In the independent activation model, we allow several patterns to be active at the same time. Consequently, an appropriate index of a pattern's activation requires normalizing the dot product by the length of the target pattern alone, rather than by the length of both the target and the output. That is, our measure of the activation of meaning  $i$  would be

$$\frac{T_i \cdot M}{\|T_i\|^2} = \frac{\sum_j T_{ij} \left[ \frac{1}{1 + e^{-2(N_j + R_j)}} \right]}{\sum_j T_{ij}^2} \quad (\text{A6})$$

Because  $T_{ij}$  can only have a value of 0 or 1, the numerator summation terms need only be evaluated for  $T_{ij} = 1$ . Consequently, the activation of meaning 1 when presented with input  $\alpha A + \beta I_1$  is

$$\begin{aligned} \frac{T_1 \cdot M}{\|T_1\|^2} &= \frac{\sum_j T_{1j} \left[ \frac{1}{1 + e^{-2(A^* + I^* + R^*)}} \right]}{\sum_j T_{1j}^2} \\ &= \frac{\left[ \frac{1}{1 + e^{-2(A^* + I^* + R^*)}} \right] \sum_j T_{1j}}{\sum_j T_{1j}^2} \\ &= \frac{1}{1 + e^{-2(A^* + I^* + R^*)}}. \end{aligned} \quad (\text{A7})$$

Similarly, it can be shown that the activation of meaning 2 when presented with  $\alpha A + \beta I_1$  is

$$\frac{T_2 \cdot M}{\|T_2\|^2} = \frac{1}{1 + e^{-2(A^* - I^* + R^*)}} \quad (\text{A8})$$

These activations are identical to those assumed in the localist representation in Equation 1.