Journal of Experimental Psychology:
Learning, Memory, and Cognition
1983, Vol. 9, No. 2, 312-327

0278-7393/83/0902-0312$00.75

Backward Masking, the Suffix Effect, and Preperceptual Storage

Howard J. Kallman
State University of New York at Albany

Dominic W. Massaro
Program in Experimental Psychology
University of California at Santa Cruz

This article considers the use of auditory backward recognition masking (ABRM) and stimulus suffix experiments as indexes of preperceptual auditory storage. In the first part of the article, two ABRM experiments that failed to demonstrate a mask disinhibition effect found previously in stimulus suffix experiments are reported. The failure to demonstrate mask disinhibition is inconsistent with an explanation of ABRM in terms of lateral inhibition. In the second part of the article, evidence is presented to support the conclusion that the suffix effect involves the contributions of later processing stages and does not provide an uncontaminated index of preperceptual storage. In contrast, it is claimed that ABRM experiments provide the most direct index of the temporal course of perceptual recognition. Partial-report tasks and other paradigms are also evaluated in terms of their contributions to an understanding of preperceptual auditory storage. Differences between interruption and integration masking are discussed along with the role of preperceptual auditory storage in speech perception.

Central to information-processing models is the concept of a preperceptual or precategorical storage structure. This structure is assumed to hold the information necessary for recognition. In most models, it is assumed that the sensory system transduces stimulus events into sensory features that are held for a short time in the preperceptual store. The analysis of the featural information takes time and leads to perceptual recognition of the stimulus. Most of the research and theoretical development concerning preperceptual storage has been in the domain of visual information processing (for recent reviews see Coltheart, 1980; Long, 1980). Within the past decade, however, there has been an increasing number of studies of auditory information processing, and this article addresses the nature of preperceptual auditory storage.

This research was supported by National Institute of Mental Health Grant MH36334 to Dominic W. Massaro and National Institute of Mental Health Grant MH36455 to Howard J. Kallman. We thank two anonymous reviewers, Robert Crowder, and Richard Shiffrin for their comments on an earlier version of the article.

Requests for reprints should be sent to Howard J. Kallman, Department of Psychology, State University of New York at Albany, 1400 Washington Avenue, Albany, New York 12222, or Dominic W. Massaro, Program in Experimental Psychology, University of California, Santa Cruz, California 95064.

The properties of the preperceptual auditory store have been studied using a backward recognition masking task. This task involves asking a listener to identify a target sound when a second sound of equal loudness is presented shortly after the target. Identification of the target sound's qualities (e.g., its pitch or duration) is impaired if the second sound is presented within approximately 250 msec of target tone onset (Hawkins & Presson, 1977; Kallman & Massaro, 1979; Massaro, 1972a, 1975a; Massaro & Idson, 1976). Within this 250-msec range, accuracy of identification generally improves with increases in the stimulus-onset asynchrony (SOA) separating the two sounds. This result has obtained in experiments using absolute identification, two-alternative forced-choice, same–different, and relative judgment tasks with both fixed and variable standards (Kallman & Massaro, 1979; Massaro, 1970a, 1975b; Massaro & Idson, 1977). The effect of the second sound on recognition of the first is termed *auditory backward recognition masking* (ABRM); the second sound is referred to as the *backward mask*. The results of the backward recognition masking experiments have been interpreted as suggesting that the readout of information from the preperceptual auditory store is completed within approximately 250 msec (Massaro, 1972a, 1975a, 1975b); it has also been argued that the results of ABRM experiments suggest that the duration of preperceptual auditory storage is approximately 250 msec, although this claim is controversial and has been challenged by the results of stimulus suffix experiments.

In a typical stimulus suffix experiment, a list of words is presented auditorily and the task is to recall the words in order. With auditory presentation of the list, recall of the terminal item is more accurate than recall of the immediately preceding items. However, if a redundant auditory item (e.g., the word *zero*) follows the terminal item on each list by approximately 2 sec or less, and recall of the redundant item (the stimulus suffix) is not required, the recency effect is diminished if not eliminated altogether. Crowder (1978) and his colleagues (Crowder & Morton, 1969; Morton, Crowder, & Prussin, 1971) have argued from the results of stimulus suffix experiments that unanalyzed features are held in preperceptual (or, precategorical) acoustic storage for at least 2 sec.

Crowder (1978) has suggested that the basic mechanism underlying the stimulus suffix effect is the same as that responsible for ABRM, namely, replacement of the preperceptual representation of a sound by a later sound's representation. The recent finding that the stimulus suffix effect is smaller with three successive suffixes rather than one (Crowder, 1978; Morton, 1976) led Crowder to propose a lateral inhibition model of backward masking, which is essentially a variant of the view that a backward mask replaces an earlier presented sound's features in preperceptual storage. According to Crowder's (1978) model, "the form of representation of information in [precategorical acoustic storage] is an entry on a matrix organized by time of arrival and, orthogonally, physical channel properties" (p. 517). Subsequent stimuli serve to inhibit entries on the matrix, and the degree of inhibition depends on the physical similarities of the stimuli and their temporal proximity. The finding that three suffixes result in less of an effect than does one is explained by Crowder by assuming that the later arriving suffixes inhibit the representation of the first suffix and in so doing disinhibit the representations of the terminal item. Although there are other explanations for the performance differences found with three suffixes versus one (Kahneman & Henik, 1981; Morton, 1976), this finding is an interesting empirical result, and we wondered whether it would obtain in an ABRM experiment as Crowder's theory would predict. Thus, in the first part of this article, we report two experiments that were designed to test whether a mask disinhibition effect would be found in an ABRM task. Following our report of the two experiments, we elaborate our views on the general issue of preperceptual storage and its relationship to the stimulus suffix effect, ABRM, and other experimental paradigms.

A Test of Mask Disinhibition in Backward Masking

Experiment 1

Method

Subjects. Ten students at the University of Wisconsin with no history of hearing disorders participated for 4 consecutive days to obtain extra credit in an introductory psychology course. Subjects were tested in groups of four or fewer and were seated in individual sound-attenuated rooms.

Stimuli and procedure. During the initial practice phase of the experiment, subjects participated in two blocks of 50 no-mask trials. The subject's task was to indicate by pressing one of two buttons whether a 20-msec test tone was higher or lower in pitch than a 20-msec, 800 Hz standard tone that preceded the test tone by 500 msec of silence. The test tone was either 770 or 830 Hz. Subjects were given as much time as necessary to respond. Two hundred and fifty msec after all the subjects had responded, feedback was presented by illuminating for 250 msec the symbol *H* or *L*, depending on whether the frequency of the test tone was greater or less than that of the standard. Following feedback, there was an intertrial interval of 1.25 sec.

Following the initial practice trials, which were designed to allow practice on a relatively simple pitch discrimination task, subjects participated in backward mask trials. Similar to the practice trials, the subject's task was to determine whether the 20-msec test tone was higher or lower in pitch than a 20-msec, 800 Hz standard tone that preceded the test tone by 500 msec of silence. However, on most of the trials (mask trials), one or three 20-msec, 800 Hz backward masking tones followed the test tone after a variable silent interstimulus interval (ISI). On mask trials the ISI separating offset of the test tone and onset of the following mask was either 10, 20, 40, 80, 160, 250, or 350 msec; on one eighth of the trials no masking tone was presented. The mask type variable (five levels) was defined by whether one backward mask followed the test tone or three masks, in which case the three 20-msec masks were each temporally separated on a trial by either 20, 50, 100, or 300 msec of silence. Given that temporal proximity of the masks is critical for disinhibition, we might expect any advantage in performance with three masks to decrease with increases in the intermask interval (IMI); in any case, the use of a wide range of IMIs maximizes the probability that we would capture the period of disinhibition if, in fact, one exists in ABRM. Within each block of trials, the 80 experimental conditions (2 × 5 × 8 [Test Tone Frequency × Mask Type × ISI]) were sampled randomly without replacement.

The frequencies of the test tones were adjusted between blocks to keep overall performance of as many subjects as possible at approximately 75% correct. On each trial, the test tone was either 800 + Δf Hz or 800 − Δf Hz. To establish reasonable Δf values for each group of subjects, approximately 12 blocks of 50 experimental trials each were presented on Day 1, and subject performance was carefully monitored to allow adjustment of the Δfs from block to block. On subsequent days of the experiment, 2 blocks of 330 trials were presented (the first 10 trials were practice). Performance on these blocks was monitored, and small adjustments in Δf were made between blocks when warranted. Only data from Days 2, 3, and 4 were included in data analysis. There were thus 24 observations per subject at each of the 80 experimental conditions.

All tones in the experiment were sine waves generated by a Wavetek Model 155 digitally controlled oscillator and amplified through a McIntosh Model MC-50 amplifier. Tones were delivered binaurally to subjects at 78 dB SPL through Grason-Stadler TDH-49 headphones. The tones began at the zero crossing and reached maximal intensity in one fourth of a cycle. Feedback was presented over a visual display of Monsanto MDA-III light-emitting diodes. Experimental events and data collection were controlled by a PDP-8/L computer.

Results

The mean Δf used to keep overall performance at approximately 75% correct was 18 Hz. For the different groups of subjects this value ranged from 8 to 27 Hz. Within a group of subjects, the range of Δf values used across blocks of trials on Days 2, 3, and 4 never exceeded 10 Hz.

Figure 1 presents percentage correct as a function of ISI under the five mask-type conditions. The percentage correct data were subjected to analyses of variance. An initial analysis was performed to evaluate the effect of ISI, test tone (i.e., relatively high or low), and mask type, which were within-subject variables. As can be seen in Figure 1, overall performance improved dramatically with increases in ISI, $F(7, 63) = 94.16$, $p < .001$. Thus the present experiment provided a con-

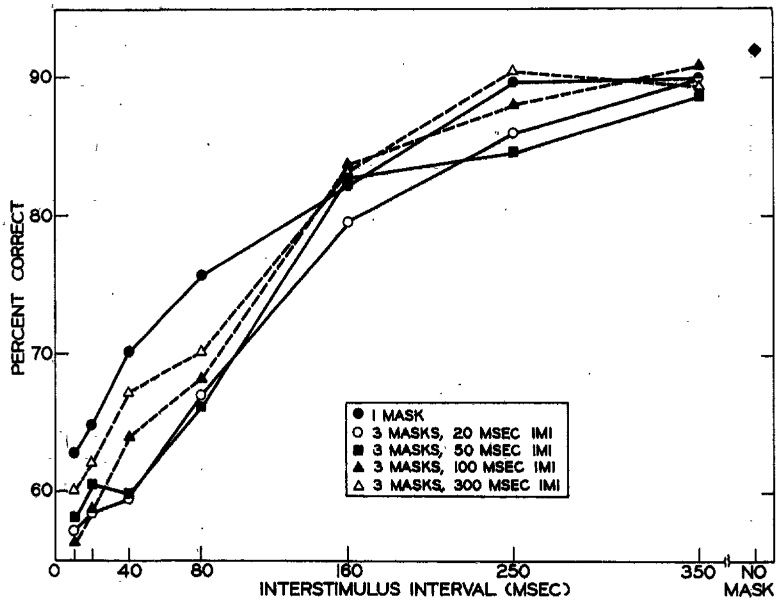

Figure 1. The percentage of correct identifications of the test tones as a function of the duration of the silent interstimulus interval in Experiment 1. (Mask condition is the curve parameter. IMI = intermask interval).

vincing demonstration of backward recognition masking. Because mask type was a dummy variable under the no-mask level of ISI, a subsequent analysis was performed in which the no-mask level of ISI was excluded. Of particular interest was a significant effect of mask type, $F(4, 36) = 8.57, p < .001$. Percentage correct was higher when one mask was presented (76.5%) rather than three masks (73.1%); a specific comparison testing this difference was significant, $F(1, 36) = 19.08, p < .001$. Although the interaction between mask type and ISI was not significant, $F(24, 216) = 1.35, p > .10$, the advantage of the one-mask condition over the three-mask conditions appears primarily at the short ISIs. But note that the direction of the effect of mask type is opposite to that predicted by Crowder's lateral inhibition model of backward masking.

The mean percentages correct under the three-mask conditions were 71.1, 72.4, 73.8, and 74.9, for IMIs of 20, 50, 100, and 300 msec, respectively. The remaining orthogonal tests indicated that although the difference between the 20- and 50-msec IMIs was not significant, $F(1, 36) = 1.40, p > .10$, nor was the difference between the 100- and 300-msec IMIs, $F(1, 36) = 1.33, p > .25$, the percentages correct were higher for the longer IMIs (100 and 300 msec) than the shorter intervals (20 and 50 msec), $F(1, 36) = 12.48, p < .01$.

There were no other significant effects except for an interaction between test tone and mask type, $F(4, 36) = 2.68, p < .05$, which reflected a bias (of approximately 5%) to respond "high" to the test tone when it was followed by three masks but not when it was followed by a single mask.

Discussion

Experiment 1 was unsuccessful in demonstrating better recognition when three backward masks rather than one backward mask followed a test tone. In fact, performance was slightly worse when three masks followed the test tone. Although Crowder contrasted a single suffix with three suffixes to test a lateral inhibition model of backward masking, it might be more appropriate to contrast a one-mask condition to a two-mask condition. It could be argued that the third mask in Experiment 1 inhibited the second mask, which thus minimized the degree to which the second mask could inhibit the first mask. If only two masks are presented, the second mask would not be subject to inhibition from subsequent masks and thus would maximally inhibit the first mask. A two-mask condition was thus contrasted with a one-mask condition in Experiment 2.

Experiment 2

Method

Subjects. Nine previously untested subjects from the same pool as in Experiment 1 participated for 4 consecutive days under conditions nearly identical to those of Experiment 1. In addition, one of the authors, Kallman, served as a subject.

Stimuli and procedure. The experimental procedure and design were identical to that of Experiment 1 except that rather than constrasting the one-mask condition with three-mask conditions, the one-mask condition was contrasted with various two-mask conditions. As in Experiment 1, the IMI when more than one mask was presented was either 20, 50 100, or 300 msec.

Results

The mean Δf used to keep overall performance at approximately 75% correct was 15 Hz. This value ranged from 7 to 38 Hz for different groups of subjects. Within a group of subjects, the range of Δf values used across experimental blocks never exceeded 10 Hz.

Figure 2 presents percentage correct as a function of ISI under the five mask-type conditions. Replicating Experiment 1, performance improved about 30% with increases in ISI, $F(4, 36) = 56.25, p < .001$. Mask type was a significant factor, $F(4, 36) = 4.22, p < .01$, although the overall advantage of the one-mask over the two-mask condition was slight (73.8% vs. 72.4% correct), $F(1, 36) = 2.13, p > .10$. Furthermore, the data for the two-mask conditions were less systematic than in Experiment 1; percentages correct were 70.7, 73.2, 71.8, and 74.0 for IMIs of 20, 50, 100, and 300 msec, respectively. As in Experiment 1, the interaction between mask type and ISI was not significant, $F(24, 216) = 1.28, p > .10$. The main point of interest is that, contrary to the prediction of Crowder's lateral inhibition model, performance under the two-mask condition was

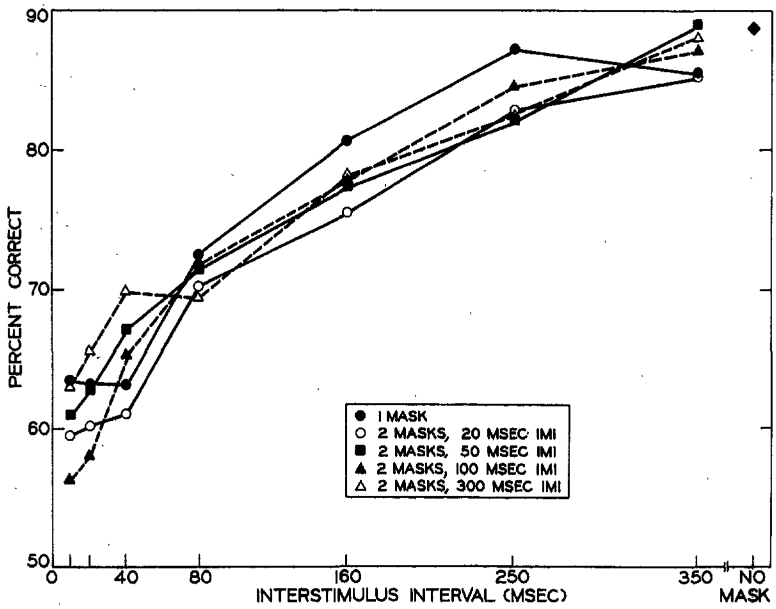

Figure 2. The percentage of correct identifications of the test tones as a function of the duration of the interstimulus interval in Experiment 2. (Mask conditions is the curve parameter. IMI = intermask interval.)

clearly not better than the one-mask condition.

The interaction between test tone and mask was significant, $F(4, 36) = 3.03$, $p < .05$. This result reflected a slight bias (3% to 4%) to respond low when the test tone was followed by either a single mask or by two masks having an IMI of 300 msec and a slight bias (on average, 2%) to respond high under the other mask conditions. The only other significant effect was the interaction between test tone, mask type, and ISI, $F(24, 216) = 2.27$, $p < .01$, but this interaction was difficult to interpret.

Discussion

As in Experiment 1, increasing the number of masking tones did not improve performance. There is thus no evidence to suggest that in backward recognition masking experiments, subsequent masks inhibit the effect of the first mask.

Performance by Sensitive Subjects

Although parenthetic to the main concerns of the present study, Cudahy and Leshowitz (1974) argued that subjects in Massaro's demonstrations of ABRM were insensitive on the pitch discrimination task. They implied that this sheds doubt on the significance of the backward recognition masking functions. According to Cudahy and Leshowitz, acute subjects would not show backward recognition masking. The mean Δfs presented to one subject from Experiment 1 (J. P.) and four subjects from Experiment 2 were in the range of 7 to 8 Hz. As can be seen in Figure 3, each of these sensitive subjects demonstrated healthy ABRM functions. This represents very sensitive performance given the short 20-msec test tones and the masking conditions. Consequently, it cannot be argued that ABRM occurs only when insensitive listeners are used as subjects.

General Discussion

Experiments 1 and 2 failed to demonstrate a mask disinhibition effect using an ABRM task. Thus it would appear that lateral inhibition is not sufficient to explain masking in the ABRM task. Given previous demonstrations of disinhibition effects in stimulus suffix experiments (Crowder, 1978; Morton, 1976), the present results might seem to suggest that

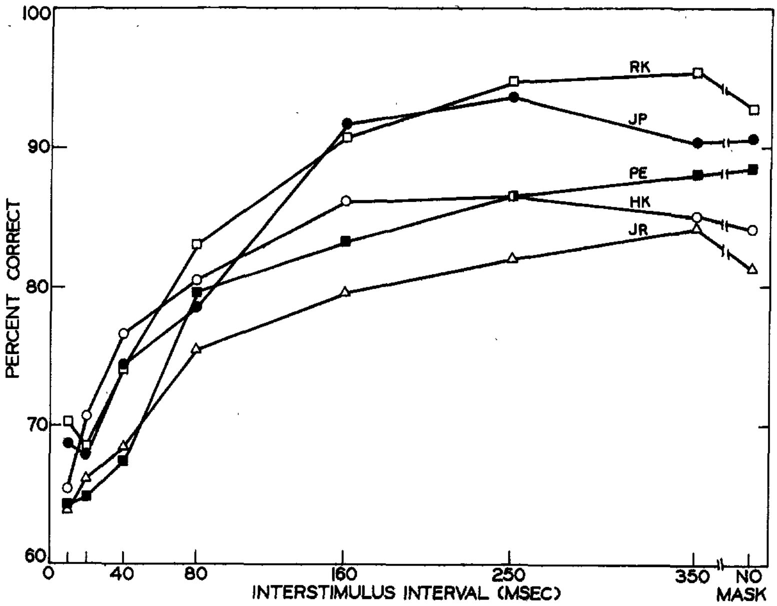

Figure 3. The percentage of correct identifications of the test tones as a function of the duration of the interstimulus interval. (The curves represent individual subject data for sensitive subjects in Experiments 1 and 2.)

different underlying processes are responsible for stimulus suffix and ABRM effects. However, a recent failure to replicate the disinhibition effect in a series of stimulus suffix experiments (Watkins & Watkins, 1982) suggests caution in concluding that mask disinhibition occurs in the suffix paradigm but not in ABRM. Until such time as Watkins and Watkins's failure to replicate is clarified, it would be premature to use the present results to argue that ABRM and stimulus suffix effects tap different underlying processes. However, we believe that the case for different underlying mechanisms in ABRM and stimulus suffix experiments can be made on the basis of a number of recent experimental findings. In the remainder of this article, we review this evidence within a broader discussion of preperceptual auditory storage.

An Evaluation of Preperceptual Auditory Storage and Its Measurement

Contemporary theorizing on human information processing has often relied on three-store information-processing models as explanatory frameworks. The major structural components incorporated in most three-store models are a sensory store (the auditory sensory store is often referred to as *echoic*), short-term store, and long-term store, and it is assumed that after transduction by the sensory apparatus, a stimulus's properties are stored in the sensory store and then transformed into short- and/or long-term codes (see Lachman, Lachman, & Butterfield, 1979, for elaboration). Despite the widespread acceptance of three-store models, we believe that they shortchange the types of operations that are necessary to process perceptually stimulus events. For illustration, consider the processing of a short sound. The unanalyzed features of the sound would need to be (a) represented in preperceptual storage and (b) synthesized into a perceptual experience. We assume that the perceptual experience (i.e., the phenomenological experience of hearing the sound's qualities) would result from evaluating and integrating the sound's features from the preperceptual store and synthesizing a percept. This percept would be stored in a structure we call synthesized auditory memory, a memory that may be thought of as something akin to an echo of the original perceptual experience (Waugh & Norman, 1965). It is important to note that the synthesized auditory memory is assumed to be precategorical because it is possible for a listener to "hear" a sound prior to categorizing it; for example, it is possible to hear a tone prior to categorizing it as either relatively high or low in pitch. Nevertheless, the synthesized memory is the outcome of having processed the unanalyzed features in preperceptual storage and is thus distinct from the preperceptual store. In addition, synthesized memory can persist after categorization of an auditory experience has taken place (Massaro, 1975b).

The distinction between preperceptual auditory storage and synthesized auditory memory is critical to any analysis of perceptual processing. A number of studies support the idea that perception and retention of auditory information are successive processes in pitch memory tasks, and these studies have been summarized in detail elsewhere (Massaro, 1975b, 1976a). For example, perception of a standard test tone is enhanced if its duration is increased, but this manipulation does not affect the rate of forgetting, whereas increasing the duration of an *interference* tone degrades the memory of the test tone but does not affect its initial perception or storage (Massaro, 1970b; Wickelgren, 1969). Indeed, quantitative models developed to describe tonal memory generally distinguish between perceptual resolution (which is limited by available processing time according to our view) and retention of the synthesized auditory memory (cf. Massaro, 1970b). Because three-store processing models fail to differentiate between preperceptual and synthesized auditory storage, a single-storage structure, usually referred to as the echoic store, subsumes the functions of both. Experiments designed to measure the duration of echoic memory are not designed to isolate the contributions of preperceptual and synthesized storage structures. The same difficulty applies to the concept of a *precategorical* store. Although Crowder's definition of the precategorical store suggests that it holds unanalyzed features and thus is essentially equivalent to what we refer to as a preperceptual store, Crowder's formulation includes nothing equivalent to our synthesized auditory memory, and indeed, as we will elaborate below, his experimental findings do not rule out the possibility that the stimulus suffix effect indexes the contributions of later stages of processing such as synthesized auditory memory rather than the persistence of preperceptual storage.

Although it might have been thought that a single underlying mechanism could explain the results of ABRM and stimulus suffix experiments, that is, replacement of unanalyzed features of a stimulus in preperceptual storage by features of a subsequent stimulus, recent research suggests that both ABRM and the stimulus suffix effect result from a number of underlying mechanisms. We now review the evidence supporting this position, first considering ABRM and then the stimulus suffix effect.

Auditory Backward Recognition Masking

Massaro (1972a, 1975a, 1975b) explained ABRM within the framework of a successive-stage model of human information processing. The relevant part of the model assumes that after transduction by a listener's sensory apparatus, properties of a presented auditory sound are stored as features in a centrally located preperceptual auditory store. The preperceptual auditory store is assumed to be a single-channel structure capable of storing a single auditory event. Perception of a sound (i.e., the phenomenological experience of hearing the sound) is assumed to result from (a) reading out the sound's features from the preperceptual store and (b) transforming the set of features into a synthesized auditory memory. However, if a second sound occurs prior to completion of the readout of the first sound's features, the second sound is assumed to replace the first in the preperceptual store and, in so doing, will terminate perceptual processing of the first sound. Because identification performance (percentage correct) in ABRM experiments reaches an asymptote at SOAs of approximately 250 msec, Massaro (1972a, 1975a, 1975b) has argued that the readout of information from the preperceptual store (assuming that no backward mask is presented) is completed within approximately 250 msec.

Although experiments on ABRM have proven illuminating with regard to perceptual processing, Massaro's original explanation of ABRM in terms of the readout of information from a preperceptual store has required some elaboration. Two major results were problematic to Massaro's explanation. First, in some ABRM experiments, test tone identification has been shown to be biased by qualities of the masking stimulus (e.g., its frequency or pitch). For example, a listener is more likely to identify a test tone as relatively high in pitch if a high- rather than low-frequency mask follows the test tone (Hawkins & Presson, 1977; Hawkins, Thomas, Presson, & Cozic, 1974; Kallman & Massaro, 1979). Massaro's explanation of backward recognition masking could not easily accommodate the mask-induced biases because it assumed that, except at the very short SOAs (i.e., approximately 60 msec and less) where perceptual integration of the two tones might occur, the sole effect of the mask on test tone identification was to replace the test tone information in the preperceptual store, thereby terminating perceptual processing of the test tone. Massaro's explanation provided no mechanism for the masking tone to change the nature of the perceptual report of the test tone. A second finding outside the domain of Massaro's original explanation of backward recognition masking was that the magnitude of backward masking depends for some subjects on the similarity of the mask to the test tone (Kallman & Massaro, 1979). This result was problematic because if the mask had its effect solely by replacing the representation of the test tone in a single-channel preperceptual store, its masking effectiveness should not depend on the similarity of the mask to the test tone.

To accommodate the findings of mask-induced biases and mask/test tone similarity effects, Kallman and Massaro (1979) presented an elaborated model of ABRM. As in Massaro's earlier model, the central assumption is that the mask terminates the readout of the test information from the preperceptual store by replacing the test sound's representation in the preperceptual store. In addition, it is assumed that during the time that mask information is read out from the preperceptual store into the synthesized auditory memory, the mask may interfere with and/or systematically modify the synthesized memory of the test tone. Because the subject would not normally arrive at a final decision about the test tone's identity until sometime after mask onset, this decision would be influenced by mask interference. Previous research has shown that auditory memory interference depends on the similarity between the interfering stimulus and the to-be-remembered stimulus (Deutsch, 1974; Massaro, 1970b). But although Kallman and Massaro did show that synthesized memory interference may play a role in ABRM, they demonstrated that after accounting for mask-induced synthesized memory interference, a substantial portion of the ABRM effect remained that could best be explained in terms of interruption of the readout of the test tone from the preperceptual auditory store. Thus ABRM appears to be the result of two underlying processes: termination of the readout of information from the preperceptual store (as was originally proposed by Massaro) and mask-induced synthesized memory interference. In the present Experiment 1, ABRM was greater when three masks rather than one were presented, and this may have been due to the greater amount of synthesized memory interference afforded by three as opposed to one mask. Similar results were found in Experiment 2, but the difference was not statistically significant, possibly because the comparison was between the effectiveness of one and just two masks.

Although Kallman and Massaro (1979) provided evidence that ABRM is the result of both interruption of perceptual processing and synthesized memory interference, their findings do not invalidate Massaro's earlier conclusions regarding the duration of perceptual processing or the persistence of the information in preperceptual storage. It appears that the asymptote in ABRM functions occurs at an SOA of approximately 250 msec regardless of whether a high interference (e.g., a tone) or a low interference (e.g., white noise) mask is used. In any case, according to Kallman and Massaro's theory, mask interference with the synthesized memory of the test tone could only serve to accentuate any masking effects obtained. Consequently, to the extent that mask interference was a factor in ABRM experiments, the duration of perceptual processing would be overestimated, not underestimated. ABRM experiments suggest that the readout of information from the preperceptual store takes no more than approximately 250 msec.

Although ABRM experiments index the duration of perceptual processing, one finding seems to suggest that ABRM experiments may underestimate the duration of preperceptual storage. In addition to the standard ABRM task, Massaro (1972b) varied the duration of a test tone from 40 to 440 msec. In this latter condition, a backward mask designed to terminate test tone processing followed the offset of the test tone after a 10-msec silent interval. Performance on the test tone identification task improved with increases in the test tone's duration up to 250 msec. This result suggests that even when the test tone remains physically present for longer than 250 msec, perceptual resolution of the test tone is completed with 250 msec. It would follow that the asymptote in performance on ABRM tasks cannot be used as a measure of the persistence of preperceptual storage; it would be possible that preperceptual storage persists for more than 250 msec, but because the perceptual process is complete within 250 msec, additional processing time has no effect on test tone identification. If Massaro's result is valid, the results of ABRM experiments cannot place an upper limit on the persistence of preperceptual storage. The finding that performance continues to improve until an SOA of approximately 250 msec is reached can only suggest that the features in preperceptual storage persist for *at least* 250 msec.

Massaro's (1972b) result should not be taken as necessarily disqualifying the ABRM task as an estimate of preperceptual storage because there is an alternative explanation of the result. The trials with long tones were intermixed with trials of short tones. In fact, only about 19% of the trials contained test tones longer than 250 msec. Subjects in this task may have adopted the strategy of arriving at a decision about the relatively long-duration test tones some time prior to their offset, given that on most trials test tone information was available for only 250 msec or less. If subjects adopted the strategy of basing their test tone decision on no more than 250 msec of information, this would have precluded an improvement in performance on test tones longer than 250 msec. In addition, the absolute-identification task might not be the best task to reveal a performance improvement when actual processing time is greater than 250 msec. Subjects generally may not have the patience to continue processing when their decision can be made at any time during the test tone presentation. A more appropriate task might be a successive-comparison task in which, on each trial, two tones are presented and subjects are required to determine whether the second-presented tone is higher or lower in pitch than the first. Given such a task, it is impossible for subjects to make the required pitch judgment until after the second tone is presented. Thus the duration of the first tone can be varied without the possibility that the manipulation would be negated by a response strategy of the kind described above. Until additional studies are conducted that clarify Massaro's (1972b) results, it is premature to dismiss ABRM experiments as a possible index of the persistence of preperceptual auditory information.

The Stimulus Suffix Effect

Crowder (1978) argued that the recency effect found when a list of words is presented auditorily occurs because the terminal list item's representation in precategorical acoustic storage serves as a source of information about the identity of the terminal item; in addition, information about the terminal item in the list that has been coded into an abstract short-term memory representation is available to aid in recall of the terminal item. Thus, at the time of recall, the subject has two sources of information available about the terminal item (i.e., precategorical acoustic and abstract short-term memory) on which to base his or her recall response. In contrast, information about the earlier presented items is available only from abstract short-term memory because presentation of each auditory item serves to mask the information about earlier items held in the precategorical auditory store. If an auditory suffix is presented subsequent to the terminal item, precategorical auditory information about the terminal item will be masked, and, consequently, there would be no advantage for the terminal as opposed to earlier items because recall of both would be based only on abstract short-term memory; as outlined in the introduction to this article, Crowder has proposed that lateral inhibitory processes underlie the masking effect. Because a stimulus suffix can interfere with recall of the terminal item if the suffix is presented within 2 sec, Crowder has suggested that information in the precategorical auditory store persists for at least 2 sec.

However, recent experiments challenge the view that the suffix effect provides an undistorted index of precategorical storage. For example, Spoehr and Corin (1978) obtained a suffix effect when the stimulus suffix was articulated silently by the experimenter and subjects were required to attend to the movement of the experimenter's lips. Based on this result, Spoehr and Corin suggested that rather than replacing previous items in precategorical storage, the suffix interfered with the formation of an articulatory code in short-term memory. Converging evidence for a short-term memory explanation of the suffix effect was provided by Campbell and Dodd (1980) who found a suffix effect when an auditory suffix followed a list of silently articulated items that the subjects were required to lip-read. This result contradicted the precategorical acoustic store explanation of the suffix effect because given that the silently articulated list items would never have entered precategorical acoustic storage, interference with memory of the list items could not have been at the level of the precategorical acoustic store.

Ayres, Jonides, Reitman, Egan, and Howard (1979) also argued that the suffix effect cannot be used as evidence to support the construct of a precategorical acoustic store. In their experiment, one of the suffixes was a plunger-muted trumpet note that sounded somewhat like a nasally spoken *wa*. Because the trumpet note could be interpreted as either a musical note or a speech sound, it was possible for Ayres et al. to manipulate the subject's categorization of the sound. Thus they compared recall of a list of words when it was followed by the trumpet note interpreted as music or the trumpet note interpreted as speech. Because in both cases the stimulus suffix was acoustically identical, an interpretation of the stimulus suffix effect in terms of precategorical acoustic storage predicts that the suffix effect should be identical in the two cases. Ayres et al. found, however, that the size of the suffix effect was substantially greater given the speech trumpet note rather than the musical trumpet note. This result clearly argues that the suffix effect is sensitive to the semantic category to which a stimulus belongs and is thus not precategorical. Salter and Colley (1977) have also provided evidence that the *categorical* similarity of the suffix to the list affects the size of the suffix effect.

In addition to the general difficulties with a precategorical interpretation of the suffix effect raised by the above studies, there are alternative explanations for the advantage of the three suffixes over a single suffix. For example, Kahneman and Henik (1981) explained the advantage in terms of a theory of perceptual groupings. They proposed that the list of words presented to a subject is typically treated as a perceptual group, the terminal item of which has a privileged status leading to superior recall. This advantage of the terminal item is eliminated with a suffix, which becomes the terminal item in the list. It is assumed that three identical suffixes are more likely to segregate into a distinct perceptual group than is a single suffix. The result is that the terminal list item is more likely to occupy the terminal position in a perceptual group and thus is more often recalled with three suffixes rather than one (see Morton, 1976, for a similar explanation).

Although Kahneman and Henik's explanation of the suffix effect in terms of perceptual groups cannot easily explain why suffix effects are more likely to be found with auditory rather than visual presentation of the stimuli and also fails to adequately address why the final items in a perceptual group have privileged status, their theory has received some empirical support (Kahneman & Henik, 1977, 1981). However, it is not clear whether the stimulus groupings proposed by Kahneman and Henik truly reflect perceptual processes or whether instead they reflect the way the items are organized in memory.

Although the stimulus suffix effect is interesting, it is clear that the effect is not due entirely to interference with preperceptual or precategorical acoustic storage. Analogous to our interpretation of ABRM as being due to both termination of perceptual processing and interference from synthesized memory, a number of processing stages may be responsible for the stimulus suffix effect. It will be necessary for researchers to isolate experimentally the contributions of the various processing stages before the stimulus suffix effect can be used as an uncontaminated index of the duration of preperceptual storage.

The Partial-Report Task

Although the main focus of this article is on ABRM and the suffix effect, one of the most influential studies on preperceptual auditory storage used the partial-report task in an attempt to clarify the duration of preperceptual storage (Darwin, Turvey, & Crowder, 1972). Three lists of three items each were simultaneously presented to the left, middle, and right sides of the head. It was reasoned that if the information was held in preperceptual form along the three spatial locations, then a partial-report cue some time after the list presentation would lead to highly accurate report of the items presented to the cued location. Under the logic of Sperling (1960) and others, the decay in partial report with the delay of the location cue would define the duration of the preperceptual auditory store.

The results revealed a very small (but statistically significant) decrease in partial report with increases in the delay of the location cue. In terms of the number of items available, there were 4.9 with an immediate cue and 4.4 with a cue delay of 4 sec. The whole report was 4.3 items correct. This result was construed as evidence for preperceptual auditory storage that lasts on the order of seconds. Unfortunately, there is a critical weakness in using the results of Darwin et al.'s experiment as a demonstration of the time course of preperceptual storage. The small decrease in performance with increases in the delay of the location cue could reflect nothing more than forgetting from short-term memory for items already recognized and categorized. Given the presentation of 9 items, it is not unexpected that subjects would tend to forget some while waiting for the partial-report cue.

Additional converging operations are necessary to use the partial-report task as an illustration of preperceptual storage. Critical to Sperling's use of the partial-report task to support precategorical visual storage was the demonstration that a category cue was an ineffective partial-report cue. That the location cue was an effective cue and the category cue was not suggested that the items were in a precategorical form at the time of cue presentation. When Massaro (1976b) provided a direct comparison between location and category cues using Darwin et al.'s basic experimental procedure, the results showed no advantage of the location cue when the partial-report cue was presented immediately after the list presentation. This result suggests that the forgetting of one half of an item over the period of 4 sec reported by Darwin et al. represented the loss of information that was categorical rather than precategorical. This is not so surprising because given the extended duration of a list of three auditorily presented words, much of the recognition and categorization of the words could occur during the list presentation and little preperceptual information would remain after list presentation was complete. It would appear that until clarification of the processing stages responsible for the partial-report advantage is forthcoming, the partial-report advantage cannot be used as evidence for preperceptual storage.

Another Paradigm

Kubovy and Howard (1976) devised a unique task to measure the duration of auditory storage. The basic stimulus was a nonharmonic chord made up of six tones. If such tones are of equal amplitude, the stimulus is heard as noise. However, if the chord is presented through headphones and one of the tones is presented with a different interaural ongoing-time disparity than the other tones, the deviant tone stands out; that is, it perceptually segregates. If each tone is given a different interaural ongoing-time disparity, the chord is again heard as noise. However, if a second chord is generated using the same time disparities of all the tones from the first chord but one, and the two chords are played in succession, the changed tone might stand out. If it did, it would mean that, at some level, information about the time disparities of the tones in the first chord was preserved after its offset. In this case, varying the interval between the two chords and assessing whether the changed tone is heard (segregated) would provide an estimate of the duration of auditory storage.

Kubovy and Howard (1976) used this basic logic, but their experiment was somewhat more complex than that described above. Instead of presenting two chords per trial, Kubovy and Howard presented 18 consecutive chords, and the only difference between chords was in the patterns of interaural time disparities of the component tones. Each chord was constructed in such a way that the interaural time disparity of one component tone, the *critical tone*, differed from the standard canonical configuration used on that trial; the rest of the tones were identical to those of the standard canonical configuration. The stimulus conditions were arranged such that the critical tones on a trial were composed of three repetitions of either a six-note ascending pitch sequence or a six-note descending sequence. The subject's task was to identify whether the sequence was ascending or descending. An adaptive procedure was used to determine for each subject the silent interval that resulted in 71% correct performance for each subject. Five of the six tested subjects performed at 71% correct with silent intervals averaging approximately 1 sec, but the other subject performed perfectly when the silent intervals were as long as 9.7 sec. These results have been interpreted as suggesting that interaural time disparity information persists in preperceptual storage for a duration that varies from subject to subject but usually is on the order of 1 sec.

Although Kubovy and Howard's paradigm is intriguing and offers the potential to clarify the persistence of auditory memory, we question their estimates of the duration of preperceptual auditory storage. First, certain aspects of their experiment would tend to yield an overestimate of the average persistence of auditory storage. To perform the experimental task, it is necessary only to segregate the critical tone on two consecutive chord presentations, yet the subject is presented a sequence of 18 chords. It would thus be possible to perform perfectly if the critical tone was segregated on only a small percentage of presentations. Combined with the fact that the estimate of persistence was based on the silent interval that yielded correct responses on only 71% of the trials, the estimates of average persistence derived by Kubovy and Howard may have been too high.

Another reason to question Kubovy and Howard's estimates of preperceptual auditory persistence is that it is likely that later processing stages contributed to their results. Two tones per chord rather than one should have segregated if the subject had available only interaural time disparity information about the most recently presented chord. Given two consecutively presented chords that each differ from the standard canonical configuration by a single tone's interaural time disparity, the two chords would differ from each other by the interaural time disparities of two of their tones. Because consecutive chords in the Kubovy and Howard experiment differed by the time disparities of two of their tones, subjects should have segregated two tones per chord rather than one. But according to Kubovy and Howard, only one tone per chord segregated, and this was the one that differed from the canonical configuration. This result suggests that the critical determinant of whether a tone segregated was that it differed from the canonical configuration, not that it differed from the previously presented chord. It is not at all clear that the representation of the canonical configuration would be held in preperceptual storage; it seems more likely that later processing stages would be responsible for its storage. Future research should clarify this issue.

Finally, Kubovy and Howard's estimates of auditory persistence ranged from less than 1 to over 9 sec for different subjects. Such a range of values seems somewhat extreme for estimates of preperceptual storage and might have been the result of the use of different coding strategies by different subjects. Kallman and Massaro (1979) suggested that the use of different strategies may affect the quality of coding at the later processing stages such as synthesized auditory memory but are unlikely to influence preperceptual storage. Thus, although Kubovy and Howard's paradigm has potential for defining properties of auditory memory, their estimates of auditory persistence need further testing and more research is necessary to clarify the processing stages responsible for their results.

Remaining Issues

We believe that ABRM remains an effective tool for exploring the temporal course of perceptual recognition. But there are a few outstanding issues about backward masking and auditory processing that should be addressed, and these will be discussed soon. One question concerns the relationship between integration and interruption masking. A second question is how it is possible to understand speech if backward masking occurs during speech processing. Before addressng these issues directly, we should clarify what might be a general misunderstanding of what backward masking does and how it is interpreted in serial-stage models of information processing. The misunderstanding is the belief that a backward mask somehow works retroactively and eliminates *all* perceptual information about the test stimulus. Thus unless subjects manage to encode the test stimulus into an abstract representation, they would not be able to report anything about what was presented. In the present view of backward masking, however, the mask simply terminates any further perceptual resolution; the resolution that occurs prior to the presentation of the mask is continuously passed on to the next processing stage, that is, synthesized auditory memory. The information in synthesized auditory memory is not eliminated by the mask although under certain conditions the mask may also function to interfere with (but not necessarily completely eliminate) information held in the synthesized memory (Kallman & Massaro, 1979). To reiterate, the gradual improvement in performance with increases in the ISI between test stimulus and mask reflects the continuous perceptual processing that occurs prior to mask presentation (cf. Massaro, 1970a).

Interruption and Integration Masking

Our discussion of the ABRM task has focused on an interruption explanation of backward recognition masking. A second sound can replace an earlier sound in preperceptual auditory storage and, therefore, terminate any further perceptual processing of the first sound. A complete account of backward masking also requires the concept of integration masking.

There is reason to believe that when two short sounds are presented with an SOA of approximately 60 msec or less, the perceptual system integrates the two sounds into a single compound percept. One line of evidence supporting this view comes from studies of forward masking. Identification of a target sound is impaired if it is preceded with an SOA of 60 msec or less by an equal amplitude *forward mask* (Massaro, 1973). This type of forward masking can be explained by assuming perceptual integration of the two sounds. If the two sounds integrate to form a compound stimulus, it would be difficult to identify the target sound against the irrelevant background of the forward mask. Effectively, the forward mask would serve to reduce the signal/noise ratio of the target sound.

If sounds with SOAs of 60 msec or less integrate, one would expect to find some effect of integration masking in backward as well as forward masking experiments. Although the data across experiments are somewhat inconsistent on this point, it is often found that the first 40–60 msec of ABRM functions fail to demonstrate the monotonic increase in percentage correct with increases in SOA that is found with SOAs greater than 60 msec. In fact, there is often a decrease in performance as the SOA increases from 0 to 60 msec. The result is that when the entire ABRM function is examined, a U-shaped function similar to that often found in visual backward masking studies is often seen. Analogous to the forward masking results, it can be assumed that at short SOAs, the two sounds tend to integrate thus forming a compound sound that can receive unlimited processing. As the SOA approaches 60 msec, the probability that the perceptual system will treat the two sounds as separate increases and, consequently, so does the probability that interruption masking will occur. It should be noted that the evidence of integration in ABRM occurs only when the target sound and the mask are presented to the same ear whereas interruption masking occurs regardless of whether the mask is presented to the same ear as the target or to a different ear (Hawkins & Presson, 1977; Massaro & Cohen, 1975). This finding is analogous to the results of visual backward masking studies that have found that integration masking occurs only when the mask and target stimulus are presented to the same eye but interruption masking does not depend on whether the target and mask are presented to the same eye (Turvey, 1973).

Although nonmonotonic masking functions provide evidence for integration, a nonmonotonic masking function is not a necessary result of an integration process in a backward masking task. If integration did not occur, then performance would improve with increases in SOA at a given rate. Integration of the test and mask would preclude interruption masking, but the quality of the integrated percept of the test and mask may not be good enough to produce a performance level that is better than that given by interruption masking at a longer SOA. Thus the failure to consistently find nonmonotonic masking functions does not argue against integration masking at short SOAs. Definitive tests for integration processing must be formulated in the context of specific quantitative models of backward masking (cf. Massaro, 1975b, Chapter 18). Furthermore, given that integration results in the mask and test stimulus forming a compound stimulus, a better understanding of integration masking requires a detailed analysis of how the type of backward mask used influences the integration process. But even without definitive tests for integration and a full understanding of the integration process, the results of the forward and backward masking studies when considered together provide at least reasonable support for the view that auditory events that occur with SOAs of approximately 60 msec or less are likely to integrate perceptually. Furthermore, converging evidence for a perceptual integration period of at least 60 msec is provided by the finding that a subject's ability to identify an isolated tone's pitch improves as the duration of the tone increases up to about 60–80 msec (Massaro, 1972b). This result suggests that the perceptual system can integrate auditory information over a temporal period of at least 60 msec.

The masking studies suggest that two physically distinct sounds may integrate if they occur with an SOA of 60 msec or less. And the finding that perception of a tone improves as its duration increases up to about 60 msec provides additional evidence suggesting that perceptual integration occurs over at least a 60-msec interval. However, it is unclear whether these results generalize directly to speech sounds. Because the speech stimulus often consists of gradual changes over time and because it is problematic to segment the speech stimulus, predicting the conditions under which integration as opposed to interruption should occur is difficult. Although we cannot at the present time clarify this ambiguity, our view of auditory processing can accommodate both the occurrence of integration and interruption in speech perception. But with regard to backward masking, it must be stressed that our interest in the phenomenon is primarily based on its utility as a tool for estimating the duration of the perceptual process; this utility is quite independent of whether backward masking naturally occurs in everyday perceptual situations. ABRM experiments have proven themselves useful because they index the readout of information from preperceptual auditory storage.

Backward Masking and Speech Perception

According to our view of ABRM, if some ABRM occurs during speech processing, this would not imply that the perceptual cues necessary to recognize a speech message would be totally unavailable to the perceiver. Rather, the resolution of a masked speech sound might be somewhat less than if the sound had been presented in isolation. But given words in a sentence, degradation of the perceptual process by ABRM would be compensated for by the effects of context. There is a growing body of evidence demonstrating that context can aid in the perception of meaningful speech, particularly when the perceptual information presented to the listener is degraded in some way. For example, a word with an absent or mispronounced phoneme is often perceived as having been correctly pronounced if the sentential or syllabic context constrains the number of linguistically or semantically permissable alternatives (Marslen-Wilson & Welsh, 1978; Warren & Obusek, 1972). Apparently, the effect of context is to supplement degraded perceptual information. Massaro (1978) and Massaro and Oden (1980) proposed that the perceptual synthesis of a speech sound by a listener depends on both the acoustic featural information and the contextual constraints in the message. In fact, a quantitative model assuming independent contributions of these sources of information provided a good fit to a number of findings of context effects in speech processing.

Although we believe that context could compensate for the effects of ABRM on speech perception, we do not mean to imply that ABRM would always be operating during the perception of speech. For example, if a 250-msec consonant–vowel (CV) syllable was followed by another CV speech sound, perceptual resolution of the first CV syllable would be relatively unaffected. The onset characteristics and the formant transitions that identify the syllable would have occurred early in the syllable and, consequently, would have been read out from preperceptual storage prior to presentation of the subsequent CV syllable. Massaro (1974, 1975c) provided a more detailed discussion of the role of preperceptual auditory storage and backward masking in speech perception.

Conclusions

Backward recognition masking remains an effective tool for exploring properties of preperceptual storage. Most of the other experimental paradigms that have been used to explore auditory memory seem to tap later processing stages and, consequently, cannot be used to explore preperceptual storage. The results of ABRM experiments suggest that the readout of information from preperceptual auditory storage takes approximately 250 msec.

References

Ayres, T. J., Jonides, J., Reitman, J. S., Egan, J. C., & Howard, D. A. Differing suffix effects for the same physical suffix. *Journal of Experimental Psychology: Human Learning and Memory,* 1979, *5*, 315–321.

Campbell, R., & Dodd, B. Hearing by eye. *Quarterly Journal of Experimental Psychology,* 1980, *32*, 85–99.

Coltheart, M. Iconic memory and visible persistence. *Perception & Psychophysics,* 1980, *27*, 183–228.

Crowder, R. G. Mechanisms of auditory backward masking in the stimulus suffix effect. *Psychological Review,* 1978, *85*, 502–524.

Crowder, R. G., & Morton, J. Precategorical acoustic storage (PAS). *Perception & Psychophysics,* 1969, *5*, 365–373.

Cudahy, E., & Leshowitz, B. Effects of a contralateral interference tone on auditory recognition. *Perception & Psychophysics,* 1974, *15*, 16–21.

Darwin, C. J., Turvey, M. T., & Crowder, R. G. An auditory analogue of the Sperling partial report procedure: Evidence for brief auditory storage. *Cognitive Psychology,* 1972, *3*, 255–267.

Deutsch, D. Generality of interference by tonal stimuli in recognition memory for pitch. *Quarterly Journal of Experimental Psychology,* 1974, *26*, 229–234.

Hawkins, H. L., & Presson, J. C. Masking and preperceptual selectivity in auditory recognition. In S. Dornic (Ed.), *Attention and performance VI.* Hillsdale, N.J.: Erlbaum, 1977.

Hawkins, H. L., Thomas, G. B., Presson, J. C., Cozic, A., & Brookmire, D. Precategorical selective attention and tonal specificity in auditory recognition. *Journal of Experimental Psychology,* 1974, *103*, 530–538.

Kahneman, D., & Henik, A. Effects of visual grouping on immediate recall and selective attention. In S. Dornic (Ed.), *Attention and performance VI.* Hillsdale, N.J.: Erlbaum, 1977.

Kahneman, D., & Henik, A. Perceptual organization and attention. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization.* Hillsdale, N.J.: Erlbaum, 1981.

Kallman, H. J., & Massaro, D. W. Similarity effects in backward recognition masking. *Journal of Experimental Psychology: Human Perception and Performance,* 1979, *5*, 110–128.

Kubovy, M., & Howard, F. P. Persistence of a pitch-segregating echoic memory. *Journal of Experimental Psychology: Human Perception and Performance,* 1976, *2*, 531–537.

Lachman, R., Lachman, J. L., & Butterfield, E. C. *Cognitive psychology and information processing: An introduction.* Hillsdale, N.J.: Erlbaum, 1979.

Long, G. M. Iconic memory: A review and critique of the study of short-term visual storage. *Psychological Bulletin,* 1980, *88*, 785–820.

Marslen-Wilson, W. D., & Welsh, A. Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology,* 1978, *10*, 29–63.

Massaro, D. W. Preperceptual auditory images. *Journal of Experimental Psychology,* 1970, *85*, 411–417. (a)

Massaro, D. W. Retroactive interference in short-term recognition memory for pitch. *Journal of Experimental Psychology,* 1970, *83*, 32–39. (b)

Massaro, D. W. Preperceptual images, processing time and perceptual units in auditory perception. *Psychological Review,* 1972, *79*, 124–145. (a)

Massaro, D. W. Stimulus information vs. processing time in auditory pattern recognition. *Perception & Psychophysics,* 1972, *12*, 50–56. (b)

Massaro, D. W. A comparison of forward versus backward recognition masking. *Journal of Experimental Psychology*, 1973, *100*, 434–436.

Massaro, D. W. Perceptual units in speech recognition. *Journal of Experimental Psychology*, 1974, *102*, 199–208.

Massaro, D. W. Backward recognition masking. *Journal of the Acoustical Society of America*, 1975, *58*, 1059–1065. (a)

Massaro, D. W. *Experimental psychology and information processing*. Chicago: Rand McNally, 1975. (b)

Massaro, D. W. *Understanding language: An information processing analysis of speech perception, reading, and psycholinguistics*. New York: Academic Press, 1975. (c)

Massaro, D. W. Auditory information processing. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 4). Hillsdale, N.J.: Erlbaum, 1976. (a)

Massaro, D. W. Perceptual processing in dichotic listening. *Journal of Experimental Psychology: Human Learning and Memory*, 1976, *2*, 331–339. (b)

Massaro, D. W. A stage model of reading and listening. *Visible Language*, 1978, *12*, 3–26.

Massaro, D. W., & Cohen, M. M. Preperceptual auditory storage in speech recognition. In A. Cohen & S. G. Nooteboom (Eds.), *Structure and process in speech perception*. Berlin: Springer-Valag, 1975.

Massaro, D. W., & Idson, W. L. Temporal course of perceived auditory duration. *Perception & Psychophysics*, 1976, *20*, 331–352.

Massaro, D. W., & Idson, W. L. Backward recognition masking in relative pitch judgments. *Perceptual and Motor Skills*, 1977, *45*, 87–97.

Massaro, D. W., & Oden, G. C. Speech perception—A framework for research and theory. In N. J. Lass (Ed.), *Speech and language: Advances in basic research and practice* (Vol. 3). New York: Academic Press, 1980.

Morton, J. Two mechanisms in the stimulus suffix effect. *Memory & Cognition*, 1976, *4*, 144–149.

Morton, J., Crowder, R. G., & Prussin, H. A. Experiments with the suffix effect. *Journal of Experimental Psychology*, 1971, *91*, 169–190. (Monograph)

Salter, D., & Colley, J. G. The stimulus suffix: A paradoxical effect. *Memory & Cognition*, 1977, *5*, 257–262.

Sperling, G. The information available in brief visual presentations. *Psychological Monographs*, 1960, *74*(11, Whole No. 498).

Spoehr, K. T., & Corin, W. J. The stimulus suffix effect as a memory coding phenomenon. *Memory & Cognition*, 1978, *6*, 583–589.

Turvey, M. T. On peripheral and central processes in vision: Inferences from an information processing analysis of masking with patterned stimuli. *Psychological Review*, 1973, *80*, 1–52.

Warren, R. M., & Obusek, C. J. Identification of temporal order within auditory sequences. *Perception & Psychophysics*, 1972, *12*, 86–90.

Watkins, O. C., & Watkins, M. J. Lateral inhibition and echoic memory: Some comments on Crowder's (1978) theory. *Memory & Cognition*, 1982, *10*, 279–286.

Waugh, N. C., & Norman, D. A. Primary memory. *Psychological Review*, 1965, *72*, 89–104.

Wickelgren, W. A. Associative strength theory of recognition memory for pitch. *Journal of Mathematical Psychology*, 1969, *6*, 13–61.

Received September 28, 1981
Revision received August 20, 1982 ■