Journal of Experimental Psychology:
Human Perception and Performance
1998, Vol. 24, No. 5, 1454–1466

0096-1523/98/$3.00

Beyond the Attentional Blink: Visual Masking by Object Substitution

Barry Giesbrecht
University of Alberta

Vincent Di Lollo
University of British Columbia

If 2 targets are to be identified among distractors displayed in rapid sequence, correct identification of the 1st target hinders identification of the 2nd. To obtain this attentional blink (AB), the 1st target must be masked with a simultaneous (integration) or a delayed (interruption) mask indifferently. In 3 experiments, it was shown that the 2nd target must also be masked, but that the precise form of masking is important: An AB occurs with interruption but not with integration masking. This nonequivalence of masking paradigms parallels that found in studies of masked priming, a phenomenon arguably related to the AB. The results are explained by a revised 2-stage model (M. M. Chun & M. C. Potter, 1995).

Visual attention can be deployed over space or over time. Deployment over space has been studied extensively with several different paradigms (see reviews by Kinchla, 1992, and by LaBerge, 1990). Deployment over time has been studied less extensively, mainly with two related paradigms. In one paradigm, known as rapid serial visual presentation (RSVP), stimuli are presented sequentially in the same location, and observers are asked to identify one or more targets within the stream. When two targets are to be reported, the first is identified almost perfectly, but identification of the second is substantially impaired. The impairment is most evident when the second target is presented with a temporal lag of about 200–500 ms after the first (Raymond, Shapiro, & Arnell, 1992). At shorter or longer lags, performance is impaired less or not at all, thus yielding a characteristic U-shaped function over lags. This second-target deficit, also known as the *attentional blink* (AB), is said to occur because processing mechanisms required commonly by the two targets are unavailable (or are less available) for processing the second target until first-target processing has been completed (e.g., Chun & Potter, 1995). A very similar second-target deficit has been obtained in an elegantly simplified paradigm in which two targets, masked by trailing pattern-masks, are displayed at different screen locations, at various temporal lags from each other (Duncan, Ward, & Shapiro, 1994). The terms *attentional blink* and *second-target deficit* are used interchangeably in the present work.

Regardless of paradigm, a second-target deficit is obtained only if the target stimuli are masked. In the simplified paradigm (Duncan et al., 1994), masking is provided by temporally trailing patterns displayed in the same locations as the targets. In the RSVP paradigm, items are displayed sequentially in the same location at the rate of approximately one every 100 ms, so each target is masked by the ensuing item in the stream. The masking effect of the trailing item on the processing of the first target has been amply documented. For example, it has been found that if the RSVP item directly following the first target is omitted, thus introducing a 200-ms mask-free period, the second-target deficit is much reduced or eliminated (Chun & Potter, 1995; Raymond et al., 1992; Seiffert & Di Lollo, 1997). The deficit is restored, however, if the first target and the mask are presented simultaneously and spatially superimposed on each other, even if that integrated display is followed by a 200-ms mask-free period (Seiffert & Di Lollo, 1997). Thus, whether the first target is masked by a simultaneous or by a temporally trailing stimulus seems to be of little consequence to the second-target deficit.

Equivalence of simultaneous and delayed masks in respect to the second-target deficit is notable because the underlying masking processes are believed to differ substantially one from the other. The processing mechanisms subserving simultaneous and delayed masking are commonly referred to as *integration* and *interruption,* respectively (see reviews by Breitmeyer, 1984; Scheerer, 1973). The masking literature suggests that the two forms of masking act in very different ways and probably affect different phases of processing. In integration masking, the mask is said to degrade the target's earliest representation through a process equivalent to adding noise (the mask) to a signal (the target). Two characteristics of integration masking must be noted. First, masking is at a maximum when the target and the mask are presented simultaneously (i.e., when the stimulus onset asynchrony [SOA] is equal to zero). Furthermore, the strength of masking diminishes symmetrically as the SOA is increased, whether the mask leads (forward masking) or trails (backward masking). Second,

Barry Giesbrecht, Department of Psychology, University of Alberta, Edmonton, Alberta, Canada; Vincent Di Lollo, Department of Psychology, University of British Columbia, Vancouver, British Columbia, Canada.

This work was supported by Grant OGP0006592 from the Natural Sciences and Engineering Research Council of Canada. Portions of Experiments 1 and 2 were presented as a poster at the 37th Annual Meeting of the Psychonomic Society, Chicago, IL.

We thank James T. Enns, Alan Kingstone, Adriane E. Seiffert, and Kimron L. Shapiro for commenting on a draft of this article. We also thank Asif Butt and Jeanette Lum for collecting the data.

Correspondence concerning this article should be addressed to Barry Giesbrecht, Department of Psychology, P-220 Biological Sciences, University of Alberta, Edmonton, Alberta, Canada T6G 2E9. Electronic mail may be sent to blgiesbr@psych.ualberta.ca.

the strength of masking is not affected by the information content of the display, as indexed by the number of items in the target set (e.g., Spencer & Shuntich, 1970). In interruption masking, on the other hand, processing of the target is said to be disrupted by the arrival of the mask, which takes over those processing mechanisms that are required in common for both stimuli. Interruption masking is found only when the mask is trailing (i.e., there is no forward masking), and masking is at a maximum not at simultaneity but when the mask follows the target after an optimal SOA. Moreover, masking increases with the information content of the display, as indexed by the number of items in the target set (e.g., Spencer & Shuntich, 1970).

A similar distinction has been drawn by Turvey (1973), who referred to the two forms of masking as *peripheral* and *central*, corresponding broadly to integration and interruption masking, respectively. Peripheral masking was said to occur at an early stage in visual processing and to respond to such factors as intensity, duration, and the relative energies of target and mask. Central masking, on the other hand, was said to occur at higher processing levels, and to be influenced by more cognitive factors. There is substantial evidence in the masking literature that integration (peripheral) masking is instantiated optimally by presenting the target and the mask simultaneously, spatially superimposed on each other, and that interruption (central) masking is instantiated optimally when the mask follows the target after a suitable SOA (Bachmann & Allik, 1976; Michaels & Turvey, 1973; Purcell & Stewart, 1970; Schiller, 1966; Spencer & Shuntich, 1970).

In the present work, we made use of the differences between the two forms of masking to study the second-target deficit. If it is the case that masking by integration and interruption are subserved by different mechanisms, we reasoned that, by using the two forms of masking selectively, it may be possible to assess the relative roles of the salient mechanisms in the second-target deficit. For example, if the second-target deficit is obtained with Masking Procedure A but not with Masking Procedure B, then we can adopt the working hypothesis that the second-target deficit probably shares more mechanisms with the former than with the latter. At a first impression, however, this reasoning seems to be at odds with the available evidence. We have noted above that whether the first target is masked by integration (with a simultaneous mask) or by interruption (with a delayed mask) seems to make little difference to the second-target deficit. On this basis it could be concluded that the precise form of masking may not be an important consideration.

Such a conclusion, however, would be premature. At the very least, it would have limited generality, being based exclusively on results obtained with masking of the first target. To date, there have been no investigations of the second-target deficit in relation to the type of masking procedure used to mask the second target. To be sure, it is generally recognized by researchers in this area that in order to obtain a second-target deficit, it is necessary to mask the second target as well as the first. This is why in studies of the AB the second target is invariably followed by at least one distractor. However, the main reason for masking the second target has been to bring performance within a measurable range, thus avoiding ceiling effects (e.g., Moore, Egeth, Berglan, & Luck, 1996). In the present work, we show that the specific procedure used to mask the second target has important consequences for the AB deficit. To wit, the deficit is obtained with interruption masking, but it is much reduced or entirely absent with integration masking. This differential effect provides novel constraints for theories of how sequential stimuli are processed in the visual system.

Experiment 1

In Experiment 1 we used a conventional RSVP paradigm in which each item was displayed for 10 ms, and was followed by the next item after a blank interval of 90 ms. Within the RSVP stream, the two targets were uppercase alphabetical characters, and the distractors were digits of approximately the same size as the targets. The first target was always masked by the next item in the stream, as in conventional AB experiments (e.g., Shapiro, Raymond, & Arnell, 1994). In contrast, the second target was masked with either a simultaneous mask or a delayed mask.

In the simultaneous-masking condition, the second target was presented at the same time as, and spatially superimposed on, one distractor. Namely, the display consisted of the combined contours of the target letter and one digit, embedded one within the other so as to form a single image, as illustrated in Figure 1. In this condition, the target was assumed to be masked by a process of integration. In order to avoid any possibility of masking by interruption, no trailing distractors were presented after the second target. Thus, in the simultaneous-masking condition, the RSVP stream ended with the presentation of the second target embedded with a digit. In the delayed-masking condition, on the other hand, the second target was presented alone, followed by a single distractor after a 90-ms interval. In this condition, the target was assumed to be masked chiefly by a process of interruption. Our intent was to juxtapose the two conditions to distinguish the effects of integration masking and interruption masking in the second-target deficit.

Two control conditions were used. The first was a no-mask condition in which there were no trailing distractors after the second target. The effects of masking could thus be evaluated against a corresponding nonmasked control. The second control was a condition with a variable number of trailing distractors presented after the second target. This provided a comparison with conventional AB experiments, in which the number of distractors presented after the second target varied inversely with the temporal lag between the two targets (e.g., Shapiro et al., 1994).

Method

Participants. Twenty-four undergraduate students (9 men) from the University of Alberta participant pool participated for class credit. All reported normal or corrected-to-normal vision. None of the participants was involved in any of the other experiments reported here.

Apparatus and stimuli. All of the stimuli used in this and subsequent experiments were displayed on a Tektronix 608 oscilloscope equipped with P15 phosphor. The viewing distance was 57 cm, set by a headrest. All stimuli subtended approximately 1° of visual angle and had a luminance of 10 cd/m², as measured by a Minolta LS-100 luminance meter. The distractor items were digits (0–9), and the target items were letters from the English alphabet. The background and surrounding visual field were dark, except for dim illumination of the keyboard and response box.

Procedure. At the beginning of each trial a small fixation dot was presented in the center of the screen, indicating where the RSVP items would be presented. Participants initiated each trial by pressing a button on the response box. After a 500-ms delay, the RSVP stream was presented. Each item was displayed for 10 ms and was separated from the next item by an interstimulus interval (ISI) of 90 ms, yielding a presentation rate of 10 items/s. On any given trial, the distractors in the stream were selected randomly with replacement from the set of digits, with the constraint that the selected digit was not one of the two immediately preceding items. The letter targets were selected randomly without replacement from all letters of the English alphabet, except I, O, Q, and Z (these letters were omitted because of their visual similarity to 0, 1, 7, and 2). The number of distractors preceding the first target was determined randomly on each trial and varied between 7 and 15. The second target was presented at one of seven lags after the first target. Participants were required to identify the two letters in the stream and to enter them on the keyboard after the stream ended. Next, the fixation dot reappeared to indicate that the next trial was ready to begin.

Participants received four blocks of trials in a single 1-hr session. The four blocks differed with respect to how the second target was masked. In the no-mask condition, the RSVP stream ended with the second target. In the stream-mask condition, the number of distractors following the second target was the same as in conventional AB studies (i.e., the first target was always followed by eight items, one of which was the second target, with the constraint that the second target was never the last item in the stream). In the delayed-mask condition, the second target was followed by only one distractor, which was presented after the regular ISI (i.e., 90 ms). In the simultaneous-mask condition, the RSVP stream ended with a display containing the second target overlaid by one distractor, so that the two stimuli formed a single configuration that was integrated both spatially and temporally. The display sequences in each of the four conditions are illustrated in Figure 1.

In each block, the temporal lag between the onsets of the first and the second targets varied between 100 and 700 ms, in steps of 100 ms. The second target was presented 12 times at each of the seven lags, resulting in four blocks of 84 trials. The order of presentation of the experimental blocks was counterbalanced across all participants, so that each participant completed 1 of the 24 possible permutations of the four masking conditions. At the beginning of each session, participants completed 14 practice trials in the stream-mask condition. A brief rest period was introduced between blocks of trials.

Results

In this and all subsequent experiments, estimates of second-target identification were based only on those trials in which the first target had been identified correctly. This procedure is commonly adopted in AB experiments on the grounds that, on incorrect trials, the source of the error is unknown, so its effect on second-target processing cannot be

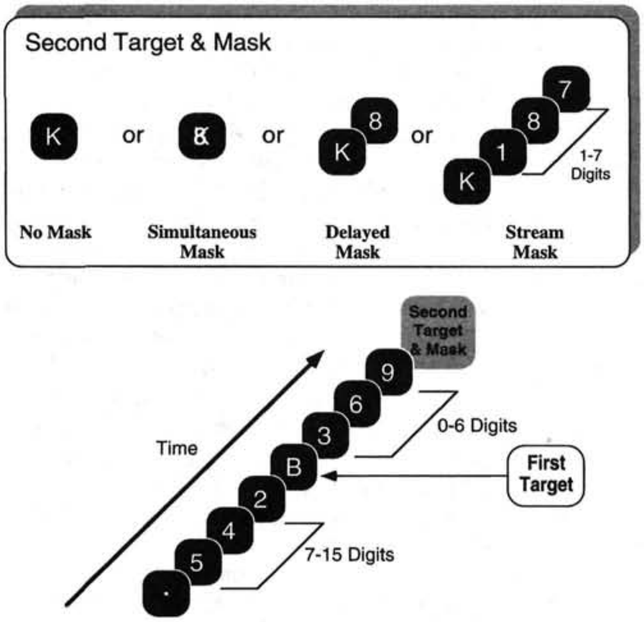

Figure 1. Schematic representation of the display sequences in Experiment 1. All stimuli were presented sequentially in the center of the screen. The first target was always masked by the next item in the sequence. The second target was masked in one of four ways: It was the last item in the stream (no mask), was presented simultaneously with a digit (simultaneous mask), was followed by a single digit (delayed mask), or was followed by between 1 and 7 digits (stream mask).

estimated. Responses were recorded as correct regardless of the order of report. Mean percentages of correct identifications of the first target, collapsed across lags, were 88.6, 88.3, 86.9, and 86.3 for the stream-mask, no-mask, delayed-mask, and simultaneous-mask conditions, respectively. Mean percentages of correct identifications of the second target as a function of lag, averaged over all participants, are presented in Figure 2. The results were analyzed with a 4 × 7 repeated measures analysis of variance (ANOVA), with two within-subjects variables: masking condition (stream mask, no mask, delayed mask, and simultaneous mask) and lag (100–700 ms). The ANOVA revealed significant effects of masking condition, $F(3, 69) = 29.07$, $p < .001$, $MSE = 797.34$, and lag, $F(6, 138) = 18.77$, $p < .001$, $MSE = 236.91$. The interaction effect was also significant, $F(18, 414) = 6.59$, $p < .001$, $MSE = 149.63$.

The comparison of major interest for the present purpose was between the simultaneous-mask and the delayed-mask conditions. Performance with the delayed mask (Figure 2A) revealed a pronounced second-target deficit whose temporal course was a U-shaped function commonly found in studies of the AB (Chun & Potter, 1995; Raymond et al., 1992; Seiffert & Di Lollo, 1997). Results with the simultaneous mask (Figure 2B) presented a fundamentally different picture: Overall performance was impaired, but the temporal course of the impairment was substantially different from that obtained with the delayed mask. A separate ANOVA performed on the data for the delayed-mask and the simulta-

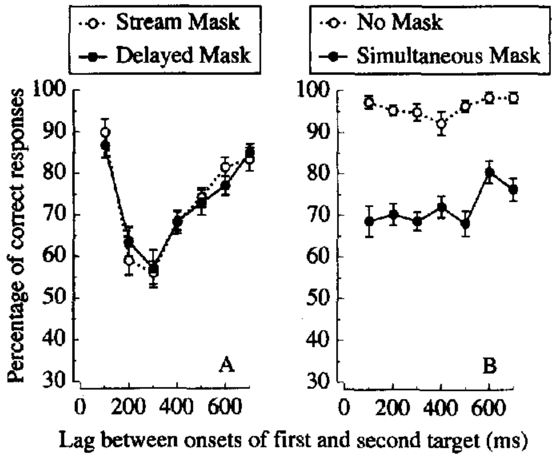

Figure 2. Results of Experiment 1. Mean percentages of correct identifications of the second target, given accurate identification of the first target. A: results of the stream-mask and delayed-mask conditions. B: results of the no-mask and simultaneous-mask conditions. Error bars represent 1 *SE* and are appropriate for within-subjects pairwise comparisons (Loftus & Masson, 1994).

neous-mask conditions revealed a significant effect of lag, $F(6, 138) = 8.93, p < .001, MSE = 237.70$, but not of mask type, $F(1, 23) < 1, MSE = 188.08$. There was a significant interaction effect between lag and mask type, $F(6, 138) = 6.71, p < .001, MSE = 188.08$, confirming that, although similar in overall mean level, the performance deficits obtained with the two forms of masking followed different time courses. At its lowest point, performance in the delayed-mask condition (see Figure 2A; lag = 300 ms) was significantly lower than performance at the corresponding lag in the simultaneous condition (see Figure 2B), $F(1, 138) = 8.71, p < .004, MSE = 188.08$. Considered separately, performance with the simultaneous mask showed a modest but significant increment over lags, $F(6, 138) = 2.83, p < .02, MSE = 184.47$. In contrast, lag had no significant effect in the no-mask condition, $F(6, 23) = 2.11, p > .05, MSE = 53.91$.

Discussion

The well-known AB deficit (Raymond et al., 1992) was replicated with a conventional RSVP paradigm in which the second target was masked by up to seven trailing distractors (see Figure 2A, filled symbols). The results of the delayed-mask condition showed that, just as in the simplified paradigm used by Duncan et al. (1994), multiple masks are not needed in the RSVP paradigm. That is, a single trailing item produced an AB deficit that was as strong as that produced by a stream of trailing items. What is more, if the trailing item was either omitted or integrated spatiotemporally with the second target, the AB deficit failed to appear. This was the key finding in Experiment 1, and deserves detailed examination.

At a strictly descriptive level, it can be said that performance on the second target was affected differently, depending on whether the mask was simultaneous or delayed. Beyond a descriptive level, it can be surmised that the pattern of results in Figure 2 may reflect a causal relationship between type of mask and second-target deficit. Specifically, it may be suggested that the deficit is mediated by processes that are triggered by delayed masks (i.e., interruption of processing), but not by simultaneous masks (i.e., stimulus degradation). Before this line of reasoning can be pursued with confidence, we must consider two aspects of the results that can create some ambiguity for interpretation: the fact that performance in the no-mask condition was near ceiling, and the finding of a significant improvement over lags in the simultaneous-mask condition. These are discussed in turn, below.

We noted earlier that performance in the no-mask condition did not vary significantly across lags. On inspection, however, the no-mask curve in Figure 2B exhibits a clear, if muted, U-shaped trend over lags, suggestive of a weak AB deficit. Because performance in this condition was near perfect, we must consider the option that a second-target deficit failed to be revealed because performance was compressed against the 100% limit imposed by the response scale. On this option, an AB deficit might well have been revealed with a less constrained response measure. Although plausible, this option is inconsistent with the results of the simultaneous-mask condition seen in Figure 2B. That is, performance with the simultaneous mask was well below ceiling, at a level where an AB deficit could have been measured, had it occurred. But the pronounced U-shaped deficit seen in Figure 2A failed to appear. Clearly, bringing performance within a measurable range cannot reveal an AB deficit unless the conditions for producing that deficit have also been met.

In a less plausible vein, it could be suggested that an AB deficit was latent in the no-mask condition and failed to appear when performance was brought within measurable range by using a simultaneous mask because the simultaneous mask itself might have prevented an AB deficit from appearing in some unspecified way. This option is dismissed in Experiment 2 in which an AB deficit was obtained under appropriate conditions even though the second target was degraded with a simultaneous mask. Rather, the pattern of results strongly suggests that the AB deficit hinges on the presence of at least one trailing item acting as a delayed mask on the second target.

Next, we consider the improvement in performance over lags in the simultaneous-mask condition (see Figure 2B). In the context of the AB, an improvement in performance over lags is thought to mirror the increasing availability of resources that can be deployed to the second target as processing of the first target nears completion. The results of Experiment 1 revealed a significant improvement over lags with both types of masks. With delayed masks (see Figure 2A), the improvement occurred at the longer lags, after a rapid decrement during the first 300 ms. With simultaneous masks (see Figure 2B), there was no initial decrement, and the improvement over lags was smaller, but it was statistically significant. The differences between the two performance curves are substantial, and justify the working hypothesis of different underlying mechanisms. By the same token, the possibility cannot be ignored that the improvement over lags in the simultaneous-mask condition might represent a recovery from some form of initial second-target deficit. This possibility would cloud the distinction between the two forms of masking because it would indicate that an AB deficit may be obtained not only with delayed masks but also when the second target is degraded with a simultaneous mask. One possible reason for the improvement over lags in the simultaneous-mask condition is examined in Experiment 2.

Experiment 2

In Experiment 1 we used two forms of masking with the intent of distinguishing between different processes underlying the second-target deficit. With simultaneous masks, we aimed at degrading the target's earliest representation by adding camouflage or noise. With delayed masks, our objective was to interfere with target identification through competition for higher level mechanisms. A critical requirement in achieving these objectives was to avoid contamination between the two types of masking processes. For example, if the simultaneous mask acted not only to degrade the low-level representation of the target but also to interfere with its processing at a higher level, then the deficit in performance could not be ascribed unambiguously to either process.

It is possible that just such a contamination may have arisen in the simultaneous-mask condition in Experiment 1. Because the mask consisted of a meaningful stimulus (a digit), the resulting interference with target identification could have arisen from at least two sources: degradation of the target by visual noise early in processing, and competition between two meaningful items (a letter and a digit) at a higher level of processing. From this perspective, the second-target deficit obtained with the simultaneous mask (see Figure 2B) cannot be ascribed solely to degradation of the target's earliest representation, but must be ascribed, at least in part, to interference at a higher processing level. On this reasoning, the improvement over lags can be understood in terms of the cost inherent in processing multiple items at the same time, as suggested by Duncan (1980; Duncan et al., 1994). To wit, on trials in which the second target was presented directly after the first target, three items (two letters and one digit) competed for the same high-level analyzers, and performance suffered accordingly. At longer lags, more processing of the first target could be accomplished, thus freeing up resources for processing the second target, with consequent improvement in performance.

One way of resolving this ambiguity is to use a masking stimulus that, although degrading an early representation of the target, does not introduce another meaningful stimulus to compete with the target at a higher processing level. In Experiment 2 this was done by using a meaningless aggregate of random dots instead of a digit in the simultaneous-mask condition.

A second objective of Experiment 2 was to evaluate the AB deficit obtained with simultaneous and delayed masks against corresponding control conditions in which the RSVP stream contained only one target. The single-target controls are needed because, at the most basic level, demonstration of an AB deficit requires a comparison between two conditions. In the experimental condition, the RSVP stream contains two targets; in the control condition, the first target is replaced by a distractor. Thus, the control stream contains only one target, whose location corresponds to that of the second target in the experimental stream. If performance on the second target is found to be lower in the experimental than in the control condition, the deficit can be ascribed to the effect of the first target.

Three masking procedures were explored in Experiment 2, each comprising an experimental and a control condition. The first was a replication of the delayed-mask condition in Experiment 1, with the addition of the corresponding control condition. On the basis of earlier results (Raymond et al., 1992; Seiffert & Di Lollo, 1997), we anticipated large differences between control and experimental conditions, with near-perfect performance in the former, and a pronounced AB deficit in the latter. The second procedure was a replication of the simultaneous-mask condition in Experiment 1, with a crucial modification: Instead of a digit, the mask consisted of an aggregate of random dots overlaid on the second target. This ensured that masking occurred through degradation of the target's earliest representation, as distinct from interference between meaningful items at a higher processing level. We anticipated that such a mask should impair identification of the target both in the experimental and in the control conditions. More important, to the extent that simultaneous masking by noise does not mediate an AB deficit, we expected the level of performance to be the same in the two conditions. The third masking procedure consisted of a combination of simultaneous and delayed masks. In that condition, the second target was overlaid with random dots, as in the simultaneous-mask condition, and was also masked by a digit displayed 90 ms later, as in the delayed-mask condition. The objective was to rule out the option, discussed in Experiment 1, that the presence of a simultaneous mask might prevent the occurrence of an AB deficit in some unspecified way. We expected the simultaneous mask to reduce the level of performance equally across lags, and the delayed mask to produce an AB deficit. To the extent that the two effects combined additively, we expected to see an AB deficit at a lower overall level than that seen with a delayed mask alone.

Method

Procedures were the same as in Experiment 1, with the following exceptions. Thirty undergraduate students viewed RSVP streams in six conditions, which were grouped in two sets. In one set the first target was present (present set), in the other it was absent (absent set). For ease of terminology, we refer to the only target in the absent set as *second target* because the only target in the absent set and the second target in the present set were presented in corresponding positions within the RSVP streams. In both the present and the absent sets, the second target was masked in three different ways: by embedded dots, by a trailing item, and by both embedded dots and a trailing item. In the simultaneous-mask conditions, the RSVP stream ended with the second target embedded in a patch of 100 dots that were positioned randomly on every trial within a notional square of 1° side. The number of dots used in the mask was determined by a pilot study in which we varied the number of dots in the mask, and the task was to identify the only target in an RSVP stream. We selected a 100-dot mask because it yielded a mean level of performance similar to that in the simultaneous-mask condition in Experiment 1. In the delayed-mask conditions, the second target was followed by a single digit at an ISI of 90 ms. This was the same as the delayed-mask condition in Experiment 1. In the combined conditions, the second target was embedded in a patch of 100 dots (as in the simultaneous-mask condition) and was followed by a single digit at an ISI of 90 ms (as in the delayed-mask condition). In all conditions, the second target was presented 10 times at each of the seven lags used in Experiment 1. The streams in the absent set were the same as in the present set, except that the first target was replaced by a digit. Thus, the lag between the two targets in the absent set should be regarded as the lag of the second target relative to when the first target would have been presented, had it been included in the stream. The display sequences in each of the six conditions are illustrated in Figure 3.

Overall, this design resulted in two blocks of 210 trials, which were completed during a single 1-hr session. Before each block, participants completed 15 practice trials in the appropriate set for that block (present or absent) and a random assortment of the three masking conditions. Order of presentation of present and absent sets was counterbalanced across participants.

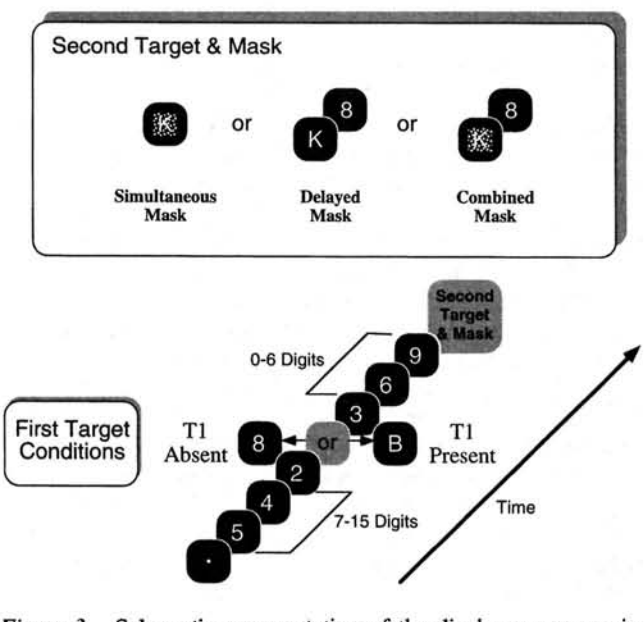

Figure 3. Schematic representation of the display sequences in Experiment 2. All stimuli were presented sequentially in the center of the screen. The first target (T1) was either present or absent, in which case the target letter was replaced with a digit. The second target was masked in one of three ways: It was presented simultaneously with a 100-dot mask (simultaneous mask), was followed by a single digit (delayed mask), or was presented simultaneously with the dot mask and was then followed by a digit (combined mask).

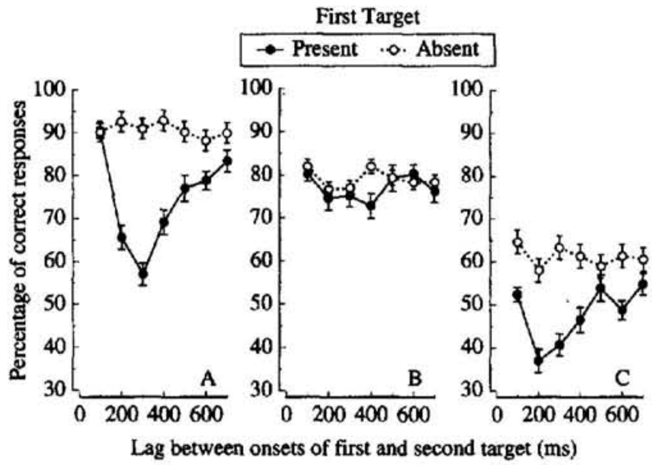

Figure 4. Results of Experiment 2. Scores in the present conditions are mean percentages of correct identifications of the second target, given accurate identification of the first target. Scores in the absent conditions are mean percentages of correct responses. Notional lags for the absent condition were devised on the basis of the way in which the rapid serial visual presentation streams were constructed (see text). A, B, and C: Results of the delayed-mask, simultaneous-mask, and combined-mask conditions, respectively. Error bars represent 1 *SE* and are appropriate for within-subjects pairwise comparisons (Loftus & Masson, 1994).

Results

Mean percentages of correct identifications of the second target as a function of lag, averaged over all participants, are presented in Figure 4. A problem arises when comparing present and absent conditions across lags. Because there was only one target in the absent conditions, by definition there could be no lag between targets. For the sake of comparison with the present condition, notional lags can be devised on the basis of the way in which the RSVP streams were constructed. To wit, the present and absent streams differed in a single detail: In the latter, the first target letter was replaced with a digit. Therefore, notional intertarget lags can be specified for the absent conditions in terms of the temporal interval that elapsed from the presentation of the digit that replaced the first target and the presentation of the second target on any given trial. This has been done in Figure 4.

Mean percentages of correct identifications of the first target, collapsed across lags separately for each of the three present conditions were as follows: simultaneous mask, 87.4; delayed mask, 88.0; combined mask, 89.3. The results in Figure 4 were analyzed in a 2 (first target present or absent [P/A]) × 3 (mask: simultaneous, delayed, or combined) × 7 (lag: 100–700 ms) repeated measures ANOVA. All main effects and interactions were significant, with the exception of one interaction effect of borderline significance: P/A, $F(1, 29) = 49.82$, $p < .001$, $MSE = 716.86$; mask, $F(2, 58) = 182.48$, $p < .001$, $MSE = 527.81$; lag, $F(6, 174) = 8.14$, $p < .001$, $MSE = 260.91$; P/A × Mask, $F(2, 58) = 27.69$, $p < .001$, $MSE = 209.84$; P/A × Lag, $F(6, 174) = 7.36$, $p < .001$, $MSE = 254.07$; Mask × Lag, $F(12, 348) = 1.75$,

$.05 > p < .06$, $MSE = 716.86$; P/A × Mask × Lag, $F(12, 348) = 3.29$, $p < .001$, $MSE = 193.30$.

Separate ANOVAs were performed on the data in each panel of Figure 4. The objective was to carry out a direct comparison between present and absent conditions, separately for each masking procedure. The results were as follows. For the results in Figure 4A (delayed mask), all effects were significant: P/A, $F(1, 29) = 82.94$, $p < .001$, $MSE = 333.37$; lag, $F(6, 174) = 10.30$, $p < .001$, $MSE = 164.60$; P/A × Lag, $F(6, 174) = 12.34$, $p < .001$, $MSE = 184.99$. For the data in Figure 4B (simultaneous mask), no effects were significant: P/A, $F(1, 29) = 1.48$, $p > .23$, $MSE = 344.01$; lag, $F(6, 174) = 1.18$, $p > .31$, $MSE = 198.04$; P/A × Lag, $F(6, 174) = 1.00$, $p > .42$, $MSE = 180.16$. For the data in Figure 4C (combined mask), all effects were again significant: P/A, $F(1, 29) = 41.76$, $p < .001$, $MSE = 459.16$; Lag, $F(6, 174) = 3.08$, $p < .007$, $MSE = 280.43$; P/A × Lag, $F(6, 174) = 2.47$, $p < .03$, $MSE = 275.52$.

Discussion

In Experiment 2 the second target was masked in three different ways. In each case, the effects of masking were examined by using two RSVP streams: an experimental stream that contained two targets, and a control stream that contained only one target whose temporal position in the display corresponded to that of the second target in the experimental stream. An AB deficit was said to have occurred if identification of the second target was poorer in the experimental than in the control stream. On the basis of this comparison, the deficit in identifying the second target could be ascribed to the requirement of having to process the first target. For an alternative control procedure, we could have used an RSVP stream containing two targets, with the participants being asked to ignore the first one and to respond only to the second one. In this procedure, however, the first target would have been processed to an unknown extent, thus confounding the comparison with the experimental procedure. This confounding was obviated by using a single-target control condition.

Performance in the control conditions was highest with the delayed mask (see Figure 4A), lower with the simultaneous mask (see Figure 4B), and lowest with the combined mask (see Figure 4C). In each case, performance remained at a steady level across lags. A vastly different pattern emerged in the experimental conditions. Pronounced AB deficits were obtained only in the two conditions in which the mask had a delayed component (see Figures 4A and 4C). In both instances, the performance curves were similar to those obtained in other AB studies (Chun & Potter, 1995; Raymond et al., 1992; Seiffert & Di Lollo, 1997). Namely, performance was closest to the control levels at the shortest and longest lags, and showed the largest deterioration at lags of 200 and 300 ms. In contrast, no AB deficit was obtained with the simultaneous mask (see Figure 4B): Performance in the experimental condition did not differ significantly from that in the control condition either in overall level or in temporal course over lags.

This pattern of results confirms the key finding in Experiment 1: A pronounced AB deficit can be obtained with delayed masking but not with simultaneous masking of the second target. Further, the joint results of Experiment 1 and 2 indicate that simultaneous masking is ineffective in bringing about an AB deficit, whether the masking stimulus is a meaningful item, as in Experiment 1, or a meaningless set of random dots, as in Experiment 2. In this respect, it is interesting to note that the significant improvement over lags seen in the simultaneous-mask condition in Experiment 1 (see Figure 2B) was not found in the corresponding condition in Experiment 2 (see Figure 4B), $F(6, 174) = 1.35$, $p > .23$, $MSE = 202.30$. Namely, the improvement occurred when the mask was meaningful but not when it was meaningless. This is what would be expected if the improvement over lags mirrored the cost inherent in processing multiple items at the same time, as was suggested in the introduction of the present experiment.

One further point should be noted. The results in Figure 4C show that an AB deficit can be obtained even if the second target is degraded with a simultaneous mask, provided that it is followed by a delayed mask. On the evidence in the three sections of Figure 4, it looks as though simultaneous and delayed masks may combine in broadly additive fashion in their effects on performance. This supports the commonly held view that the two forms of masking occur at different processing stages (e.g., Breitmeyer, 1984; Turvey, 1973). In addition, this finding is inconsistent with the option, noted in the *Discussion* section of Experiment 1, that simultaneous masks may act to suppress the AB deficit.

Experiment 3

Simultaneous masking of the second target did not result in any AB deficits in Experiments 1 and 2, whether the mask was meaningful or meaningless. On the other hand, pronounced AB deficits were obtained in both experiments if the mask was delayed. In every case, however, the delayed mask consisted of a meaningful stimulus, namely a digit. This raises the question of whether the delayed mask needs to be meaningful, or whether an AB deficit can be obtained with a meaningless mask, provided that it is delayed. An answer to that question was sought in Experiment 3.

In Experiment 3, meaningless sets of random dots were used to mask the second target in RSVP streams. In any given stream, the mask was either embedded with the second target, as in the simultaneous condition in Experiment 2, or it trailed the second target by a variable interval. Consistent with Experiment 2, no AB deficit was found when the target and the mask were displayed simultaneously. However, significant AB deficits were obtained when the mask was delayed.

Method

Procedures were the same as in the previous experiments, with the following exceptions. Thirty undergraduate students from the University of British Columbia viewed RSVP streams in four conditions. The conditions differed in respect to the SOA between the second target and a dot mask. The SOAs between the second target and the mask were 0, 50, 100, and 200 ms. The mask was the same as in the simultaneous-mask condition of Experiment 2: It consisted of a patch of 100 dots positioned randomly on every trial within a notional square of 1° side. Thus, the 0-SOA condition in this experiment was exactly the same as in the simultaneous-mask experimental condition of Experiment 2. Similarly, the 100-ms SOA condition was the same as the delayed-mask experimental condition in Experiment 2, except that the mask consisted of a meaningless group of dots instead of a digit. In all conditions, the second target was presented 15 times at each of the seven lags used in the previous experiments. The order of presentation of the different conditions was randomized.

Overall, this design resulted in 420 trials, which were completed during a single 1-hr session. At the beginning of the session, participants completed a block of 15 practice trials that consisted of a random assortment of the different conditions. Within the testing block, the order of presentation of the different conditions was randomized for each participant. The experimental trials were separated into four blocks, and participants were offered a break between blocks.

Results and Discussion

Mean percentages of correct identifications of the second target as a function of lag, averaged over all participants, are presented in Figure 5, separately for each SOA. Mean percentages of correct identifications of the first target, collapsed across lags, were 90.3, 89.4, 87.6, and 87.3 ms for mask SOAs of 0, 50, 100, and 200 ms, respectively. The results in Figure 5 were analyzed with a 4 (SOAs) $\times$ 7 (Lags) repeated measures ANOVA. All effects were significant: SOA, $F(3, 87) = 58.41, p < .001, MSE = 349.69$; lag, $F(6, 174) = 9.10, p < .001, MSE = 122.39$; SOA $\times$ Lag, $F(18, 522) = 1.76, p < .03, MSE = 88.93$.

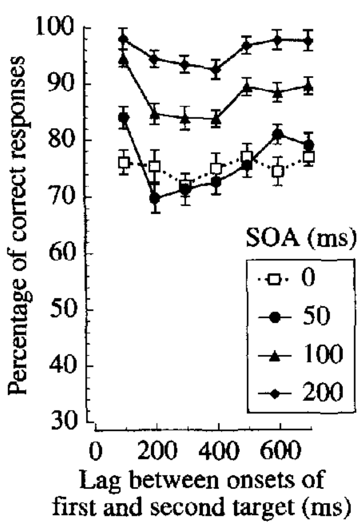

Figure 5. Results of Experiment 3. Mean percentages of correct identifications of the second target, given accurate identification of the first target. Stimulus onset asynchrony (SOA) refers to the temporal interval that elapsed from the onset of the second target to the onset of the mask. Error bars represent 1 *SE* and are appropriate for within-subjects pairwise comparisons (Loftus & Masson, 1994).

No AB deficit was found when the target and the mask were displayed simultaneously (see Figure 5, SOA = 0). A separate ANOVA showed the effect of lag in the 0-SOA condition to be nonsignificant, $F(6, 174) < 1$. This parallels the results obtained with the simultaneous mask in Experiment 2. On the other hand, significant AB deficits were obtained at all other SOAs, including 200 ms, $F(6, 174) = 3.18, p < .01, MSE = 48.07$. Notably, a separate ANOVA carried out on the data for the 0-SOA and 50-ms SOA conditions revealed a significant interaction effect between SOA and lag, $F(6, 174) = 2.78, p < .01, MSE = 132.10$. This shows that the introduction of a 50-ms SOA is sufficient to bring about an AB deficit.

On the basis of these results, it can be concluded that to bring about an AB deficit, a delayed mask must be presented after the second target; but the mask itself need not be a meaningful stimulus. This said, it must be noted that although meaningful masks are not necessary, they do yield larger AB deficits. This is revealed by a direct comparison between the delayed-mask condition in Experiment 2 (see Figure 4A, filled symbols) and the 100-ms SOA condition in Experiment 3. The SOA in the two conditions was the same, but the mask was meaningful in Experiment 2 and meaningless in Experiment 3. A between-subjects ANOVA revealed a significant interaction effect between lag and mask meaningfulness, confirming that the two performance curves followed different time courses, $F(6, 348) = 5.32, p < .001, MSE = 174.14$.

Before concluding the exposition of the empirical data, a curious coincidence should be noted. It pertains to the temporal course of performance in the simultaneous-mask and no-mask conditions across all three experiments. Consider the 0-SOA curve in Figure 5: The trend over lags resembles that in an AB deficit. Namely, performance declined over the first three lags and it recovered thereafter. Similar trends can be detected in the simultaneous-present curve in Figure 4B and in the no-mask curve in Figure 2B. This raises the possibility that a minimal AB-like deficit may have occurred when the second target was masked with a simultaneous mask or when it was not masked at all. Admittedly, the effect of lag was not statistically significant in any of these instances, and the U-shaped trend was entirely missing in the simultaneous-mask curve in Figure 2B. Besides, the magnitude of the trend, when it occurred, was negligible in comparison to that obtained with delayed masks. Nevertheless, we felt that the coincidence was worth noting in case it reappears in future studies.

General Discussion

Investigations of the AB deficit have concentrated primarily on variables that affect the processing of the first target. This is understandable because processing of the first target is essential for obtaining an AB deficit. The present work shows that variables that affect the processing of the second target are just as important. Two key findings emerged from the experimental work. First, to obtain an AB deficit, the second target needs to be masked. Second, the form of masking is important: An AB deficit is obtained when the mask is delayed but not when it is presented simultaneously with the second target. To be sure, accuracy in identifying the second target is impaired whether the mask is simultaneous or delayed, but a second-target deficit time-locked to the first target is obtained only with delayed masking.

Two Stages of Processing

A comprehensive account of these results can be provided in terms of a two-stage model proposed by Chun and Potter (1995). In that account, processing is said to occur in two sequential stages. The first is a rapid detection stage, where potential targets are detected on the basis of specific features (e.g., color, letter case) or on the basis of category. The second is a capacity-limited stage in which items are processed serially for subsequent report. Potential targets detected in Stage 1 gain access to Stage 2 only if the latter is not busy. If the second target arrives while Stage 2 is busy, it is delayed in Stage 1 until Stage 2 is free. During the period of delay, the representation of the second target is subject to deterioration through passive decay and through erasure by subsequent items. The representation becomes immune from deterioration once it is selected for further processing and for consolidation in Stage 2. The AB deficit is said to stem from the deterioration that occurs while the second target is delayed in Stage 1.

This account has proved capable of explaining the major findings in the AB literature. For example, it is known that the magnitude of the AB deficit increases with the difficulty of the first target (Chun & Potter, 1995). This is explained on the assumption that the period for which Stage 2 is kept busy with the first target increases as the target's difficulty is increased. In turn, the second target is delayed in Stage 1 for a correspondingly longer period, during which it is subject to deterioration. Another aspect of the results, which at first seems inconsistent with this account, can be explained on an additional assumption. Namely, a simple version of the two-stage model predicts that the AB deficit should be greatest at the shortest lag because, all other things being equal, the period of delay in Stage 1 should vary inversely with lag. Yet, it has been found that the second-target deficit is greatest not at the shortest lag (typically 100 ms) but at considerably longer lags. To account for this result, it has been hypothesized that, from the instant at which an item enters Stage 2, access is denied to subsequent items not immediately but gradually, over a period that is typically longer than 100 ms. Thus, when the second target is presented directly after the first, there is a finite probability that both targets may be processed concurrently in Stage 2. Similar principles governing the relationship between successive stimuli have been proposed by Raymond et al. (1992) and by Weichselgartner and Sperling (1987). These temporal contingencies give rise to the U-shaped function over lags that is commonly found in AB experiments.

According to the two-stage model, there are two sources of deterioration for any given item in Stage 1: passive decay and erasure by temporally trailing items (Chun & Potter, 1995). The present results suggest that the main source of deterioration in Stage 1 is not passive decay but erasure by the item presented directly after the second target. Passive decay would be evidenced by an impairment at the shortest lag, followed by a gradual improvement over lags. Had passive decay been an important factor, this trend should have been observed in the no-mask condition in Experiment 1, and the simultaneous conditions in Experiments 1 and 2, in which there were no trailing items to erase the second target. However, the salient curves in Figures 2B and 4B show little evidence of such a trend over the first few lags, suggesting that passive decay does not play a major role in the deterioration of the second target.

Erasure—or its absence—provides a consistent account of the results obtained both with the simultaneous and with the delayed masks. With simultaneous masks, performance was impaired because the figural properties of the stimulus were impoverished, thus making it harder to extract the target from the noise. However, the impairment was not time-locked to the first target, namely, there was no evidence of an AB deficit. Within the two-stage model, this result can be explained by noting that the second target and its mask form a unitary stimulus that remains available in Stage 1 throughout the period of delay (i.e., while the first target is being processed in Stage 2) because there is no trailing stimulus to erase it. Thus, what gains access to Stage 2 after the delay is a unitary representation of the second target, albeit embedded in noise. In contrast, performance with delayed masks was clearly time-locked to the first target. Namely, performance was impaired at lags of 200 and 300 ms, and improved progressively thereafter. According to the two-stage model, performance was impaired at the shorter lags because the target was erased by the trailing mask during the period of delay in Stage 1. The probability that the second target could enter Stage 2 before being erased by the mask increased at the longer lags, and performance improved accordingly.

When the two forms of masking are combined, as was done in the combined-mask condition in Experiment 2, the two-stage model provides a consistent account of the results. In that condition, the second target was either absent from the RSVP stream, or it was present, in which case it was embedded in a patch of dots and was followed by a single digit. When the first target was absent, overall performance was relatively low (see Figure 4C), and no AB deficit was obtained. In terms of the two-stage model, the low level of performance was the result of a degraded stimulus gaining entry into Stage 2. Notably, entry was direct because, given the absence of the first target, Stage 2 was free. Thus, not having been delayed in Stage 1, the target was never vulnerable to erasure, and an AB deficit was avoided. On the other hand, when the first target was present, not only was performance quite low, but it was clearly time-locked to the first target. This result is explained in the same terms as the results of the delayed-mask condition. Namely, while Stage 2 was occupied by the first target, the degraded second target was delayed in Stage 1. At the shorter lags, while delayed in Stage 1, the second target remained vulnerable to erasure by the delayed digit. At the longer lags, the probability that the second target could enter Stage 2 before being erased increased, and performance improved accordingly. This improvement in performance, however, could not exceed the level set by the simultaneous mask, namely the level achieved when the first target was absent from the RSVP stream.

Another finding that can be readily accommodated within the two-stage model is an asymmetry in the way in which the two targets may be masked to obtain an AB deficit. That is, we have shown that an AB deficit is obtained with delayed but not with simultaneous masking of the second target. In contrast, an AB deficit is obtained with either type of masking of the first target (Seiffert & Di Lollo, 1997). Why is type of masking important for the second target but not for the first? In the preceding discussion one can see how the two-stage model accounts for the masking asymmetry in respect to the second target. The masking equivalence in respect to the first target is handled on the principle that any procedure that increases the difficulty of first-target processing will also increase the period for which Stage 2 remains busy, and therefore the period for which the second target remains vulnerable to erasure while delayed in Stage 1. This principle has been amply demonstrated by Chun and Potter (1995), and is supported by the negative correlation between level of performance on the first target and the magnitude of the AB deficit (Grandison, Ghirardelli, & Egeth, 1997; Seiffert & Di Lollo, 1997). It is plausible to expect that the difficulty of processing the first target was increased with either type of mask. The prolongation of Stage-2 processing then led to a longer delay for the second target with a corresponding increment in the probability of erasure and consequent AB deficit.

It is clear that the two-stage model can account adequately not only for the pattern of results obtained in the present work but also for the major results in the AB literature. We should note, however, that results obtained in some very recent studies may not be as readily interpretable, and that the model may need to be revised to accommodate them. Those findings, and their implications for the two-stage model, are reviewed below, after two alternative accounts of the AB deficit have been considered.

Alternative Accounts

Two other accounts of the AB deficit have been proposed. One is based on competition amongst items in visual short-term memory (VSTM; Raymond et al., 1992; Raymond, Shapiro, & Arnell, 1995). The other is based on the concept of attentional dwell time (Duncan et al., 1994). According to the competition model, not all items from the RSVP stream gain entry into VSTM. To enter VSTM, items must match preset templates corresponding to the two targets. This parallels the process of rapid detection (Stage 1) in the two-stage model (Chun & Potter, 1995). In addition, the items directly following the two targets also gain entry because of temporal contiguity. Items in VSTM are assigned weights in accordance with the goodness of match with their corresponding template. Finally, attentional resources are allocated to individual items according to two criteria: the item's weight, and the item's order of entry into VSTM. Thus, according to this scheme, the first target is assigned the largest amount of resources, and the item next to the second target receives the least. The finding that the AB becomes smaller as the intertarget lag is increased is explained as follows. If the lag is relatively short, the leading items (the first target and the next item) are still active in VSTM. This prevents sufficient attentional resources from being allocated to the second target, even though its weight is relatively high. In this case, an AB deficit occurs because the meager attentional resources allocated to the second target do not allow it to compete effectively with the other items in VSTM. Conversely, if the lag is long, decay of the first target in VSTM releases attentional resources that can be used for processing the second target, and the AB is reduced.

With reference to the present findings, the competition model would predict correctly that an AB deficit should be obtained with delayed masks because the mask itself would compete with the two targets in VSTM. The competition model would also predict a reduction in the AB deficit when the second target is not masked, because there would be no mask item in VSTM to compete with a representation of the second target. But there are at least two important aspects of the results that are problematic for the competition model.

First, the competition model has no provision for explaining why an AB deficit occurs with delayed masking but not with simultaneous masking of the second target. To be sure, this does not constitute disconfirmation, but it does show that the model is incomplete. This incompleteness is further emphasized by the predictions of the model that run afoul of the empirical results. For instance, the model predicts that items such as digit masks that do not match the target template can nevertheless gain access to VSTM provided that they are presented in close temporal proximity to the target. The probability of the mask being admitted to VSTM diminishes as its temporal separation from target is increased. Thus, a simultaneous mask should gain access to VSTM far more easily than a delayed mask, because of its temporal contiguity with the target, and therefore produce a larger AB deficit than a delayed mask. This is the opposite of what was found.

A second finding that the competition model is strained to explain is that there were no significant differences between the simultaneous-mask and the corresponding control condition in Figure 4B. The model would predict lower performance in the experimental than in the control condition on the basis of the contents of VSTM on each condition. At a first approximation, VSTM in the experimental condition should contain about four items: the two targets, plus the item directly following each target. In contrast, VSTM in the control condition should contain only half that number of items because the first target is never presented. It follows that competition for retrieval of the second target in VSTM should be far greater—and performance correspondingly lower—in the experimental condition. But that was not the case.

In a more recent account of the AB deficit, Duncan et al. (1994; Ward, Duncan, & Shapiro, 1996) have suggested that attention remains concentrated on the first target for several hundred milliseconds before it can be shifted to the second target. During this attentional dwell time, the second target cannot be processed adequately, and its identification suffers correspondingly. This account shares broad similarities with that of Chun and Potter (1995) in that processing of the second target is said to suffer while the system is busy with the first target. However, in its present form, the account must be regarded as incomplete because it has no means of explaining why an AB deficit is obtained with delayed but not with simultaneous masks, or why no deficit is found if the second target is not masked.

Parallels Between the AB Deficit and Masked Priming

Throughout the present work, we distinguished between two masking processes: integration (peripheral) and interruption (central). Following convention (Bachmann & Allik, 1976; Breitmeyer, 1984; Scheerer, 1973; Spencer & Shuntich, 1970; Turvey, 1973), we instantiated the process of integration with simultaneous masks, and the process of interruption with delayed masks. An AB deficit was consistently obtained with delayed but not with simultaneous masking of the second target. We inferred from these results that an AB deficit occurs when the processing of the second target is disrupted by a trailing mask.

A logical next step is to consider what attributes of the target might be affected, and what stages of processing might be disrupted by the arrival of the mask. In the traditional view, backward masking is said to interfere with—or terminate the processing of—information at a precategorical level of stimulus representation (e.g., Scheerer, 1973; Turvey, 1973). The recent masking literature, however, suggests otherwise. The evidence (some of which is reviewed below) strongly suggests that under conditions of backward masking, processing of the target can continue beyond precategorical levels to lexical and semantic levels. What appears to be disrupted by the mask is not precategorical information, but the kind of information needed for direct report of the stimulus. Much of the evidence comes from studies of a phenomenon known as masked priming, which is described below. Upon comparison, we found clear points of contact between the experimental literatures on masked priming and on the AB deficit. Below, we suggest that the two sets of outcomes may provide converging evidence toward an understanding of how backward masking affects stimulus processing and, more generally, how the visual system handles rapid sequential inputs.

In conventional priming experiments, a brief display of a temporally leading word (the prime) facilitates the identification of a trailing target word, provided that the two words are semantically related (Meyer, Schveneveldt, & Ruddy, 1975). This is taken as evidence that the semantic activation produced by the prime facilitates the processing of a semantically related target. A similar stimulus sequence is used in masked-priming experiments, except that the prime is backward-masked so that the participant cannot report it. The fascinating and counterintuitive finding is that much the same type of facilitation is found in masked priming as in conventional priming experiments (Carr & Dagenbach, 1990; Cheesman & Merikle, 1986; Dagenbach, Carr, & Wilhelmsen, 1989; Marcel, 1983a, 1983b). This equivalence of outcomes suggests that the prime is capable of producing semantic activation even if it is masked. In turn, as noted by Marcel (1983a), this strongly suggests that backward masking disrupts the conscious registration of a stimulus but not its visual analysis.

A similar conclusion was reached on the basis of electrophysiological evidence in a recent study of the AB deficit. Luck, Vogel, and Shapiro (1996) recorded event-related brain activity during an AB experiment. They were particularly interested in the N400 wave, whose amplitude increases with the degree of incongruity between a test word and its semantic context. The presence of an N400 wave indicates two things: first, that the test word was perceived as being incongruous with its semantic context. Second, and more important for the present purpose, that the test word had indeed been processed to a semantic level. In Luck et al.'s study, the test word occupied the position of the second target in an RSVP stream, thus the participants were unable to report it on at least some of the trials. Yet, the N400 component was very much in evidence when the semantic context was incongruous even when the participant was unable to report the test word. This strongly suggests that the second target was processed to a semantic level even though it could not be reported accurately. More important, this outcome suggests that what is disrupted in the AB deficit is the information required to make an overt identification response, not the information accrued in the course of processing the second target.

The parallel between masked priming and the AB deficit is compelling. In both paradigms, backward masking has been shown to disrupt the overt response to—and perhaps the conscious registration of—the target, but not its processing to high lexical and semantic levels. Bearing in mind that the evidence is suggestive rather than definitive, we can formulate a working hypothesis that masked priming and the AB deficit may be mediated by mechanisms with a good deal of commonality. This hypothesis is strengthened by a second parallel between masked priming and the AB deficit, a parallel that stems from the main outcome of the present work. Just as the AB deficit is obtained with delayed but not with simultaneous masking of the second target, so does masked priming occur with interruption (central) but not with integration (peripheral) masking of the prime (Marcel, 1983b; Experiment 4).

These parallels justify—and invite explicit testing of—the hypothesis that masked priming and the AB deficit are different expressions of the same thing. A test that readily suggests itself is one in which the second target in an AB stream is used to prime a third target presented later in the stream. In fact, evidence for just this kind of priming has been reported by Shapiro, Driver, Ward, and Sorensen (1997), who found that the second target, even though unreportable because of its position in the RSVP stream, was capable of priming a third target presented shortly afterwards. That is, performance on the third target was better when the second and third targets were semantically related than when they were not. Thus, just as in masked priming, a sizable amount of information about the second target remained available within the system, even though the target itself could not be reported because it was backward-masked by the next item in the stream. Results entirely consistent with this conclusion have been reported by Maki, Frigen, and Paulson (1997). Pursuing the parallel between masked priming and the AB a step further, and bearing in mind the outcome of the present work, it should be expected that priming of the third target would occur when the second target is masked by interruption (as in Shapiro et al.'s, 1997, study) but not when it is masked by integration. Such a study remains to be done.

A Revised Two-Stage Model

It can be inferred from the preceding evidence that in studies of the AB, the information accrued during the visual analysis of the second target is not totally erased by a trailing mask. This creates a problem for the two-stage model proposed by Chun and Potter (1995). Especially intractable within the model is the finding by Shapiro et al. (1997) that a second target that cannot be reported accurately can nonetheless act as a prime for a third target. In considering how this might be handled within the two-stage model, the following question needs to be asked: If it is true that the second target could not be reported because it had been erased while waiting in Stage 1, how come it was still capable of priming a semantically related third target?

As presently stated, the two-stage model cannot provide a plausible answer. A simple revision, however, enables the model to handle all of the salient evidence. The revision introduces an intermediate stage that could be regarded as a holding buffer where the output of Stage 1 can be stored if Stage 2 is busy. Except for the addition of the holding buffer, the two-stage model would remain unchanged. Within this system, the sequence of processing events is as follows. Incoming stimuli are processed in Stage 1 in the manner proposed by Chun and Potter (1995). Namely, processing is in parallel and can include some lexical and semantic as well as sensory attributes of the stimulus. The encoded representation is then transferred to the holding buffer where it replaces (erases) the previous contents. The representation remains in the buffer until it gains access to Stage 2 or until it is replaced (erased) by the next input from Stage 1. In this fashion, a trailing mask can erase the representation of the second target in the holding buffer, but need not interfere with the residual activity of Stage-1 mechanisms that had been triggered while the second target was being processed at that level. It goes without saying that such residual activity may then be used to mediate priming of a third target. In passing, it is worth noting that although our suggestion of an intermediate buffer is speculative, it is not entirely ad hoc. Homologous intermediate stages between low-level processing and response programming have been proposed in earlier models. Two examples are the informational-persistence stage proposed by Irwin and co-workers (Irwin & Brown, 1987; Irwin & Yeomans, 1986), and the schematic memory buffer proposed by Di Lollo and Dixon (1988; Dixon & Di Lollo, 1991).

A broad parallel could be drawn between the revised two-stage model and the competition model of the AB (Raymond et al., 1995). Stages 1 and 2 in the two-stage model could be likened to the template-matching and the report stages, respectively, in the competition model. And the holding buffer in the revised two-stage model could be likened to the VSTM store. Beyond this superficial level, however, the parallel breaks down because implicit in the revised two-stage model is a stimulus-substitution theory of backward masking that differs sharply from competition in VSTM. That is, while delayed in the holding buffer, the representation of a leading target is vulnerable to erasure by a trailing mask. When that happens, the representation of the mask replaces that of the target in the buffer and eventually gains access to Stage 2. The upshot is that the mask is substituted for the target as the object for eventual conscious registration. This object-substitution account differs sharply from an account based on competition in VSTM. According to the competition model, the trailing mask is added to the contents of VSTM. This increases the number of items from which the target must be selected, and the probability of a correct response is reduced correspondingly. Thus, although both models postulate a temporary store, the processing events that take place within the store are vastly different in the two schemes.

A view of backward masking akin to object substitution was held by Marcel (1983a) who surmised that "at the relevant SOAs the [target] and the mask are parsed into the same [temporal] segment and the relative recency of the mask is sufficient to grant it figural status for recovery" (Marcel, 1983a, p. 269). This view is well supported in the masking literature. There is ample evidence to show that when two targets are presented sequentially at an optimal SOA, it is the second one that is perceived to the detriment of the first (Bachmann & Allik, 1976; Schiller, 1966). This effect has been found to be more pronounced in unattended visual locations, suggesting that stimuli displayed outside the focus of attention are more likely to be delayed in the holding buffer, thus remaining vulnerable to substitution over a longer period (Enns & Di Lollo, 1997).

More important to the present argument, a tendency toward increased stimulus substitution has been obtained not only when attention is distributed over space (Enns & Di Lollo, 1997) but also when it is distributed over time, as in the AB deficit. This was revealed with remarkable clarity in two recent studies of the AB deficit (Chun, 1997; Martin, Isaak, & Shapiro, 1995). In the study by Martin et al., all items in the RSVP stream were alphabetical characters, with the two targets differing in size from the distractors. The principal issue under investigation was the nature of the errors made when the second target was identified incorrectly. It was found that the most common misidentifications of the second target arose from reporting the next item instead. In the study by Chun, the task was to report colored letter targets presented among black-letter distractors. In agreement with the findings of Martin et al., Chun reported that the proportion of reports of the item directly following the second target increased during the AB. Both outcomes are concordant with the tenets of the revised two-stage model. In each case, the trailing item in the RSVP stream erased and replaced the second target while the latter was delayed in the holding buffer because Stage 2 was busy processing the first target. One might add that according to the revised model, this process of substitution did not interfere with the pattern of activation that processing of the second target had produced in Stage 1. In accordance with the results of Shapiro et al. (1997), that activation could mediate priming of a related third target.

References

Bachmann, T., & Allik, J. (1976). Integration and interruption in the masking of form by form. *Perception, 5,* 79–97.

Breitmeyer, B. G. (1984). *Visual masking: An integrative approach.* New York: Oxford University Press.

Carr, T. H., & Dagenbach, D. (1990). Semantic priming and repetition priming from masked words: Evidence for a center-surround attentional mechanism in perceptual recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 341–450.

Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology, 40,* 343–367.

Chun, M. M. (1997). Temporal binding errors are redistributed by the attentional blink. *Perception & Psychophysics, 59,* 1191–1199.

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 109–127.

Dagenbach, D., Carr, T. H., & Wilhelmsen, A. (1989). Task-induced strategies and near-threshold priming: Conscious effects on unconscious perception. *Journal of Memory and Language, 28,* 412–443.

Di Lollo, V., & Dixon, P. (1988). Two forms of persistence in visual information processing. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 671–681.

Dixon, P., & Di Lollo, V. (1991). Effects of display luminance, stimulus meaningfulness, and probe duration on visible and schematic persistence. *Canadian Journal of Psychology, 45,* 54–74.

Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review, 87,* 272–300.

Duncan, J., Ward, R., & Shapiro, K. L. (1994). Direct measurement of attentional dwell time in human vision. *Nature, 369,* 313–315.

Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. *Psychological Science, 8,* 135–139.

Grandison, T. D., Ghirardelli, T. G., & Egeth, H. E. (1997). Beyond similarity: Masking of the target is sufficient to cause the attentional blink. *Perception & Psychophysics, 59,* 266–274.

Irwin, D. E., & Brown, J. S. (1987). Tests of a model of informational persistence. *Canadian Journal of Psychology, 41,* 317–338.

Irwin, D. E., & Yeomans, J. M. (1986). Sensory registration and informational persistence. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 343–360.

Kinchla, R. A. (1992). Attention. *Annual Review of Psychology, 43,* 711–742.

LaBerge, D. L. (1990). Attention. *Psychological Science, 1,* 156–162.

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subjects designs. *Psychonomic Bulletin & Review, 1,* 476–490.

Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature, 383,* 616–618.

Maki, W. S., Frigen, K., & Paulson, K. (1997). Associative priming by targets and distractors during rapid serial visual presentation: Does word meaning survive the attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1014–1034.

Marcel, A. J. (1983a). Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology, 15,* 238–300.

Marcel, A. J. (1983b). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15,* 197–237.

Martin, J., Isaak, M. I., & Shapiro, K. L. (1995, August). *Probe identification errors support an interference model of the attentional blink in rapid serial visual presentation.* Poster presented at the 103rd Annual Convention of the American Psychological Society, New York.

Meyer, D. E., Schveneveldt, R. W., & Ruddy, M. G. (1975). Loci of contextual effects on visual word recognition. In P. M. A. Rabbitt & S. Dornic (Eds.), *Attention and performance V* (pp. 98–118). New York: Academic Press.

Michaels, C. F., & Turvey, M. T. (1973). Hemiretinae and nonmonotonic masking functions with overlapping stimuli. *Bulletin of the Psychonomic Society, 2,* 163–164.

Moore, C. M., Egeth, H., Berglan, L. R., & Luck, S. J. (1996). Are attentional dwell times inconsistent with serial visual search? *Psychonomic Bulletin & Review, 3,* 360–365.

Purcell, D. G., & Stewart, A. L. (1970). U-shaped backward masking functions with nonmetacontrast paradigms. *Psychonomic Science, 21,* 362–363.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18,* 849–860.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1995). Similarity determines the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 653–662.

Scheerer, E. (1973). Integration, interruption and processing rate in visual backward masking. *Psychologische Forschung, 36,* 71–93.

Schiller, P. H. (1966). Forward and backward masking as a function of relative overlap and intensity of test and masking stimuli. *Perception & Psychophysics, 1,* 161–164.

Seiffert, A. E., & Di Lollo, V. (1997). Low-level masking in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1061–1073.

Shapiro, K. L., Driver, J., Ward, R., & Sorensen, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science, 8,* 95–100.

Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 357–371.

Spencer, T. J., & Shuntich, R. (1970). Evidence for an interruption theory of backward masking. *Journal of Experimental Psychology, 85,* 198–203.

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

Ward, R., Duncan, J., & Shapiro, K. (1996). The slow time-course of visual attention. *Cognitive Psychology, 30,* 79–109.

Weichselgartner, E., & Sperling, G. (1987). Dynamics of controlled visual attention. *Science, 238,* 778–780.

Received May 13, 1996
Revision received July 21, 1997
Accepted August 19, 1997 ■