The attentional blink is immune to masking-induced data limits

Elizabeth N. McLaughlin, David I. Shore, and Raymond M. Klein

Dalhousie University, Halifax, Canada

The attentional blink is the robust finding that processing a masked item (T1) hinders the subsequent identification of a backwards masked second item (T2), which follows soon after the first one. There has been some debate about the theoretically important relation between the difficulty of T1 processing and the ensuing blink. In Experiment 1 we manipulated the difficulty of T1 in such a way as to affect the quality of data without altering the amount of resources allocated to its identification. We found no relation between the accuracy of T1 identification and the blink. In Experiment 2, the same difficulty manipulation was applied to T2, and we observed an additive pattern with the blink. Together, this pattern of results indicates that a data-limited difficulty manipulation does not affect the blink, whether applied to T1 or T2. In Experiment 3 we used an individual differences methodology to show that performance in the traditional "stream"-like presentation (rapid serial visual presentation) was highly correlated with performance in our modified "target mask, target mask" paradigm, thus allowing for comparisons beyond the present methodology to much of the previous literature that has used the stream paradigm.

It is now well known that processing a target item (T1) hinders an observer's ability to detect or identify a subsequent item (T2) when the two items are both backwards masked and separated by less than a few hundred milliseconds. This robust effect has been demonstrated using a variety of procedures, including presenting words (e.g., Broadbent & Broadbent, 1987), letters (e.g., Raymond, Shapiro, & Arnell, 1992), or pictures (Boucart, Moroni, Fuentes, & Belin, 1998) in rapid serial visual presentation (RSVP) in the same location or in different locations (e.g., Duncan, Ward, & Shapiro, 1994). The relatively poor performance on T2 following T1 processing has been called cognitive dwell time by some (Duncan et al., 1994), and, more popularly, the attentional blink (AB; Raymond et al., 1992). In their initial studies, Raymond et al. (1992) demonstrated that

when participants were instructed to ignore the first target, there was no subsequent effect on T2 accuracy. By doing so, they confirmed their hypothesis that it is the act of having to process T1 for subsequent report that causes the reduction in performance on T2. As such, there is a common assumption that one or more components of T1 processing require a limited capacity resource, which is then unavailable for T2 processing. Theories that have been advanced to explain the attentional blink differ, however, with regard to the nature of this underlying information-processing limit.

Models of the blink

There are two broad classes of model that have been used to account for the basic attentional blink result—interference models (e.g., Raymond, Shapiro, & Arnell, 1995; Shapiro & Raymond, 1994; Shapiro, Raymond, & Arnell, 1994) and bottleneck models (e.g., Chun & Potter, 1995; Giesbrecht & Di Lollo, 1998; Seiffert & Di Lollo, 1997). According to the interference models, when a series of items is presented in RSVP, each item is processed to some degree, but only a few items are admitted into the higher level of visual short term memory (VSTM). Items may be admitted into VSTM if they match a preset template of the target or probe, and if they are temporally contiguous to the target or probe. Thus, in the traditional RSVP stream, four items should be admitted to VSTM: the first target (T1), the item immediately following the first target (T1+1), the second target (T2), and the item immediately following the second target (T2+1). Each item is assigned a weighting depending on factors such as the degree to which it matches a preset template of T1 or T2, its temporal contiguity with items that match a template, and its order of entry (with earlier items being given higher weightings). Target 2 errors occur when the wrong item is retrieved from VSTM, and this is presumed to be due to interference or competition from other highly weighted items in VSTM at the time of retrieval. (Because of the competition in VSTM, the interference model has also been referred to as the "competition hypotheses", cf., Seiffert & Di Lollo, 1997.) The characteristic AB occurs because T2 errors are most likely to occur at short (200–500 ms) T1–T2 stimulus onset asynchronies (SOAs), when VSTM is still occupied by T1 and T1+1. In contrast, at longer T1–T2 SOAs, T1 and T1+1 have a chance to decay, to be flushed from VSTM, or to move on to a reporting stage, and thus there is less resulting competition within VSTM for the second target.

The second class of model presumes that the AB occurs because of a processing bottleneck, such that resources are occupied with T1 when T2 occurs. Duncan et al. (1994) used the term "attentional dwell time" to describe the period during which attentional resources are occupied with the first target. Chun and Potter's (1995) two-stage model of the blink encompasses a similar notion. They hypothesized that representations of each item in the RSVP stream are initially formed in a high-capacity short-term storage system (Stage 1), which is highly subject to visual interference. When a target item must be selected from among a stream of distractor items for subsequent report, it must be consolidated in a second, capacity-limited stage. In Chun and Potter's (1995) model, the AB is due to a bottleneck at this second stage, which can handle only one item at a time. In an RSVP stream, if Stage 2 resources are occupied with T1 when T2 is presented, T2 is likely to be missed because it decays or is overwritten while awaiting access to Stage 2.

Comparing the models

Although slightly different versions of these two models can be found in the literature (e.g., the Central Limitation Theory by Jolicoeur, 1998, the revised two-stage model by Giesbrecht & Di Lollo, 1998, and an extension of this model using the concept of object substitution by Brehaut, Enns, & Di Lollo, 1999), at this point it is reasonable to regard the VSTM interference and bottleneck approaches as exemplars of two broad classes of model, and then to see if we can use the literature or new empirical evidence to favour one class over the other (Broadbent, 1958, p. 307). There are several similarities and differences between these two classes of theory that should be highlighted. According to both the interference and the bottleneck models, all items in the stream are processed to some degree, and only some items are selected for further processing. In both theories, the blink-causing limitation occurs during this latter stage of processing. The main difference between the theories lies in the mechanism whereby T2 identification is impeded at short T1–T2 SOA's. According to the interference model, it is competition among items in VSTM that decreases the probability of correct retrieval of the second target. In contrast, according to the bottleneck model, it is the fact that a later stage of processing can only process one item at a time, that causes the second target to be lost while it waits for access to this stage.

One method to evaluate, validate, and further develop these models is to explore the factors that modulate the magnitude of the effect. Thus far, both theories have been able to account for much of the data generated by following this strategy. For example, both offer a compelling explanation for the consistent observation that when T1 is followed by a blank screen (i.e., when there is nothing in the T1+1 position), the AB is markedly reduced or absent (Chun & Potter, 1995; Grandison, Ghirardelli, & Egeth, 1997; Raymond et al., 1992; Seiffert & Di Lollo, 1997). According to the interference model, when there is no item immediately following T1, fewer items are admitted into VSTM. (Recall that the T1+1 item is usually admitted to VSTM because of its temporal relation to T1.) As a result, there is less competition and thus less interference in VSTM when T2 arrives (Shapiro & Raymond, 1994). The bottleneck model also accommodates this finding. When there is no distractor item immediately following T1, T1 is easily discriminable from the distract stream and is processed rapidly in Stage 2 (Chun & Potter, 1995). This leaves Stage 2 available by the time T2 must be processed. Thus, the observation that a blank following T1 leads to a reduction in the AB is unable to distinguish between the two classes of model.

In contrast, the two classes of model lead to explicitly different predictions with regard to the relation between the difficulty of the first target and the size of the ensuing blink. Whereas the interference model predicts no relation, the bottleneck model predicts a strong negative relation between T1 accuracy and the magnitude of the blink. For the interference model, the main modulating factor in the blink is the weighting assigned to the T1+1 item. If, for example, T1+1 is similar to other items in the stream, it gets into VSTM with a higher weighting and interferes more with the subsequent probe, relative to a dissimilar T1+1 item. According to this model, the difficulty of T1 processing should have no direct impact on the magnitude of the blink (Raymond et al., 1995). In contrast, according to a

widely held version[1] of the bottleneck view, whatever makes T1 identification more difficult should lead to a bigger blink (Chun & Potter, 1995; Giesbrecht & Di Lollo, 1998; Grandison, Ghirardelli, & Egeth, 1997; Seiffert & Di Lollo, 1997; see also Potter, Chun, Banks, & Muckenhoupt, 1998). The more difficult the T1 task, the more Stage 2 processing is delayed. As a result, when T1 is difficult, T2 is impaired at longer T1–T2 SOAs, leading to a larger AB.

It should be noted that Shapiro, Arnell, and Raymond (1997) have put forward a unified model of the blink, incorporating the elements of several models, including those described earlier. Although Shapiro et al. (1997) do not make an explicit prediction concerning the effect of T1 difficulty upon the magnitude of the blink, their second tenet asserts that "As less attention is available for T2 by virtue of T1's demands, T2 cannot be consolidated. . . . " (p 293). If by this they meant that the available attention would be related to T1's demands, then this unified theory, like bottleneck theories, predicts that blink magnitude should vary with T1 difficulty. Empirically, there has been support for and against a relation between T1 difficulty and the size of the AB.

Blink magnitude as a function of T1 difficulty

Shapiro et al. (1994) were the first to systematically investigate the relation between T1 difficulty and the blink. They found a slight improvement in T2 performance when T1 set size was reduced from 25 to 3 letters. However, they found no differences in T2 performance when the T1 task was made less difficult by changing it to a detection (vs. identification) task, even when T1 had the same identity on each trial. Also, they compared the relatively easy T1 task of detecting a white letter in a stream of black letters with the harder task of detecting a specific black letter in a stream of black letters. Despite lower performance in the latter task, the blink magnitude was the same in both tasks. Overall, when they combined the results from their four experiments, Shapiro et al. found no significant relation between T1 target detectability (d′) and blink magnitude,[2] leading them to conclude that the blink occurred in "all or none fashion" (p 370).

Later, Raymond et al. (1995) hypothesized that one potential determinant of the size of the AB could be the similarity between the T1+1 item and other items in the stream; a T1+1 item that is highly similar to T1 or T2 should lead to more competition in VSTM. In support of this hypothesis, a reduction in the magnitude of the AB was observed by making the T1+1 item featurally or spatially dissimilar to other items in the stream. The T1 task was to identify a white letter in a stream of black letters, and the T2 task was to detect the presence or absence of a black X later in the stream. Comparing across experiments, the size of the blink was attenuated when the T1+1 item was a dot pattern (featural dissimilarity) or when it was a displaced letter (spatial dissimilarity), relative to

---

[1] It is worth noting that not all versions of the bottleneck model predict such a relation. Ward, Duncan, and Shapiro's (1996) version of the bottleneck view predicts, at least in the skeletal (target–mask, target–mask) paradigm, that varying the amount of information to be reported from a single T1 object will not affect the blink, whereas the blink can be increased by increasing the number of objects in the T1 frame.

[2] More recently, however, both Grandison et al. (1997) and Seiffert and Di Lollo (1997) have suggested that the lack of a significant correlation in this case was due to low statistical power; see the meta-analysis (Seiffert & Di Lollo, 1997) discussed later.

conditions for which the T1+1 item was either a specific black letter or one of a set of black letters. Important for the present discussion is the relation between T1 accuracy and the magnitude of the AB—accuracy on T1 was higher in the dissimilar (T1+1 = dots or displaced letters) conditions. However, the correlation between T1 accuracy and blink magnitude was non-significant in each experiment. This led Raymond et al. to conclude that "target task difficulty and AB magnitude are unrelated" (p 658). Consistent with this conclusion, Ward, Duncan, and Shapiro (1997) failed to affect the blink when they varied the difficulty of the size discrimination required for T1 (hard = medium vs. small; easy = large vs. small). Even though their difficulty manipulation yielded a 14% difference in T1 accuracy, the resulting blink was identical in the two conditions. Similarly, Ward, Duncan, and Shapiro (1996, Experiment 2) asked subjects to discriminate the shape (identify) or the size of T1. Again, despite a significant difference in T1 accuracy levels (higher accuracy for the size discrimination), T2 performance and the pattern of T2 performance across SOA (the blink) remained unchanged. The blink was also unaffected when participants were required to report both stimulus attributes (size and identity) of T1. In Ward et al. (1997), participants were instructed to prioritize the first task, and it is possible that this caused homogenous allocation of resources to T1 and hence there was no effect on blink magnitude. However, this speculation cannot account for the finding in Ward et al. (1996).

Others, however, have found support for an inverse relation between T1 accuracy and the size of the blink (or correspondingly, a positive relation between T1 difficulty and the blink). Chun and Potter (1995) demonstrated that when T1 was made harder to discriminate from the distractor stream, T1 accuracy was reduced, and the subsequent AB was larger. For example, in their Experiment 5, T2 accuracy was significantly lower when T1 (a letter) was embedded among a stream of visually similar distractors (numbers), than when the target letter was embedded among a stream of visually distinct distractors (symbols). Furthermore, the blink was significantly attenuated (relative to other experiments) in the latter condition. Whether this finding is due to visual or conceptual similarity between T1 and its neighbours in the stream cannot be determined from this study. Supporting the possibility that conceptual similarity is an important factor, Boucart et al. (1998) showed that increasing the semantic relatedness between a T1 picture and those in the stream resulted in a larger blink.

The "T1+1 blank" effect mentioned previously has also been used by some as evidence to support the role of T1 difficulty in the blink. Seiffert and Di Lollo (1997) argued that the reason no blink occurs following a blank screen at T1+1 is that T1 is not effectively masked. They supported this claim empirically by showing an AB when they used a simultaneous mask on T1 instead of using a T1+1 item. This finding (which has been replicated by Brehaut et al., 1999, using a location-switching paradigm) demonstrates that the degradation of T1 through masking (and not the inclusion of a T1+1 item, per se) is necessary to cause the subsequent blink. Furthermore, comparing different types of mask (superposition, metacontrast, backward), Seiffert and Di Lollo (1997) found a significant relation between the amount of degradation due to masking (as reflected by low T1 accuracy) and the size of the AB.

Grandison et al. (1997) also maintain that the magnitude of the AB is negatively related to T1 accuracy, when the latter is varied by the effectiveness of T1 masking. They used

different types of mask (pattern, luminance, and metacontrast) in the T1+1 position and found an AB only when T1 was effectively masked. Across these studies, Grandison et al. found a significant negative correlation ($r = -.48$, $p < .01$) between T1 accuracy and the magnitude of the AB. Similarly, Seiffert and Di Lollo (1997) combined the results of five published studies (a total of 27 experiments, not including those of Grandison et al.) and found a negative correlation ($r = -.73$, $p < .001$) between percentage correct on T1 and AB magnitude.[3]

Finally, using a modification of Duncan et al.'s (1994) dwell time paradigm (with location and set switching, but without task switching), Moore, Egeth, Berglan, and Luck (1996) showed that when the first target was made easy by either delaying (Experiment 1) or not presenting (Experiment 2) its mask, T2 performance reached asymptotic levels sooner than when T1 was immediately masked. Moore et al. (1996) observed this main effect of T1 difficulty on the blink whether T1 difficulty was blocked (Experiment 2) or mixed (Experiment 1). It is important to note that in their Experiment 2, a significant blink was observed even when T1 was not masked, challenging the assertion that T1 must be masked in order for a blink to occur (e.g., Giesbrecht & Di Lollo, 1998, p. 1454; Grandison et al., 1997, p. 271). We believe that these conflicting results may be related to procedural differences such as the presence or absence of location switching and set switching. As is discussed later, we believe (as others have suggested, e.g., Ward et al. 1997, with respect to location shifting and Potter et al., 1998, with respect to task switching), that switching in and of itself may cause a delay in the ability to process the second target.

Scope of the present study

Clearly then, there is evidence both for and against a relation between T1 difficulty and the size of the AB. As the two competing theories (interference and bottleneck) make different predictions regarding this relation (the former predicting no such relation, the latter predicting a negative correlation), this topic is still very much worthy of investigation. We believe that the reason for the mixed findings to date may be found in the methods used to implement the difficulty manipulation and those used to study the blink. First, unlike previous studies, we restricted our difficulty manipulation to one of perceptual quality, rather than of resource allocation, and second, we designed our experimental paradigm with the intention of avoiding unnecessary and potentially confounding features that have been used in many prototypical blink studies.

A difficulty manipulation based only on perceptual quality

In much of the previous research on target difficulty and the blink, the difficulty manipulation probably had its effect by requiring more or fewer resources. In light of recent theories suggesting that utilization of central resources might be an important factor in the AB (Jolicoeur's, 1998, central limitation theory is discussed further in the

---

[3] The data from Shapiro et al. (1994) and Raymond et al. (1995), which were mentioned earlier and alone showed no relation, were included in this meta-analysis.

General Discussion), we chose to manipulate difficulty by varying, in a continuous and quantifiable fashion, the data quality of the targets. Specifically, we altered the quality of the evidence from the first target and, as such, altered only its encoding difficulty. We made T1 detection easy or hard by varying reciprocally the duration of the target (T1) and mask (M1) (easy = long T1 and short M1; hard = short T1 and long M1).

Others have attempted to look at a data-limited approach to this issue. For example, as reviewed earlier, many have used the masking of T1 as their data manipulation (e.g., Grandison et al., 1997; Seiffert & Di Lollo, 1997). However, in these studies, the different degrees of T1 difficulty were manipulated between blocks of trials or between subjects. When aware in advance that T1 encoding might be difficult, the subject may allocate more resources to T1 encoding, leaving less for T2. In order not to confuse resource allocation with target difficulty, we randomly intermixed target difficulty within blocks so that subjects could not anticipate T1 difficulty in advance. Note that this approach assumes either that participants will not be able to differentially allocate resources "on line", or that if they could, it would take too long to lead to improved encoding.

An alternative method used to manipulate the difficulty of T1 has been to manipulate the similarity between the target and the distracting stream (cf. Boucart et al., 1998; Chun & Potter, 1995). For example, Chun and Potter (1995) manipulated T1 difficulty by changing the category of the distractor stream. The modulated blink magnitude in this case could be explained by the interference model as being due to the difference between T1 and the T1+1 item (with more highly similar items being assigned higher weightings) and not to the difference in T1 difficulty (e.g., Raymond et al., 1995; Shapiro et al., 1994). To avoid this potential ambiguity, our difficulty manipulation did not involve a featural or spatial modification to the T1+1 item. Target 1 was always a letter, and the T1+1 item was always a pattern mask.[4]

Other studies that have used masking to implement their difficulty manipulation (Grandison et al., 1997; Seiffert & Di Lollo, 1997) compared qualitatively different types of mask (integration, interruption, luminance, metacontrast). Different types of mask may interfere with different resources or at different levels of processing and thus not affect the blink in the same way (see Brehaut et al., 1999, Giesbrecht & Di Lollo, 1998, for an examination of the role of various types of mask in the AB). Here, we used the same kind of mask (interruption pattern mask), but varied the signal strength by co-varying the target and mask durations.

A design with no switching

Many of the experiments reviewed earlier involved a location switch, such that T1 was in a different location from that of T2 (e.g., Ward et al., 1996); a set switch, such that the category (e.g., colour) for T1 was different from the category for T2 (e.g., Grandison et al.,

---

[4] Another factor that may influence the weighting of the T1+1 item in VSTM is its temporal contiguity to the target (Shapiro et al., 1994). In our design, the T1–M1 SOA varies with the difficulty condition, but the T1–M1 interstimulus interval (ISI) remains fixed. As temporal contiguity is not operationalized in the interference model to the degree that it distinguishes between these two values, for now we will assume that our paradigm does not confound T1+1 weighting (via temporal contiguity) with difficulty.
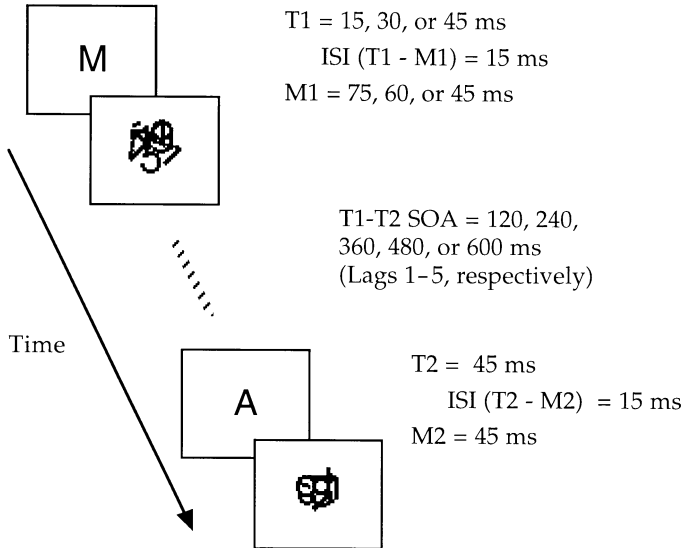
1997); or a task switch, such that the T1 task (usually identification) was different from the T2 task (usually detection; e.g., Seiffert & Di Lollo, 1997). The precise nature of the relation between the AB and switching is still not completely understood. However, Potter et al. (1998) have shown empirically that the T2 deficits observed in traditional AB tasks (with task and set switching) are probably due to both task switching and a true cognitive dwell time. To be sure that we would be studying a "blink" that is elicited because of the processing of T1, and not because of a switch between T1's and T2's locations, stimulus sets, or tasks, we decided to eliminate switching altogether. This is in line with John Duncan's pronouncement of a methodological "new world standard" for studies of the so-called attentional blink (Duncan, 1998; see also, Arnell & Duncan, 1999).

All of the previous studies exploring the sensitivity of the AB to T1 difficulty included at least one potentially complicating factor (e.g., confounding difficulty with similarity of T1+1 to T1; blocking of T1 difficulty thus permitting advanced adjustments of resource allocation; location, task, or set switching; instructions biasing resource allocation). It is unclear at this point how or whether one or more of these factors could act as modulators in the relation between T1 difficulty and the blink. Indeed, in our review of the literature, we were unable to find a consistent pattern between a specific feature (or a specific combination of features) and whether or not a relation was observed between T1 difficulty and the blink. However, as each of these potentially complicating factors could influence the use of central resources, we designed our paradigm with the exclusion of each of these factors.

To do so we combined features of Chun and Potter's (1995) RSVP and Ward et al.'s (1997) modification of Duncan et al.'s (1994) skeletal RSVP paradigms. We refer to our task as a TM–TM task (with the name denoting the sequence of items in a stream: target mask–target mask). As in Chun and Potter's (1995) paradigm, there was no task or set switch because the T1 task and the T2 task were the same: Identify a black letter. As in Duncan et al.'s (1994) skeletal paradigm, we removed the distractor stream, and each target was followed by a pattern mask. Whereas the sequence of events was similar to Duncan et al.'s (1994) skeletal paradigm, like Ward et al. (1997) we removed location switching; and, moreover, we also removed task and set switching. Unlike Ward et al. (1997), we did not instruct participants to make T1 the priority. In order to alter the quality of perceptual evidence for T1 (i.e., T1 difficulty) the duration of the first target and its mask were reciprocally co-varied. That is, the overall duration of the target–mask combination was held constant (105 ms) but the target duration was made shorter or longer as the mask duration was made longer or shorter (see Figure 1). Difficulty level was randomly mixed within a block. As in all previous paradigms used to elicit the AB, T1 and T2 were separated by variable SOAs, and we were interested in T2 performance across different SOAs.

In Experiment 1 we tested the effect of T1 difficulty (i.e., reduced perceptual quality) on the AB. In Experiment 2, we further examined the relative independence of perceptual quality and the AB by manipulating the difficulty of T2. In Experiment 3, we strengthened the link between our new paradigm and previous research by administering our paradigm and a more traditional paradigm to the same group of participants and conducting a within-subject correlation on the blink magnitudes collected from them. Findings will be discussed in terms of their implications for the two classes of theory.

A. Sequence of Events

T1 = 15, 30, or 45 ms
ISI (T1 - M1) = 15 ms
M1 = 75, 60, or 45 ms

T1-T2 SOA = 120, 240,
360, 480, or 600 ms
(Lags 1–5, respectively)

Time

T2 = 45 ms
ISI (T2 - M2) = 15 ms
M2 = 45 ms

B. Difficulty Manipulation

Difficulty   T-ISI-M = 105 ms

Hard

Medium

Easy

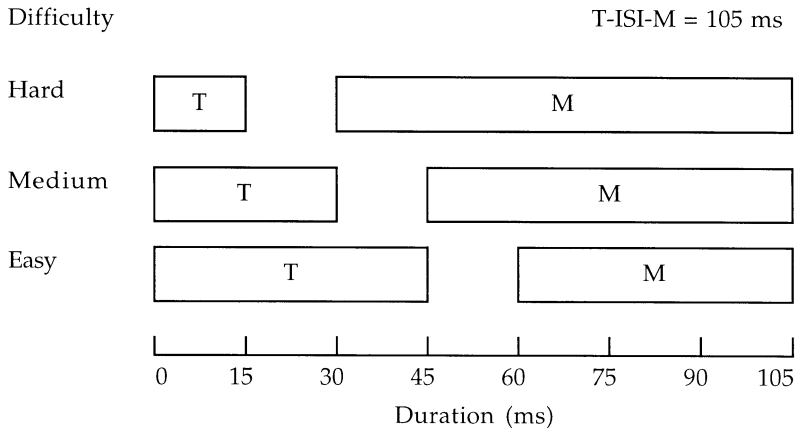0   15   30   45   60   75   90   105

Duration (ms)

Figure 1.   A. Sequence of events and presentation rates in the target mask–target mask task. T1, T2, M1, and M2 refer to the first and second targets and first and second masks, respectively. ISI refers to the interstimulus interval between targets and their respective masks. SOA refers to the stimulus onset asynchrony between the first and second targets. Lags 1–5 correspond to the five different SOAs.
B. The data-limited difficulty manipulation, which was implemented by co-varying the duration of the target (T) and mask (M). In all three difficulty conditions, the total duration of the target, ISI, and mask was 105 ms. In Experiment 1, T1 and M1 durations were co-varied, whereas T2 and M2 were fixed at the medium durations. In Experiment 2, T2 and M2 were co-varied, and T1 and M1 were fixed.

EXPERIMENT 1

Method

Participants

A total of 17 students (14 female, 3 male, ranging in age from 20 to 29 years) in a cognitive psychology class at Dalhousie University participated in this experiment. Participation was required as part of the laboratory component of the class, but students had the option of whether or not they gave their consent to have their data used as part of this study. All participants were naive as to the nature of the paradigm and to our hypotheses. All had normal or corrected-to-normal vision.

Stimuli and apparatus

A trial consisted of two upper-case letters, each followed by a pattern mask. Letters were chosen randomly from the alphabet, with the exception of I and O. They were presented in black on a white screen, in Helvetica font (18 point). The largest letter (W) measured approximately 0.62 by 0.65 degrees when viewed at a distance of approximately 55 cm. Pattern masks were composed of jumbled pieces of digits (see Figure 1 for examples). We used two different pattern masks, measuring 1.02 by 1.06 degrees and 0.83 by 0.85 degrees, each rotated in four different orientations. The fixation stimulus was a small (0.40 by 0.65 degrees) black cross in the centre of the screen.

Procedure

Participants were tested in a group setting, in a room with 25 Power PC Macintosh computers. Lighting in the room was slightly dimmed. Viewing distance was not fixed; participants were instructed to sit at a comfortable distance from the computer screen.

Each trial began with the fixation cross. Participants initiated each trial by pressing the space bar, and the four-item sequence (T1, M1, T2, M2) began after a variable delay (100 ms, 200 ms, or 300 ms). Following the four-item sequence, participants were required to indicate which two letters they had seen by typing their responses on the computer keyboard. They were given no time limit to respond, were asked to name the items in the order that they saw them (only responses given in the proper order were recorded as correct), and were encouraged to guess when uncertain. Responses were recorded by the computer.

In order to manipulate the difficulty of T1, we systematically varied the duration of T1 and M1. In order not to affect the remaining RSVP lag durations, the total duration for T1, M1, and a 15-ms interstimulus interval (ISI) between them was always maintained at 105 ms. (See Figure 1 for a pictorial representation of the timing in the different conditions, as well as the sequence of the four events.) There were three levels of difficulty: easy, medium, and hard. In the easy condition, T1 was presented for 45 ms and M1 was presented for 45 ms. The T1 duration in the medium condition was shorter (30 ms), and the M1 duration was longer (60 ms) than in the easy condition. Similarly, T1 in the hard condition was shorter (15 ms), and M1 was longer (75 ms) than those in the other two conditions. The T1–M1 ISI for each of the difficulty conditions was 15 ms. Therefore, the T1–M1 SOA was 60 ms, 45 ms, and 30 ms in the easy, medium, and hard conditions, respectively. The second variable of interest was the SOA between T1 and T2. It varied between 120 ms and 600 ms, in steps of 120 ms (these are referred to as Lags 1–5: Lag 1 = 120 ms SOA, Lag 2 = 240 ms SOA, etc.). Target 2 durations paralleled those in the medium condition of T1. That is, T2 was presented for 30 ms, M2 was presented for 60 ms, and the T2–M2 ISI was 15 ms. Each participant completed a block of 30

practice trials, followed by a block of 300 test trials. The test trials were composed of 20 in each of the 15 levels of Difficulty × Lag.

Results and discussion

Figure 2 represents T1 accuracy and conditional T2 accuracy[5] by T1 difficulty and T1–T2 SOA (lag). The pattern of results is quite clear: Although our masking manipulation had a dramatic impact on T1 accuracy, it had no effect on the AB. A 3 (difficulty) × 5 (lag) repeated measures analysis of variance (ANOVA) was performed on T1 accuracy. There was a significant main effect of difficulty, $F(2, 32) = 53.04$, $p < .0001$. Accuracy was a monotonic function of the quality of the evidence. Collapsed across lags, the mean percentage correct was 89.82 ($SD = 12.21$), 79.52 ($SD = 18.09$), and 59.71 ($SD = 21.34$), in the easy, medium, and hard conditions, respectively. There was also a significant main effect of lag on T1 performance, $F(4, 64) = 2.84$, $p < .05$. Despite this main effect, there was no consistent linear trend associated with lag. The Difficulty × Lag interaction was not significant, $F(8, 128) = 0.724$, $p = $ ns.

A 3 (difficulty) × 5 (lag) repeated measures ANOVA was also performed on T2 accuracy. The main effect of lag on T2 was significant, $F(4, 64) = 66.13$, $p < .0001$, reflecting the characteristic blink pattern. That is, performance is poor at the shortest lags and improves across lag. It is important to note that the U-shaped function typically observed (e.g., Raymond et al., 1992) was not seen here. This is likely because there was always a mask following the first target, whereas in the typical stream task, the T1+1 item could be the T2 item, in which case "Lag 1 sparing" is observed. Lag 1 sparing has been attributed to the slow closing of an attentional gate, which permits items immediately following T1 (within a window lasting approximately 200–250 ms) to be processed (see Visser, Bischof, & Di Lollo, 1999, for a complete discussion of the phenomenon). If temporal contiguity of T2 to T1 were the only factor responsible for Lag 1 sparing then, as our shortest T1–T2 SOA was 120 ms, we should probably have seen Lag 1 sparing. However, in every case that Lag 1 sparing has been observed to date, the second target has been the item immediately following T1, and has been the only item appearing within the 200–250 ms window. In our case, we have an intervening item, T1's mask, which occurs after T1 and before T2. To our knowledge, Lag 1 sparing has not been observed under these conditions.

Our main question in Experiment 1 was whether the difficulty of the first target would affect performance on the second target. Neither the main effect of difficulty, $F(2, 32) = 0.04$, $p = $ ns, nor the Difficulty × Lag interaction, $F(8, 128) = 1.77$, $p = .09$, were significant, indicating that the difficulty of T1 had no effect on the subsequent blink. Overall, this pattern of performance indicates that there was no change in the magnitude of the AB (in terms of either absolute performance or performance by lag) corresponding to T1 difficulty. To put this finding in perspective it is instructive to consider Seiffert and Di Lollo's (1997) meta-analysis of the effect of T1 difficulty on blink magnitude. In their

---

[5] As noted by others (e.g., Giesbrecht & Di Lollo, 1998), we can be certain that the interference on T2 identification is due to the processing of T1 only on trials for which T1 was correct. Therefore, all measures of T2 performance are conditional values—that is, T2 accuracy when T1 was correct.
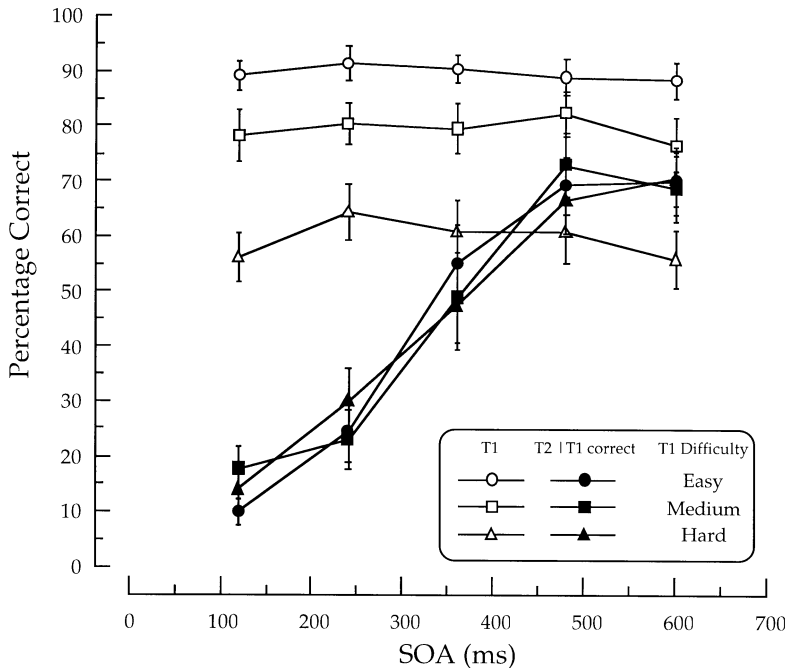
Figure 2. Results from Experiment 1. Target 1 (T1) accuracy and conditional Target 2 (T2|T1) accuracy at three levels of T1 difficulty (easy, medium, and hard) and five levels of lag (corresponding to T1–T2 SOA). Target 2 difficulty was held constant (medium difficulty). Error bars represent the between-observer standard error of the mean.

review (see their Figure 6) it was shown that the performance decrement on T2 during the 180–540 ms time interval following T1 increased five-fold as T1 performance decreased from near 100% to about 65%. In our experiment T1 accuracy varied from 90% to 60% with virtually zero effect on the blink timing and magnitude.

EXPERIMENT 2

Since the initial studies of the AB, the focus of attention has been on the role of the T1 and T1+1 items in eliciting the blink. The second target was viewed simply as a way to measure the resources that were available following T1, T1+1, and the variable SOA. Recently, however, researchers have turned their attention to the role of T2 processing in the occurrence and magnitude of the blink. Giesbrecht and Di Lollo (1998) were the first to demonstrate that the variables surrounding the T2 item were critical to the AB. Specifically, they showed that in order for the AB to occur, T2 must be masked with an interruption mask. If T2 is not masked, or if it is masked with a simultaneous/ integration mask, there will be no AB. Giesbrecht and Di Lollo (1998) examined the difference between integration (simultaneous) and delayed (interruption) masking of T2. Their delayed masks were administered at three different SOAs: 50 ms, 100 ms, and 200 ms. For their task, participants were required to name the two target letters in a stream of

digit distractors. As such, there was a distractor stream, but there was no task switching. The different conditions (reflecting different SOAs between T2 and T2+1) were inter-mixed within a block. They found an AB at each SOA. They also found an overall main effect of masking difficulty, such that the overall accuracy level was highest in the 200 ms SOA condition, lowest in the 50 ms condition, and intermediate in the 100 ms condition. Brehaut et al. (1999) also showed that an AB does not occur following integration masking of T2. The study of T2 characteristics promises to elucidate as much about the AB as did earlier studies on T1 and T1+1.

In Experiment 2, we performed the same difficulty manipulation on T2 as we did on T1 in Experiment 1. We randomly varied the perceptual quality of T2. As in Experiment 1, we used our TM–TM paradigm, which does not entail a distractor stream and is free of task, set, and location switching.

Method

Unless otherwise noted, methods were identical to Experiment 1.

Participants

A total of 15 participants from Experiment 1 (13 female, 2 male) took part in this experiment.

Procedure

Target 1 durations remained constant throughout this experiment at T1 = 30 ms and M1 = 60 ms. The durations of T2 and M2 were manipulated to yield three different conditions of T2 masking strength (easy: T2 = 45 ms, M2 = 45 ms; medium: T2 = 30 ms, M2 = 60 ms; hard: T2 = 15 ms, M2 = 75 ms). The T2–M2 interstimulus interval remained constant at 15 ms, producing T2–M2 SOAs of 60 ms (easy), 45 ms (medium), and 30 ms (hard). As in Experiment 1, the T1–T2 SOA varied in steps of 120 ms from 120 ms to 600 ms. Participants completed 30 practice trials and 300 test trials (20 in each Difficulty × Lag condition).

Results and discussion

Figure 3 shows that T1 was unaffected by either the T2 masking manipulation or lag. This was confirmed by a non-significant 3 (difficulty) × 5 (lag) repeated measures ANOVA on T1 accuracy: main effect (ME) of difficulty, $F(2, 28) = 0.19$, $p = $ ns; ME of lag, $F(4, 56) = 0.34$, $p = $ ns; Difficulty × Lag interaction, $F(8, 112) = 1.77$, $p = .09$. In contrast, T2 accuracy was influenced by both the difficulty manipulation and lag. A 3 (difficulty) × 5 (lag) repeated measures ANOVA revealed a main effect of difficulty, $F(2, 28) = 39.21$, $p < .0001$, with easy masks yielding higher accuracies than medium masks, which, in turn, yielded higher accuracies than hard masks. Collapsed across lags, T2 accuracy was 63.48 ($SD = 30.34$), 55.36 ($SD = 33.14$), and 41.31 ($SD = 28.65$) in the easy, medium, and hard T2 conditions. There was also a main effect of lag, $F(4, 56) = 83.16$, $p < .0001$, reflecting an attentional blink; this was the same pattern of T2 performance across SOA as that observed in Experiment 1. The Difficulty × Lag interaction
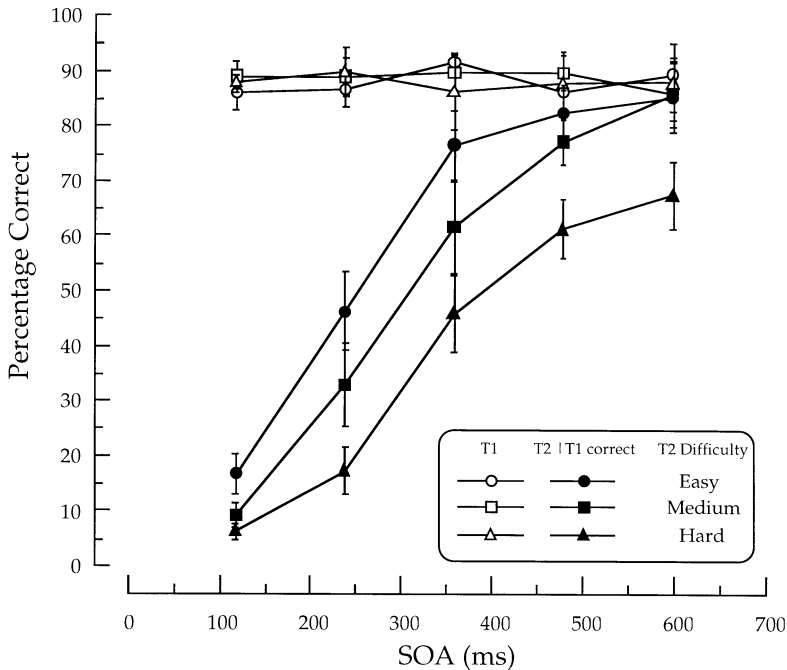
Figure 3. Results from Experiment 2. Target 1 (T1) accuracy and conditional Target 2 (T2|T1) accuracy at three levels of T2 difficulty (easy, medium, and hard) and five levels of lag (corresponding to T1–T2 SOA). Target 1 difficulty was held constant (medium difficulty). Error bars represent the between-observer standard error of the mean.

was also significant, $F(8, 112) = 2.97$, $p < .01$. An inspection of Figure 3 suggests that this interaction is probably due to floor effects at short lags (with 24 different target stimuli, chance on this task is 4%) and asymptotic effects at longer lags. To confirm this interpretation of the data, we conducted a repeated measures ANOVA on T2 performance excluding the shortest and longest lags (Lags 1 and 5). Similar to the results obtained with all five lags, the analysis on the intermediate lags revealed main effects of both difficulty, $F(2, 28) = 36.07$, $p < .0001$, and lag, $F(2, 28) = 72.76$, $p < .0001$. However, in this new analysis, the Difficulty × Lag interaction was no longer significant, $F(4, 56) = 0.75$, $p = $ ns. At the more intermediate lags, the slopes of the three lines were indistinguishable, suggesting that despite the significant Difficulty × Lag interaction (when all five lags are considered), the attentional blink was actually unaffected by the T2 difficulty manipulation.

This result replicates that of Giesbrecht and Di Lollo (1998), who also showed a change in overall T2 accuracy, but not in AB, when they altered the T2–M2 SOA. As our T2–M2 SOAs (30 ms, 45 ms, 60 ms) were much shorter than those used by Giesbrecht and Di Lollo (1998; 50 ms, 100 ms, and 200 ms), it is not surprising that overall we observed a larger blink than that in their experiment.

It is also important to reinforce the fact that we obtained a blink even when the T2–M2 SOA was only 30 ms. Both Giesbrecht and Di Lollo (1998) and Brehaut et el. (1999) make a categorical distinction between the effects of integration masking at T2 (where

T2–M2 SOA = 0 ms and no blink is observed) and interruption masking at T2 (where T2–M2 SOA > 0 ms and a blink is observed). Until now, the shortest T2–M2 SOA to be tested was 50 ms (as reported by Giesbrecht & Di Lollo). The present study adds support to the notion that even very small (30 ms) T2–M2 SOAs are sufficient to produce an AB.[6]

EXPERIMENT 3

Insofar as our paradigm requires participants to process and report on two targets that occur close in time, our TM–TM task should be tapping into the same attentional limitation that has been measured in other AB tasks. In fact, in Experiments 1 and 2, we replicated the characteristic reduction in accuracy of identifying T2 at T1–T2 SOAs below 500 ms, with accuracy increasing with SOA between 200 and 500 ms. In a similar situation, other researchers (e.g., Ward et al., 1997) have simply made the assumption that they were measuring the same underlying process with their revised tasks. To increase our confidence in this assumption, we sought to quantify the equivalence of the two paradigms. We administered a stream paradigm and our simplified task to the same group of participants. Our stream paradigm was similar to Chun and Potter's (1995) in that it required that participants identify two letters embedded in a stream of digits. We predicted that if our task measures the same underlying mechanism as in the stream task, and if the blink is a stable construct that is consistent within an individual (a second major assumption that has been made by researchers in the past), then performance on both tasks (i.e., strength of blink) should be correlated within participants.

Method

Unless otherwise noted, the methods used in Experiment 3 were the same as those in Experiments 1 and 2.

Participants

A total of 18 students (15 female, 3 male) from Experiment 1 participated in at least one part of this two-part (RSVP stream and TM–TM) experiment. Of the students, 16 (14 female, 2 male) completed both tasks.

Order of administration

The order of the two paradigms was fixed, such that all participated in the stream paradigm in one session, followed by the TM–TM 4 weeks later.[7]

---

[6] This finding may say something about the distinction between integration and interruption masking. The fact that we do obtain a blink under our parameters suggests that our mask has at least some component of interruption masking, because according to Brehaut et al. (1999) and Giesbrecht and Di Lollo (1998), an AB occurs only when T2 is masked with an interruption mask. And, operationally of course, our mask occurs *after* the target, which is a hallmark feature of an interruption mask. However, as some of our T2–M2 SOAs were very short, there is likely to be some component of integration masking as well.

[7] Because we were interested in correlating performance on both tasks, we chose not to counterbalance order of administration in order to minimize between-subjects variation due to practice effects.

RSVP stream task

Stimuli and apparatus.   Each trial consisted of an RSVP stream consisting of two upper-case letters embedded in a stream of digits.[8] Letters were chosen randomly from the alphabet, with the exception of the letters I and O. Digits were chosen randomly from the set of numbers 2 to 9. Letters and numbers were black and were presented in the centre of a light-grey screen. They measured 0.32 degrees by 0.64 degrees when viewed at a distance of approximately 60 cm. The fixation stimulus was a small white circle presented in the centre of the screen.

Procedure.   Each trial began with the presentation of the fixation stimulus. Participants initiated the RSVP sequence by pressing the space bar. A brief tone accompanied the start of each stream. Each RSVP stream consisted of a pre-T1 stream of 4–11 digits, T1 (a letter), and a post-T1 stream of six items: 5 digits and 1 letter (T2). T2 occurred randomly in one of the six post-T1 positions (referred to as Lags 1–6). Each item in the stream was presented for 15 ms, and the SOA between items in the stream was 90 ms. Following the stream, the screen remained blank as the participants used the computer keyboard to indicate which two letters they saw. They were instructed to respond in the order in which they saw the items and were told that their responses were untimed. Responses were recorded by the computer.

The T1–T2 SOA ranged from 90 ms to 540 ms, in steps of 90 ms, corresponding to six different lags (Lag 1 = 90-ms SOA or immediately following T1). Participants completed one practice block of 30 trials, followed by a test block of 120 trials (20 trials at each SOA). Participants were instructed to proceed through the experiment at their own pace.

TM–TM task

The TM–TM task was the same as that in Experiments 1 and 2. The timing of the T1 and T2 sequences corresponded to the medium condition in the earlier experiments (that is, T1 = 30 ms, ISI = 15 ms, M1 = 60 ms; T2 = 30 ms, ISI = 15 ms, M2 = 60 ms).

Results and discussion

RSVP stream task

Note that in our design, when T2 occurred at Lag 6, it was the last item in the stream. Given recent reports (Giesbrecht & Di Lollo, 1998) that the T2 item must be followed by a mask in order for the blink to occur, we removed these items from the analysis.

The mean percentage of correct identification of T1 (collapsed across Lags 1–5) was 85% (SD = 9.4). A repeated measures ANOVA on T1 accuracy revealed a main effect of lag, $F(4, 60)$ = 14.89, p < .0001. As can be seen in Figure 4, where T1 is plotted as a function of lag, T1 performance was worse at shorter lags than at longer ones. That is, when T2 closely followed T1, T1 identification was impaired. That T1 accuracy is affected by T1–T2 Lag has been found before, but with some inconsistency. Raymond et al. (1992) found a lower T1 accuracy at short SOAs than at long ones. Seiffert and Di Lollo (1997) found the opposite pattern: T1 accuracy was best at the shortest (90-ms)

---

[8] For the RSVP task, software was provided courtesy of Kimron Shapiro and Jane Raymond. To obtain a copy of the software, write to Kimron Shapiro at University of Wales, Bangor; Email: k.shapiro@bangor.ac.uk
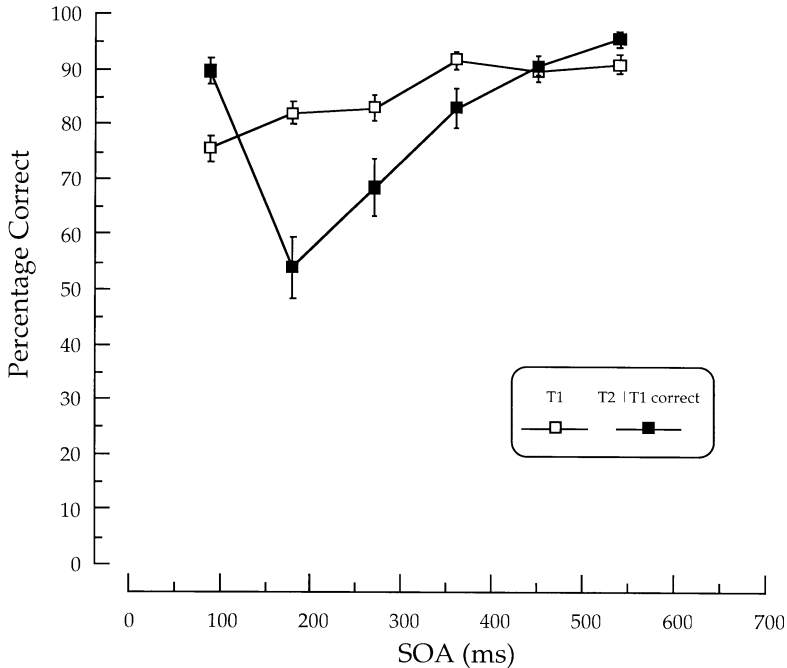
Figure 4. Results from Experiment 3, RSVP stream task. Target 1 (T1) accuracy and conditional Target 2 (T2|T1) accuracy at six levels of lag (corresponding to T1–T2 SOA). Error bars represent the between-observer standard error of the mean. At Lag 1, T2 immediately follows T1, leading to "Lag 1 sparing".

SOA. However, in their study, the linear relation between T1 accuracy and SOA disappeared when this point was removed from the analysis, and it was not apparent in any of their later studies.

The mean percentage of correct T2 (collapsed across lags 1–5) was 75.87 (SD = 21.33). Accuracy on T2 as a function of lag is plotted in Figure 4. These data were analysed with a repeated measures ANOVA, which revealed a significant main effect of lag on T2 accuracy, $F(4, 60) = 25.02$, $p < .0001$. Figure 4 depicts the characteristic reduction of T2 at short lags (Lags 2–4) as well as the "Lag 1 sparing" effect (described earlier), which can occur when T2 is presented immediately after T1.

TM–TM paradigm

Results from the TM–TM paradigm are presented in Figure 5. On this task, average T1 percentage correct (collapsed over Lags 1–5) was 81.16 (SD = 11.02). There was a marginally significant main effect of lag on T1 accuracy, as revealed by a repeated measures ANOVA, $F(4, 60) = 2.44$, $p < .06$. Target 2 accuracy (collapsed over Lags 1–5) was 54.88 (SD = 34.40). A repeated measures ANOVA revealed a significant effect of lag on T2 accuracy, $F(4, 60) = 114.03$, $p < .0001$. As in Experiments 1 and 2, T2 accuracy was lowest at the shortest lag and gradually improved across lags.
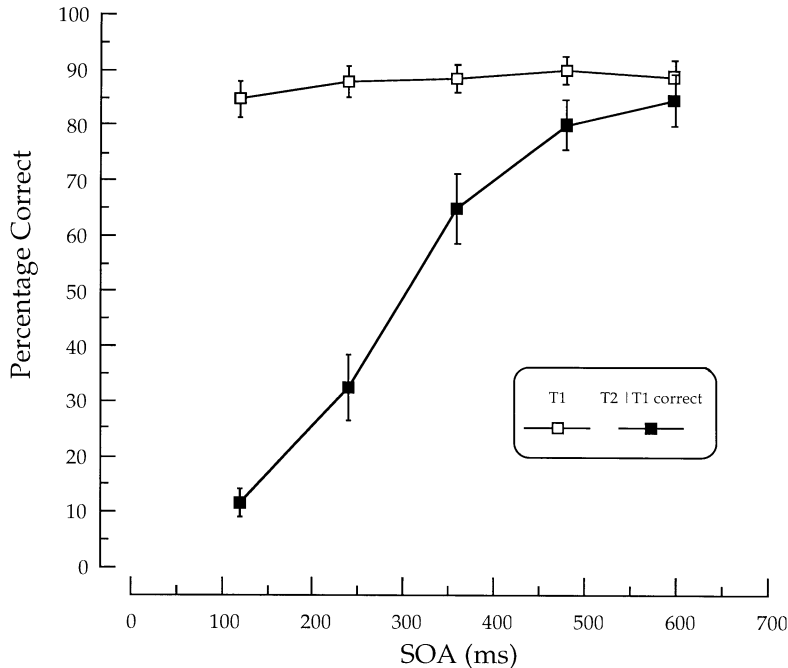
Figure 5.   Results from Experiment 3, target mask–target mask task. Target 1 (T1) accuracy and conditional Target 2 (T2|T1) accuracy at five levels of lag (corresponding to T1–T2 SOA). Error bars represent the between-observer standard error of the mean.

Comparison of the two paradigms

In examining the stream task and the TM–TM task, the first thing that becomes apparent is that whereas T1 performance is comparable across the two paradigms, T2 accuracy is much higher in the stream task (compare Figures 4 and 5). Whereas T2 performance (i.e., the group mean for any one lag) never goes below 52% in the stream task, it falls as low as 11% in the TM–TM task. This pattern might be explained as follows. In the TM–TM paradigm there is only 45 ms of uninterrupted time to process T1, whereas in the stream paradigm there is 90 ms. All other things being equal, one might therefore expect T1 performance to be worse in the TM–TM paradigm. However, because the two paradigms were blocked subjects could try to optimize their resource allocation to identify T1. For this purpose, the TM–TM paradigm provides an opportunity to aim resources at the perceptual moment of T1 because it is the first item presented. Such effective "aiming" of resource allocation would not be possible in the stream paradigm because the first item is never T1 and the position of T1 within the stream is randomized. Therefore, with attention more tightly focused and more intensely allocated toward T1 in the TM–TM paradigm, the performance disadvantage that one might have expected for T1 is counteracted and, following the well-accepted notion that increased attention to T1 will result in an increased blink, that is precisely what we find when comparing the two paradigms.

The critical comparison for our purposes is not with regard to overall accuracy, but with regard to the relation between T2 accuracy and lag. In order to test empirically the assumption that both paradigms are measuring the same underlying mechanism, we conducted a within-subjects correlation of blink strength or magnitude. Blink strength has been quantified before (Grandison et al., 1997; Raymond et al., 1995; Seiffert & Di Lollo, 1997; Shapiro et al., 1994). In all previous instances, however, blink magnitude was calculated on an experiment-wide basis by subtracting mean T2 performance at each lag from 100 (the presumed asymptotic level of T2 performance) and adding each of those values together. We chose not to focus on this method for two reasons.[9] First, because T1 accuracy on both tasks was less than 100%, we thought that 100 was not a realistic asymptote value in our task. Second, because we were interested in comparing individuals' patterns of responding, we used each person's blink and his or her own individual asymptotic. To determine an individual's strength of blink score we first examined his or her unique pattern of performance to determine at what lag(s) his or her blink occurred (explained in more detail later). We then subtracted each subject's mean accuracy during his or her blink from his or her own mean asymptotic performance.

We used a subject's mean performance on T1 as his or her asymptote. As T1 and T2 contained exactly the same perceptual information, we considered mean T1 performance to be a good reflection of a subject's ability to identify stimuli with our parameters. In order for a particular lag to be considered as part of the blink, performance at that lag had to be 15% below the subject's mean T1 accuracy. Using these criteria, the modal "length of blink" (number of lags judged to constitute part of the blink) was 3 in the stream task and 4 in the TM–TM task. Because of the Lag 1 sparing effect in the stream task, it is not surprising that the mean length of blink in this task was one lag shorter than in the TM–TM task. Whereas T1 performance was used to identify lags that were considered to be part of the blink, the strength of blink score was computed by subtracting T2 performance on each lag within the blink from average T2 performance outside of the blink. To further clarify this procedure, two examples are fully calculated in the Appendix.

As a final step, we correlated this single value, representing strength of blink, within subjects across both tasks. We found that the two values were significantly correlated, $R = .66$, $p < .05$ (see Footnote 9). We interpret this significant correlation to mean that our TM–TM paradigm is probably assessing the same underlying construct as that measured in an RSVP stream paradigm. Although we might have simply assumed this based on having found the same apparent relation between T2 and lag in our TM–TM task as in previous published reports of the AB, we have been able to quantify this association. As

---

[9] To be certain that this correlation was not a reflection of the specific procedures that we used to determine the strength of blink, we also calculated blink strength based on the procedures used by Grandison et al. (1997), Raymond et al. (1992), Seiffert and Di Lollo (1997), and Shapiro et al. (1994). That is, blink magnitude was calculated by subtracting a subject's performance on each lag from 100 (an idealized symptotic value) and summing those five values. The within-subjects correlation between strength of blink scores calculated in this way for each of the two tasks was significant, $R = .50$, $p < .05$. Although still statistically significant (indicating that performance on the two tasks is correlated), the slightly weaker correlation using this method probably reflects the fact that 100 is not a reasonable asymptote under our testing conditions. For this reason, we prefer the first method of computing the strength of blink, for which each subject's asymptote is based on his or her own performance on our task.

well, we have provided preliminary evidence that the blink may be a stable within-subjects individual difference characteristic.

GENERAL DISCUSSION

This study was designed to examine the role of target-processing difficulty (via a data limitation) on the attentional blink, in the absence of switching. In Experiment 3, we found a significant within-subjects correlation between participants' performance on our TM–TM task and a more traditional stream task. Thus, we confirmed that the TM–TM paradigm introduced here is likely to be tapping into the same underlying construct as that measured in previous studies of the attentional blink.

Using this paradigm, we manipulated the perceptual quality of T1 (Experiment 1) and T2 (Experiment 2), by co-varying the duration of the target and the mask. Importantly, this was a strictly data-limited manipulation. As difficulty was randomly manipulated within a block (vs. the blocked approach employed in most previous studies in this area), this difficulty manipulation could not affect resource allocation in advance of the target. It would not be until after presentation of the target that the participant would know whether it was difficult to encode. Strategic recruitment of resources to deal with a difficult target would likely be too late to be effective. Thus even if possible, such rapid recruitment would not be an efficient strategy. Indeed, one implication of our finding that the blink is unaffected by our random manipulation of the target's perceptual quality might be that this inefficient strategy is not used by our participants in this study.

Results from Experiment 1 showed that increased perceptual difficulty had the expected effect on T1 accuracy, but no effect on the subsequent blink. That is, as long as a durable representation of T1 was encoded (as reflected by correct performance on T1), there was a fixed effect on subsequent T2 processing. We obtained similar results in Experiment 2—perceptual difficulty had an effect on overall accuracy, but not on the blink function. Together, these results demonstrate that perceptual quality is independent of backward masking in producing the blink. That is, degradation in the perceptual quality of either target is additive with whatever post-perceptual mechanism is proposed to produce the blink.

Implications for theories of the blink

One goal of this study was to investigate the relation between target difficulty and the blink to see if the results could help to determine which of the two predominant classes of model (interference and bottleneck) was better supported and to place constraints on these models. Our findings do not undoubtedly favour one model over the other, but they do place significant constraints and require modifications of both models.

As reviewed earlier, the interference model of the blink predicts that the relative weightings in VSTM (and not target difficulty) modulate the AB. Insofar as we found no significant relation between target difficulty and the blink, our results are consistent with the interference hypothesis. However, the results are less clear with regard to weightings in the VSTM. The interference model maintains that one factor that affects weight assignment is similarity. Items that are similar to the target templates should

receive high weightings in VSTM and interfere more with correct retrieval of the target from VSTM (see Shapiro & Raymond, 1994). In our TM–TM paradigm, the similarity of the distractors to the targets was consistent in the different difficulty conditions: The targets were always letters, and the distractors were always pattern masks. In this way, our paradigm might be regarded as an improvement over studies that mixed target-distractor similarity with the difficulty manipulation (e.g., Chun & Potter, 1995; Raymond et al., 1995; Shapiro et al., 1994). As we did not manipulate similarity in our paradigm, the interference model would have predicted our results.

However, similarity is not the only factor that is presumed to influence target weightings in the interference model. Weightings are also affected by the temporal contiguity between targets and distractors (Shapiro et al., 1994). As temporal contiguity is not clearly operationalized in the interference model, we are unclear as to whether our results are consistent with this or not. In our paradigm, SOA between the targets and their masks varied systematically with the difficulty manipulations (SOA was longer in the easy condition than in the hard condition). On the other hand, the ISI between each target and its mask remained fixed for all difficulty manipulations. Clearly then, depending on how the interference model operationalizes "temporal contiguity", the relation between temporal contiguity and weightings in the VSTM is either supported or not supported by this study.

The interference model also presumes that there is a fixed amount of weighting, and that at short T1–T2 SOAs, T2 may receive a lesser weighting because much of the total has already been allocated to T1 and the item immediately following it (Shapiro et al., 1994). If the duration of an item is related to its weighting, the interference model would predict exactly the opposite pattern from the one that we found. In our easy condition, the T1 item is longer in duration than M1. As it is longer, and more similar to the target template, T1 should receive a high proportion of the available weightings, leaving relatively less for T2. On the other hand, in the hard condition, T1 is short and M2 is long. The first target would receive a weighting for being similar to the target template, but this would be less than that in the easy condition because its duration was shorter, and its similarity to the target template would be less obvious because of the reduction in the quality of the evidence. Such a scenario, which could be consistent with the interference model, depending on how "weightings in VSTM" were operationalized, would actually predict that larger weightings would be available for T2 in our hard condition than in our easy condition. This would predict a bigger blink in the easy condition, a result not found here.

We have illustrated a number of scenarios whereby our results could be either consistent with or inconsistent with the interference model, depending on what features can alter weightings in VSTM. Although we have not been able to support or reject the interference model, we have certainly provided some factors that must be accounted for. Our findings show that when target–distractor similarity and ISI are held constant, but target and distractor duration and SOA are permitted to vary, the blink is unaffected. Note that despite these possible limitations in the operationalization of the interference model in its original form, the interference model nevertheless predicts the results found here, that T1 difficulty is not always related to size of blink.

As one of our goals was to compare the models, it is important to note that whereas slight modifications of the interference model will enable it to handle our findings, at least one pattern of results has been very problematic for it. The interference model has no way of accounting for Giesbrecht and Di Lollo's (1998) and Brehaut et al.'s (1999) observation that T2 must be masked with a backward (but not simultaneous) mask in order for an AB to occur. In contrast, the bottleneck theories have been able to account for most of the findings to date.

Recall that the bottleneck model predicts that anything that increases T1 difficulty should lead to a larger blink (e.g., Grandison et al., 1997; Potter et al., 1998; Seiffert & Di Lollo, 1997). When T1 is difficult, Stage 2 processing is delayed, making it more likely for T2 to decay while waiting for access to Stage 2. In its traditional form, our results clearly do not support the bottleneck model. We believe, however, that with one crucial modification to the bottleneck model, our results could be accommodated. Under the bottleneck model, any manipulation of T1 difficulty is presumed to lead to a delay of Stage 2 processing of T1. We have already noted, however, that to date all of the evidence in support of this assertion has been found using a difficulty manipulation that implicates resource allocation, either because of the nature of the difficulty manipulation itself, by virtue of the fact that the difficulty manipulation was blocked, or by the fact that resources were engaged by other task features (e.g., switching). Here, in the first study using only a data limitation to implement the difficulty manipulation, we have demonstrated that target difficulty has no effect on the blink. This finding can be accommodated within the bottleneck class of models merely by assuming that only difficulty at the post-perceptual level should affect the blink. In contrast, any manipulation that affects perceptual processing and encoding only should not affect the blink (see Posner, 1978; Posner & Boies, 1971; Posner & Klein, 1973, for arguments that encoding processes typically do not interfere with other mental activities). Notice that the main theoretical modification is that the bottleneck is explicitly defined as being post-perceptual. Vogel, Luck, and Shapiro (1998) have provided converging electrphysiological evidence for a post-perceptual locus of the AB.

Jolicoeur's (1998) central limitation theory (see also Jolicoeur & Dell-Acqua, 1998) is a recent version of the bottleneck model that does specifically place the bottleneck at a post-perceptual stage. According to this model, the bottleneck in the traditional paradigm occurs after sensory encoding and perceptual encoding of the first target, in the stage during which T1 is encoded into short-term memory, called short-term consolidation (STC). The STC stage is serial and capacity limited. Jolicoeur (1998) draws a parallel between sensory and perceptual encoding in the central limitation theory with Stage 1 in the two-stage model, and the STC in the central limitation theory with Stage 2 in the two-stage model. Insofar as this new model explicitly states that the bottleneck is post-perceptual, it may be able to account for the present results better than does the earlier two-stage model.

Also, another important feature of the central interference theory is that whatever makes use of central resources before the onset of T2 could contribute to the blink. This is not limited to STC, but may also include response selection of the first target, as well as task switching (Jolicoeur, 1998). Earlier we hypothesized that task, location, and set switching, as well as global interference from the distractor stream and task instruc-

tions, may have modulated the relation between T1 difficulty and the blink observed in previous studies. Jolicoeur provides a context for these arguments: Whatever makes use of central resources prior to T2 can have an effect on the blink. In our study, central resources (which could be allocated in advance of T1 to its processing) were held constant across the difficulty conditions. Also, our TM–TM task, in and of itself, did not recruit any central resources other than those required for identifying two masked targets (unlike previous paradigms that included, for example, task switching). Under these conditions, we saw no resulting effect on the blink. So then, insofar as the central interference theory proposes that the bottleneck is post-perceptual, and that it can be modified by utilization of resources before the bottleneck, it is entirely consistent with our data and interpretation.

However, before accepting Jolicoeur's (1998) model outright, it must be noted that Jolicoeur does include one caveat that makes it inconsistent with our results. In order to account for previous findings in the literature demonstrating a relation between target identification difficulty and the blink, Jolicoeur stated that "increasing the duration of STC_1 itself *or of any stage of processing before STC_1* would likely lengthen the period of postponement of STC_2, leading to a larger and longer AB effect. Thus, the fact that manipulations believed to have an effect on the perceptual procesing of T1 (such as different levels of masking) modulate the magnitude of the AB effect is entirely consistent with the Central Interference Theory" (p. 1030, emphasis in original).[10] In contrast to this proposal about T1 manipulations, we have demonstrated here that, in fact, increasing the difficulty of perceptual processing through a data-limited manipulation does not, in and of itself, postpone the bottleneck.

We propose an important modification to the bottleneck models of the blink. The critical factor that causes the bottleneck is the utilization of central resources before T2 onset. Factors that use central resources could include task switching (Allport, Styles, & Hsieh, 1994; Potter et al., 1998; Rogers & Monsell, 1995), response selection (Jolicoeur, 1998), and set or expectancy. In contrast, manipulations of the perceptual quality of data should have no impact on utilization of central resources (Norman & Bobrow, 1975), unless the subject is aware of the manipulation and is able to prepare appropriately. As such, factors related to sensory encoding (including time taken to encode a stimulus) are less important than utilization of resources in modulating the AB.

The role of object substitution in the attentional blink

Masking by object substitution (Enns & Di Lollo, 1997) has recently been proposed as an explanatory mechanism for the attentional blink (Brehaut et al., 1999). This form of masking requires two conditions: Attention must not be focused on the target, and some stimulus (a mask) must be presented following the target (Enns & Di Lollo, 1997; see also Di Lollo, Enns, & Rensink, in press). Under these conditions, the target is replaced by the trailing item and is thus unavailable for report. This concept is qualitatively different from the similar-sounding theory of masking by interruption (cf. Turvey, 1973) in that the underlying information-processing hierarchy consists of

---

[10] STC_1 is the STC for the first target, STC_2 is the STC for the second target.

iterative feed-back loops. That is, with masking by interruption it is assumed that processing proceeds from early stages to later stages in sequence, whereas with object substitution, processing involves the establishment of re-entrant connections, which develop a representation of the target (Di Lollo et al., in press). These iterative loops form such a representation by comparing models from later stages with the more veridical information present in earlier processing stages. The main support for an object substitution explanation of backward masking over the more traditional erasure account is the demonstration of zero–SOA masking (Bischof & Di Lollo, 1995). The mask is presented coincident with the target and simply remains on the screen after the target has been extinguished. Robust masking is observed, even when the mask—something as simple as four dots—does not share any contours with the target. As the target and mask onset at the same time, and the mask simply continues beyond the target, there is no opportunity for the mask to interrupt processing of the target. Instead, it is assumed that the crystallizing representation of the target is replaced by the mask, because at some point there is no evidence for the target and a great deal for the mask. Importantly, it has been shown that the allocation of attention to the location of a target can rescue the target from the effects of object substitution or conversely, that a precondition for object substitution to occur is that the target's location not be attended (Di Lollo & Enns, 1998). To apply object substitution to the AB, it must be assumed that the attentional resources recruited by T1 processing perform a function analogous to spatial attention, and hence when they are unavailable for processing T2 a blink will occur.

In the context of the attentional blink, support for the object substitution account comes from the finding that T2 must be masked by a backwards mask (Giesbrecht & Di Lollo, 1998); a simultaneous mask will not produce a blink (Brehaut et al., 1999). With regard to the present experiments, all of the masks are presented 15 ms after the relevant target and are thus backwards masks and should produce a blink. As discussed earlier (see Footnote 6), the fact that we observe a blink implies that our mask is indeed a backwards mask, albeit with some amount of integration masking also occurring. The main support for the role of object substitution in the present data consists of the additive pattern observed in Experiment 2 between our difficulty manipulation and the blink (see Figure 3). Consider that masking by object substitution is an "all or none" phenomenon—either the target is replaced, or its representation is computed. If object substitution of the second target occurs, there is a blink. In the event that no object substitution occurs, any letter identification process would be hindered by a degraded representation, which would result from the difficulty manipulation that we employed. As our results consist of an average, we cannot distinguish which targets were substituted and which were not. As such, this additive pattern should not be taken as strong support for, but rather a set of results that is entirely consistent with, the model proposed by Di Lollo et al. (in press). Support for the model would come from a direct test of the hypothesis by masking T2 with an object substitution mask (e.g., four dots). Just such a test has recently been reported (Giesbrecht, Bischof, & Kingstone, 1999) and claims to disprove the object substitution hypothesis. A resolution of this issue will have to await further testing.

The object substitution account of the AB is very similar in nature to the bottleneck models discussed earlier, with the exception that the mechanism whereby the identity of the target is computed has been made computationally explicit. Di Lollo et al., (in press)

have outlined a closed-loop control model, which implements the re-entrant architecture discussed earlier. This is a clear advance over other models with constructs that can be difficult to interpret. We feel that this framework has a great deal of potential in understanding specifically the AB and masking in general.

Summary and implications

The main finding from the present study is that making T1 and T2 more difficult to identify had no influence on the blink function. Importantly, we used a difficulty manipulation based on perceptual quality, mixed the difficulty conditions within a block, and avoided switching of all kinds. With these key features, our paradigm did not lead to the utilization of different amounts of central resources across the different difficulty conditions, or to the utilization of resources prior to the bottleneck other than for the task at hand (i.e., target identification).

We also provided correlative evidence that our modified TM–TM paradigm was tapping into the same construct as that assessed in the more traditional stream paradigm. The strategy used here (i.e., correlating the measured strength of blink across paradigms) could be used to confirm that paradigms that study cross-modal and auditory attentional blinks, or those that use a psychological refractory period approach, are, in fact, all measuring the same underlying processing limitation.

These data require modifications of both classes of theories of the blink. Although the interference model can predict our main findings, further operationalization is required to account for the weightings of items in VSTM in this new paradigm. With regard to bottleneck models, their tenet that blink magnitude should always be correlated with T1 difficulty must be weakened. Instead, only those difficulty manipulations that alter utilization of Stage 2 resources should have an effect on the subsequent blink. Perhaps most importantly, these data supports Jolicoeur's (1998) notion (and Vogel et al.'s, 1998, electrophysiological evidence) that the bottleneck that creates the blink is post–perceptual.

We have undoubtedly shown that when a masking-induced data-limited difficulty manipulation is mixed within a block and examined with a blink paradigm without switching, T1 difficulty has no effect on the magnitude of the attentional blink, and T2 difficulty has an additive, but not interactive, effect. It still remains unclear as to which features of our study were most influential in eliminating the relation between T1 difficulty (as manipulated by masking of T1) that has been demonstrated in previous studies (e.g., Grandison et al., 1997; Seiffert & Di Lollo, 1997). As mentioned earlier in our review of the literature, we were unable to identify with confidence any one feature (e.g., blocking the difficulty manipulation, task switching) or set of features that modulates the relation. This is likely because each study incorporated a different set of task features or difficulty manipulations, making it hard for a clear pattern to emerge.

The next step in this line of research should be to discover under which conditions a relation between T1 or T2 difficulty and the blink is observed. For example, what will happen if task, set, or location switching is added to our TM–TM paradigm? If we were to block the difficulty manipulation used here, but retain our TM–TM task as it is, would this be sufficient to see the effect observed by others? Our speculative explanation for the pattern of results in Experiment 3 (T1 performance is about the same in the two

paradigms, whereas the blink is much larger in TM–TM) suggests a "yes" in answer to this question. Conversely, if previous studies that found a relation were administered with critical changes (e.g., mixing the difficulty manipulation used by Grandison et al., 1997, with or without the features of task and set switching, which existed in their original studies), would the observed relation then disappear? The more that we can discover about the conditions under which the blink is dependent on T1 or T2 difficulty, the better we will understand the nature of the limited-capacity system thought to be implicated in the blink.

REFERENCES

Allport, A., Styles, E.A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 421–452). Cambridge, MA: MIT Press.

Arnell, K.M., & Duncan, J. (1999). *Specific and general sources of dual-task cost in stimulus identification and response selection*. Manuscript submitted for publication.

Bischof, W.F., & Di Lollo, V. (1995). Motion and metacontrast with simultaneous onset of stimuli. *Journal of the Optical Society of America, A, 12*, 1–14.

Boucart, M., Moroni, C., Fuentes, L.J., & Belin, C. (1998). Variations of AB as a function of task and semantic similarity. *Abstracts of the Psychonomics Society 39th Annual Meeting, Dallas, Texas*, U.S.A. (p. 72).

Brehaut, J.C., Enns, J.T., & Di Lollo, V. (1999). Visual masking plays two roles in the attentional blink. *Perception & Psychophysics, 61*, 1436–1448.

Broadbent, D.E. (1958). *Perception and communication*. London: Pergamon Press.

Broadbent, D.E., & Broadbent, M.H. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception & Psychophysics, 42*, 105–113.

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.

Di Lollo, V., & Enns, J.T. (1998). Perceiving with and without attention: Competition for consciousness among visual events. *Abstracts of the Psychonomics Socieity 39th Annual Meeting, Dallas, Texas*, U.S.A. (p. 53).

Di Lollo, V., Enns, J.T., & Rensink, R.A. (in press). Competition for conciousness among visual events: The psychophysics of reentrant visual processes. *Journal of Experimental Psychology. General*.

Duncan, J. (1998). Public remarks made at the XVIIIth meeting of the Interntional Society for the Study of Attention and Performance, Englewood Green, U.K.

Duncan, J., Ward, R., & Shapiro, K. (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.

Giesbrecht, B., Bischof, W., & Kingstone, A. (1999, November). *Masking in the attentional blink: Failure of the object substitution hypothesis*. Paper presented at the 40th meeting of the Psychonomic Society, Los Angeles, U.S.A.

Giesbrecht, B., & Di Lollo, V. (1998). Beyond the attentional blink: Visual masking by object substitution. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 1454–1466.

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.

Jolicoeur, P. (1998). Modulation of the attentional blink by on-line response selection: Evidence from speeded and unspeeded Task 1 decisions. *Memory & Cognition, 26*, 1014–1032.

Jolicoeur, P., & Dell'Acqua, R. (1998). The demonstration of short-term consolidation. *Cognitive Psychology, 36*, 138–202.

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.

Norman, D.A., & Bobrow, D.G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology*, *7*, 44–64.

Posner, M.I. (1978). *Chronometric explorations of mind*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Posner, M.I., & Boies, S.J. (1971). Components of attention. *Psychological Review*, *78*, 391–408.

Posner, M.I., & Klein, R.M. (1973). On the functions of consciousness. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 21–35). New York: Academic Press.

Potter, M.C., Chun, M.M., Banks, B.S., & Muckenhoupt, M. (1998). Two attentional deficits in serial target search: The visual attentional blink and an amodal task-switch deficit. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 979–992.

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.

Rogers, R.D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, *124*, 207–231.

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., Arnell, K.M., & Raymond, J.E. (1997). The attentional blink: A view on attention and a glimpse on consciousness. *Trends in Cognitive Science*, *1*, 291–296.

Shapiro, K.L., & Raymond, J.E. (1994). Temporal allocation of visual attention: Inhibition or interference? In D. Dagenbach & T.H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 151–188). San Diego, CA: Academic Press.

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.

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.

Visser, T.A.W., Bischof, W.F., & Di Lollo, V. (1999). Attentional switching in spatial and non-spatial domains: Evidence from the attentional blink. *Psychological Bulletin*, *125*, 458–469.

Vogel, E.K., Luck, S.J., & Shapiro, K.L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1656–1674.

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

Ward, R., Duncan, J., & Shapiro, K.L. (1997). Effects of similarity, difficulty, and nontarget presentation on the time course of visual attention. *Perception & Psychophysics*, *59*, 593–600.

APPENDIX

Calculating strength of blink in the stream and TM–TM tasks: Sample calculations for one subject.

Stream task

Step 1

Lag	1	2	3	4	5	6
T1 percentage correct	80	95	80	95	100	95

Average T1 performance over 6 lags = 90.83%. This is this subject's asymptotic performance, given our parameters. To be considered part of the blink, T2 performance would have to be 15% below the asymptote value, that is 75.83% or less (90.83 − 15 = 75.83).

Step 2

Lag	1	2	3	4	5	6
T2|T1 percentage correct	100	42.1	56.3	52.6	85	94.7

Performance at Lags 1, 5, and 6 at not or near asymptotic T1 performance (as defined in Step 1). Therefore they are not considered to be part of the blink. Performance at Lags 2, 3, and 4 are below the cutoff defined in Step 1 (75.83%) and are therefore considered to be part of the blink. This subject's length of blink is 3 lags.

Step 3

The strength of blink score is computed by subtracting performance on each blinked lag from the average on non-blinked lags. In this case, the average non-blinked performance is (Lag 1 + Lag 5 + Lag 6)/3 = (100 + 85 + 94.7) / 3 = 93.23. Therefore, strength of blink = (93.23 − Lag 2) + (93.23 − Lag 3) + (93.23 − Lag 4) = (93.23 − 42.1) + (93.23 − 56.3) + (93.23 − 52.6) = 128.69.

TM–TM task

Step 1

Lag	1	2	3	4	5
T1 percentage correct	85	86.7	95	96.7	91.7

Average T1 performance over 5 lags = 91.02%. This is this subject's asymptotic performance, given our parameters. To be considered part of the blink, T2 performance would have had to be 15% below the asymptote value, that is 76.02% or less (91.02 − 15 = 76.02).

Step 2

Lag	1	2	3	4	5
T2|T1 percentage correct	19.7	22.3	58	81	82.3

Performance at Lags 4 and 5 are at or near asymptotic T1 performance (as defined in Step 1). Therefore, they are not considered to be part of the blink. Performance at Lags 1, 2, and 3 are below the cutoff described in Step 1 (76.02%), and are therefore considered to be part of the blink. This subject's length of blink is 3 lags.

Step 3

The strength of blink score is computed by subtracting performance on each blinked lag from the average on non-blinked lags. In this case, the average non-blinked performance is (Lag 4 + Lag 5) / 2 = (81 + 82.3) / 2 = 81.65. Therefore, strength of blink = (81.65 − Lag 1) + (81.65 − Lag 2) + (81.65 − Lag 3) = (81.65 − 19.7) + (81.65 − 22.3) + (93.23 − 58) = 144.95.