

Research Article

Temporal Selection Is Suppressed, Delayed, and Diffused During the Attentional Blink

Edward Vul, Mark Nieuwenstein, and Nancy Kanwisher

Department of Brain and Cognitive Science, Massachusetts Institute of Technology

ABSTRACT—*How does temporal selection work, and along what dimensions does it vary from one instance to the next? We explored these questions using a phenomenon in which temporal selection goes awry. In the attentional blink, subjects fail to report the second of a pair of targets (T1 and T2) when they are presented at stimulus onset asynchronies (SOAs) of roughly 200 to 500 ms. We directly tested the properties of temporal selection during the blink by analyzing distractor intrusions at a fast rate of item presentation. Our analysis shows that attentional selection is (a) suppressed, (b) delayed, and (c) diffused in time during the attentional blink. These effects are dissociated by their time course: The measure of each effect returns to the baseline value at a different SOA. Our results constrain theories of the attentional blink and indicate that temporal selection varies along at least three dissociable dimensions: efficacy, latency, and precision.*

Selective attention is the process of choosing a subset of the sensory input for detailed perceptual analysis. Although most research on attention has focused on the ability to select locations in space and the objects that occupy them, people can also select moments in time and the events that unfold over them. In the experiment reported here, we investigated the nature of temporal selection by testing a case in which it goes awry: the attentional blink (AB).

Mark Nieuwenstein is now at the Department of Cognitive Psychology, Vrije Universiteit, Amsterdam, The Netherlands. Address correspondence to Edward Vul, Department of Brain and Cognitive Science, Massachusetts Institute of Technology, 77 Massachusetts Ave. 46-4141, Cambridge, MA 02139, e-mail: evul@mit.edu.

The selection of a single item from a rapid serial visual presentation (RSVP) sequence can be modeled as the operation of a linear system: In such a model, attention is a gate that can be described as a spatiotemporal cue-impulse response function (Reeves & Sperling, 1986; Shih & Sperling, 2002). This function can be characterized empirically by noting which items from the visual stream are reported given a particular cue (Botella, Arend, & Suero, 2004; Botella, Garcia, & Barriopedro, 1992); from this information, one can infer the temporal and spatial extent of selection (Reeves & Sperling, 1986; Shih & Sperling, 2002; Weichselgartner & Sperling, 1987). In the more complex case in which two targets (T1 and T2) must be selected in rapid succession (within 200–500 ms of each other), observers often miss the second, a phenomenon known as the AB (Raymond, Shapiro, & Arnell, 1992). Most research on the AB has not characterized the temporal profile of selection of T2, but has instead simply reported the conditions under which T2 is reported correctly (Chun & Potter, 1995; Raymond, Shapiro, & Arnell, 1995; but see Botella et al., 2004, and Chun, 1997a). Despite a recent demonstration that the AB reflects a failure of selection rather than memory (Nieuwenstein & Potter, 2006), single-item models of selection have not generally been applied to the AB, which has instead been interpreted within symbolic-processing models (Chun, 1997b; Chun & Potter, 1995) as a limit in type-token binding (Bowman & Wyble, 2007; Kanwisher, 1987) or as a broader processing constraint on attention (Di Lollo, Kawahara, Shahab Ghorashi, & Enns, 2005). To bridge the two research traditions, we measured the extent of temporal selection during the AB by applying the methods of single-item temporal selection to AB conditions.

We considered three dimensions of temporal selection in order to examine how it is altered for T2 to produce the AB. Selection could be less effective (i.e., suppressed; see Fig. 1a),

meaning that fewer items in the vicinity of the target are selected. This would result in a concomitant increase in the rate of random guessing (observable in more frequent reports of items from remote serial positions in the stimulus sequence). This scenario is consistent with several accounts of the AB (Olivers, van der Stigchel, & Hulleman, 2007; Raymond et al., 1992). Alternatively, or in addition, the time of engagement of selection could be less accurate (e.g., delayed; see Fig. 1b), meaning that subjects systematically report letters from incorrect temporal positions. This scenario is consistent with another group of AB theories (Chun, 1997a; Chun & Potter, 1995; Nieuwenstein, Chun, van der Lubbe, & Hooge, 2005). Finally, temporal selection could be less precise, or spread out over time (i.e., diffused; see Fig. 1c), so that subjects reliably report letters further in time from the cue. In the limit, of course, diffusion and suppression would produce identical patterns of reports (equally

frequent guesses from all list positions), but the two possibilities differ outside this limiting case.

Because each stimulus letter occurred in only one list position in each RSVP sequence of our experiment, the serial position of any reported letter was known. Therefore, we were able to precisely measure all three dimensions of temporal selection as a function of the stimulus onset asynchrony (SOA) between T1 and T2. We found that all three dimensions of selection are altered at short SOAs: Selection of T2 is suppressed, delayed, and diffused. Moreover, these three changes are dissociable by their time courses: Diffusion lasts about 400 ms after T1, suppression lasts about 500 ms, and delay persists for longer than 800 ms.

Our data are not consistent with the published accounts of the AB, none of which entail multiple processes with different time courses. More broadly, our findings reveal that temporal selection can vary along three independent dimensions: efficacy, latency, and precision.

METHOD

Participants

Twelve participants (7 female, 5 male) between the ages of 18 and 35 were recruited from the Massachusetts Institute of Technology subject pool and were paid for participation.

Materials

On each trial, subjects saw an RSVP stream composed of one instance of each of the 26 English letters in a random order. Each letter was presented for 25 ms and was followed by a 58.3-ms blank, resulting in an RSVP rate of 12 items/s. The letters were white, uppercase, and in 48-point Courier font. They were presented on a black background. With our resolution (1024×768), monitor (Viewsonic G90f), and viewing distance (roughly 50 cm), the letters subtended roughly 2.5° of visual angle.

On each trial, two cues in the RSVP stream indicated which of the letters were targets. The cues were white annuli with an inner diameter of 2.8° and an outer diameter of 3.2° ; thus, they appeared as rings around the cued letters. When a cue appeared, it was shown for 25 ms, concurrently with the target letter (see Fig. 2).

Onset of the first cue (T1 onset) was randomly counterbalanced to occur equally often in Positions 6 through 10 in the RSVP stream. The delay between the first and second cues (T1-T2 SOA) was manipulated over 10 levels, from 1 to 10 items (83–833 ms).

The experiment was programmed in Psychophysics Toolbox (Brainard, 1997; Matlab 7, Windows XP).

Procedure

Each participant began the experiment with 2 practice trials that had a randomly selected T1-T2 SOA; the results of these trials were discarded. Following the practice trials, participants completed four blocks of 100 trials each.

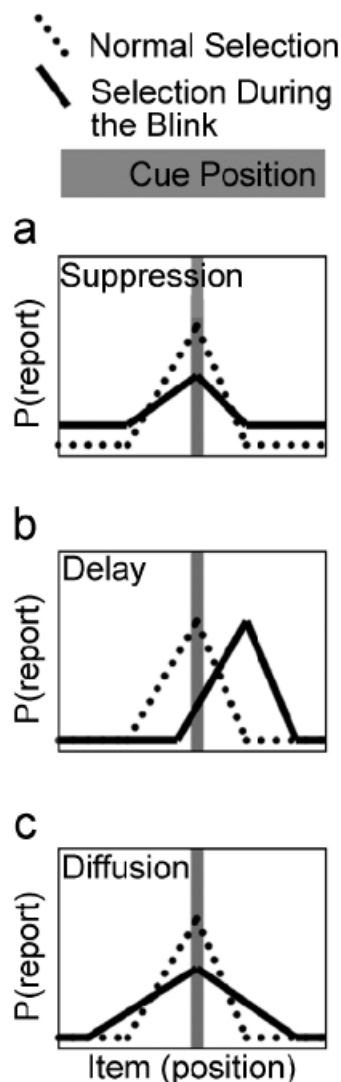


Fig. 1. Illustration of the three ways in which temporal selection may change during the attentional blink: (a) suppression, (b) delay, and (c) diffusion. Each graph shows the probability of an item being reported as a function of the item's position in the stimulus stream.




Fig. 2. Experimental design. On a given trial, all 26 English letters were presented in a random order without repetition. Annulus cues were presented simultaneously with the 2 letters designated as targets (T1 and T2). Subjects were asked to report the identities of the 2 cued letters in order. Stimulus onset asynchrony (SOA) varied across trials.

Each block contained 2 instances of each of the 50 possible combinations of T1 onset (five levels: Positions 6–10) and T1-T2 SOA (10 levels: SOAs of 1–10 items). For instance, in a trial with T1 onset at Position 6 and a T1-T2 SOA of 4 items, the two cues would appear during the 6th and 10th items in the sequence. These conditions appeared in a random order within each block.

At the end of each trial, subjects were asked to indicate which two letters they thought had been cued by the annuli. Subjects reported the letters by pressing the corresponding keys on the computer keyboard. Duplicate letters were not accepted, and subjects were told to report the first letter first and the second letter second. After each trial, subjects were told how many points they had earned (1 point possible for each target). In addition to receiving the flat rate of \$10 for participation, participants were offered bonus cash awards for performance: \$0.01 for each letter reported correctly (on average, subjects reported about 200 letters correctly, receiving a \$2 bonus). This bonus was provided as an incentive for subjects to try to report the cued letters correctly, given the particularly difficult conditions.

RESULTS

Our analyses are all based on the serial position of the reported letters at different T1-T2 SOAs. We computed the distribution of reports for T1 from the first responses and the distribution of reports for T2 from the letters reported second. Figure 3 shows the distribution of the serial position of responses for all of the SOA conditions (1–10 items), shifted to the onset of the first cue. The distribution of letters reported for T1 (see Fig. 3a) was unaffected by the T1-T2 SOA, indicating that the AB indeed affects the attentional response to the second cue only. The familiar curve of T2 accuracy in the AB is clear for both T2

report and T2 report when T1 was reported correctly (the thick black lines in Figs. 3b and 3c, respectively). Lag-1 sparing is more clearly seen in the latter case, presumably because the constraint that T1 be correct excludes trials that contain swap errors, known to occur commonly when T2 follows T1 in quick succession (Chun & Potter, 1995). Otherwise, the distribution of T2 reports was the same on trials when T1 was reported as item -2 , -1 , 0 , 1 , or 2 , which indicates that it was the cue—the initialization of the selection episode—not which item was reported, that determined the latency between T1 and T2 that produced the AB.

As did Chun (1997a), we saw many posttarget intrusion errors in the T2 reports. However, our data contained substantially more posttarget intrusions, and the distance of the intrusions from the target was greater than Chun found, presumably because of the faster RSVP rate we used. Indeed, at SOAs of three, four, five, and six items, the posttarget item was reported more often than the target itself, and intrusions from distractors two or more serial positions after T2 were also substantially increased above chance. The increase in posttarget intrusions at a higher RSVP rate suggests that the observed effects of delay reflect a delay in time, rather than serial position.

As described in the introduction, the properties of selection around each target can be estimated from the distribution of reported letters. We estimated efficacy of selection as the average probability of reporting an item from a seven-item window around the target (spanning three items before to three items after the target¹). This efficacy measure (see the appendix) is independent of delay and diffusion, as it does not depend on which item is reported (the target or earlier or later items within the window). If selection were totally ineffective, such that the cues had no effect on which items were reported, this quantity would correspond to chance ($1/26$ for T1 and $1/25$ for T2: .0385 and .04, respectively). If the cues were so effective that only items from within this window were reported, this quantity would be $1/7$, or .1429. The degree to which this efficacy measure is lower for T2 than for T1 is the degree to which selection is suppressed (rendered less effective) because of the AB.

Figure 4a shows efficacy as a function of SOA for both T1 and T2 report. It is clear that reports were above chance (close to ceiling for T1). However, for T2, a trend corresponding to the AB curve is evident. During the AB window (SOAs of 166–417 ms), there was a clear drop in the probability of reporting an item from the window around the target. Efficacy was greater at an SOA of 83.3 ms (a finding that reflects some lag-1 sparing), and recovered to T1 levels at an SOA of 500 ms. These results show that

¹Varying the size of the window within a reasonable range did not result in qualitative changes to the results. Also, note that given a limited measurement window, there is a possibility that delay and suppression are not orthogonal measures: If delay is so great that subjects report items well outside the window, then our measure of suppression will reflect this change. We chose our selection window to capture all of the items around each target that were reported at above-chance frequencies, thus ensuring that all the effects of delay and suppression were captured within the window, and were thus orthogonal.

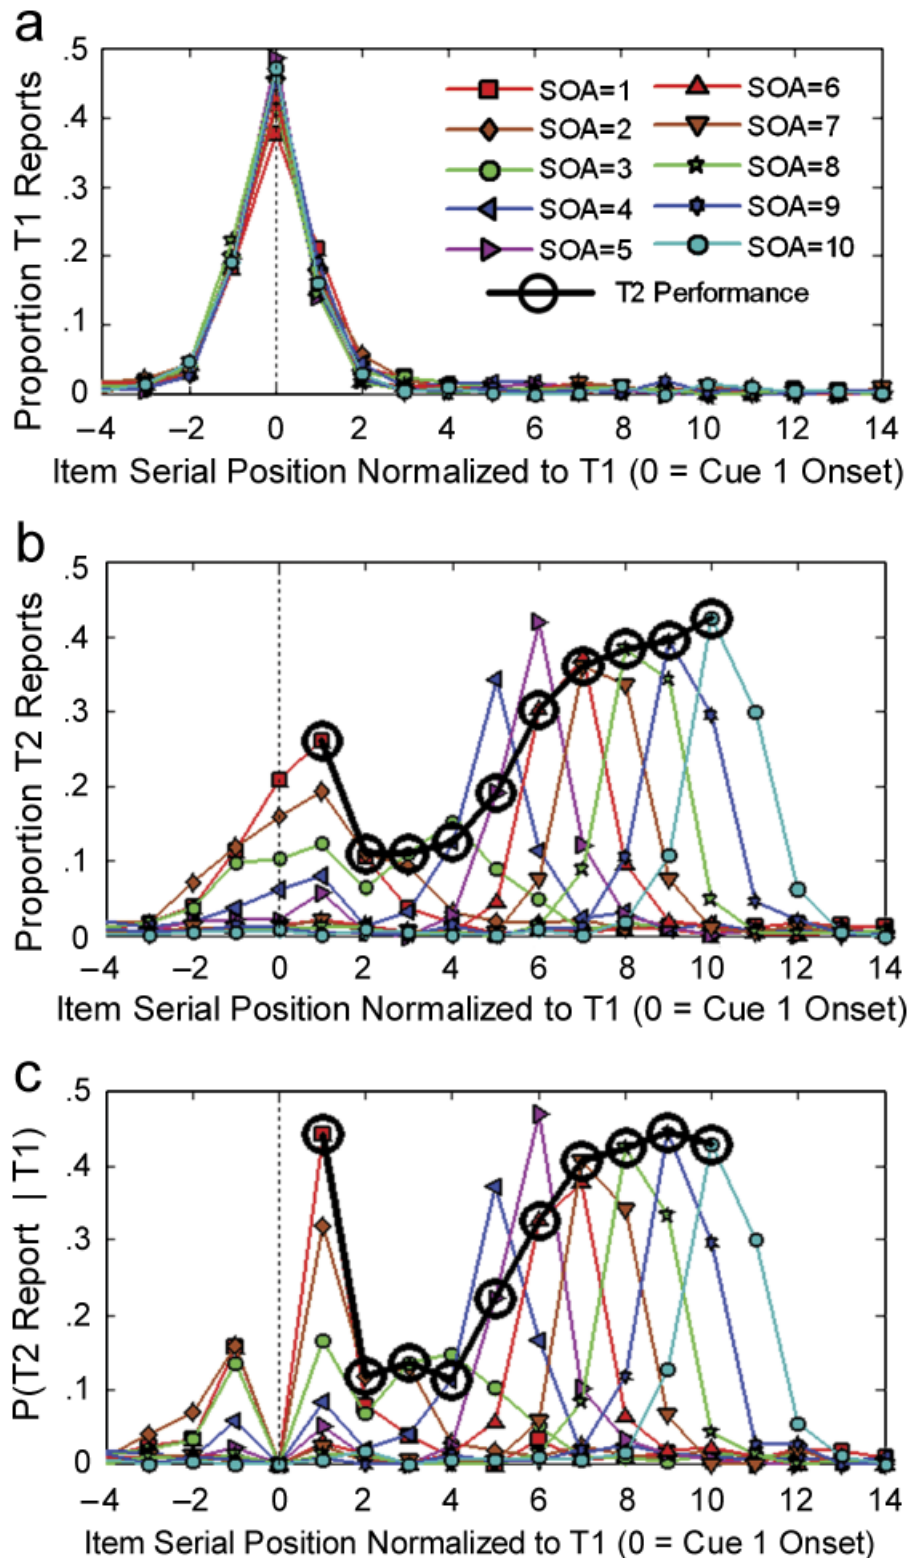


Fig. 3. Probability of reporting a letter from a particular position in the rapid serial visual presentation stream. Results are shown separately for (a) the first target (T1), (b) the second target (T2), and (c) T2 given that T1 was reported correctly. The thick black lines in (b) and (c) indicate the probability of accurate T2 report as a function of stimulus onset asynchrony (SOA). Item position is normalized to T1 (i.e., T1 is always at $x = 0$).

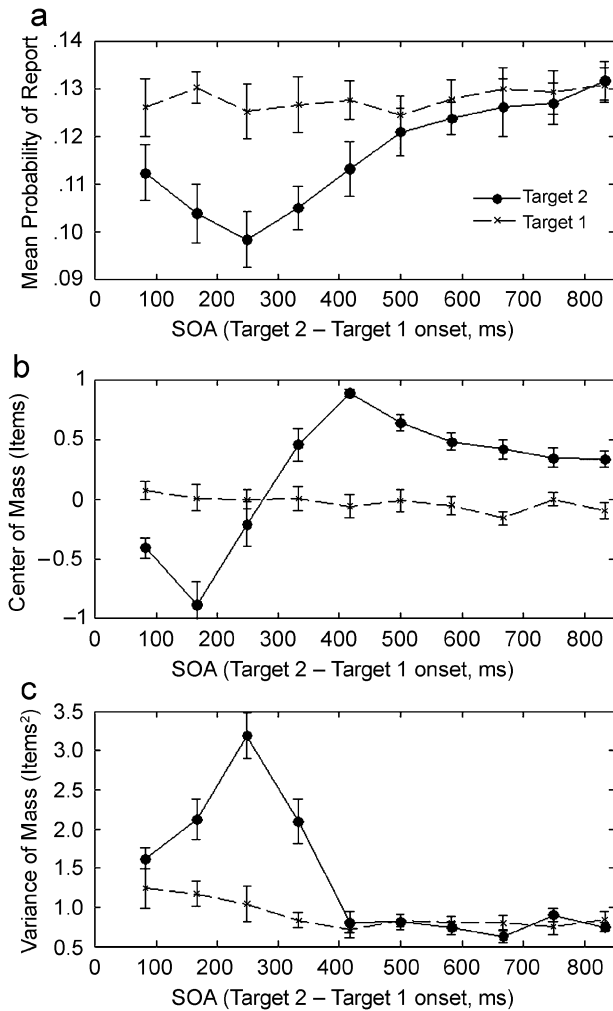


Fig. 4. Results for the three selection-window measures: (a) average probability of reporting an item within the seven-item selection window around the target (suppression), (b) center of mass of reports in the window around the target (delay), and (c) variance of the center of mass in the window around the target (diffusion). Each measure is shown as a function of stimulus onset asynchrony (SOA) between the two targets. All error bars correspond to 1 SEM.

during the AB, not only T2 itself, but also other items around T2, are less likely to be seen and reported. The AB is not caused simply by elevated levels of posttarget intrusions, because even those intrusions are less frequent during the AB than at SOAs outside the range of the AB. Instead, selection is less effective during the blink. An attempt to select letters from the RSVP stream yields fewer educated guesses: Selection is suppressed, and subjects resort to greater levels of random guessing.

The second relevant metric of temporal selection is its latency, which we estimated as the center of mass of reports in the window around a given target. This measure (originally employed by Chun, 1997a) corresponds to the average reported serial position, relative to the target (see the appendix), and is independent of both the efficacy and the precision of selection. Nonzero values of the center of mass correspond to systematic inaccuracies of selection: A positive center of mass means subjects

were more likely to report items that followed the target than items preceding the target, whereas a negative center of mass means subjects were more likely to report items preceding the target than those following the target. If the center of mass is more positive for T2 than T1, this indicates that selection is delayed for T2, relative to T1.

Figure 4b plots this measure of delay for T1 and T2 reports as functions of SOA. At short SOAs (83, 167 ms), the center of mass for T2 reports was negative, indicating that items before the target were reported more often than items after the target; this may have occurred because items between T1 and T2 benefit from the imprecise selection of both targets and are thus reported more often than other items. The negative center of mass for T2 reports at short SOAs became a delay (relative to the center of mass for T1 reports) at SOAs greater than 250 ms. This result indicates a substantial shift to a pattern of posttarget intrusions. The delay of selection persisted longer than the suppression of selection, as the measure of delay did not return to T1 levels in any of the SOAs we investigated (up to 833 ms). At the long SOAs, we found a slowly decaying exponential trend, extrapolation of which suggests that subtle traces of delay may last as long as 2 s. Given that items were separated by 83.3 ms, we were able to compute the delay of the center of mass in time. These calculations revealed that selection was delayed by roughly 75 ms at an SOA of 417 ms and was still delayed by about 30 ms at the longest SOA we measured (833 ms).

A third metric of selection is the precision of selection around the center of mass. This can be measured by the variance of the center of mass of reported items (see the appendix). Figure 4c displays this measure as a function of SOA for T1 and T2 report. The time course of the variance in the position of T2 report was different from that for both delay and suppression: Variance was much greater for T2 than for T1 at SOAs of 167, 250, and 333 ms, but this difference was gone by 417 ms. Because in the limit, suppression and diffusion yield identical distributions of report, we checked that our measure of diffusion does not merely reflect a side effect of suppression. By estimating the level of random guessing (i.e., the probability of reporting items from serial positions distant from the target), we could correct for the effects of suppression; this correction resulted in no qualitative changes to the reported results.

To assess whether the three measures of selection follow different time courses, we first conducted an analysis of variance with the factors of SOA (1–10 items), measure (efficacy, latency, or precision), and target (1 or 2). All effects and interactions were significant at a p value of .0001; most important, the three-way interaction was highly significant, $F(18, 660) = 16.31, p < .00001$. Of course, the three-way linear interaction is not the most rigorous test of the claim that the three effects follow different time courses, because relationships between measures are not expected to be linear. The essence of our hypothesis is that each of the three measures takes a different amount of time (in terms of SOA) to return to T1 levels. Thus, we could rigor-

ously test this claim with pair-wise comparisons of the time courses of the three effects. First, suppression and diffusion recovered to T1 levels long before the maximal SOA (833 ms); however, delay was still significantly greater for T2 than for T1 at this lag, $t(11) = 5.98, p < .0001$. Second, we tested the less obvious time-course difference between suppression and diffusion. To test the claim that diffusion recovered to T1 levels earlier than suppression, for each subject we found the first SOA greater than 1 item at which each measure did not differ between T2 and T1. We then ran a t test on the values of these SOAs across subjects. Just as the graphs suggest, this statistic showed a significant difference between the two measures: Diffusion recovered earlier than suppression, $t(11) = 2.7, p < .05$.

Thus, all three measures of selection—corresponding to efficacy, latency, and precision—change when two selection episodes interact. Under normal AB conditions, selection is delayed, diffused, and suppressed (see Fig. 1). Each of these effects alone would diminish the probability of reporting T2; however, as normally measured, the AB reflects all three effects. Furthermore, given that each of these effects follows a different time course, they are likely to emerge from separate processes, which may be independently affected by manipulations known to alter T2 report.

DISCUSSION

We analyzed the pattern of distractor intrusions around two targets from an RSVP stream and showed that the occurrence of a first selection episode causes suppression, delay, and diffusion of selection during a second episode. These three effects are dissociable by their time courses. Our results indicate that although the three effects have a common cause (proximity of two temporal-selection episodes), they reflect different processes that combine to form the commonly reported AB.

The delay effect we obtained replicates Chun's (1997a) results, showing that during the AB, intrusions from posttarget items are elevated. At the faster RSVP rate we employed, this effect was amplified, which indicates that the delay of selection is a function of time, rather than number of items. Because our faster RSVP rate provided a more sensitive measure of the distribution of intrusions, we could compare the time courses of the delay, suppression, and diffusion of selection. This analysis revealed that the delay effect disappears more slowly than does the suppression effect, which persists longer than the diffusion effect.

These findings cannot be accounted for in terms of the mechanisms currently proposed for the AB. In some circumstances, suppression may cause a delay in the center of mass, but suppression alone cannot account for the finding that delay and diffusion follow different time courses. One possible way to account for these data would be to assume that suppression has an effect at multiple levels of the visual system (Visser, Williams, Cunningham, & Mattingley, 2005). Such a mechanism of suppression, accom-

panied by a particular pattern of processing delays between visual areas might result in the particular time courses of delay and diffusion we observed. A key goal for future research will be to determine if a central modulatory mechanism with a single time course could result in the three different time courses we observed by affecting multiple stages of the visual hierarchy, or if multiple mechanisms must be invoked to account for these data.

More broadly, these data inform current understanding of temporal selection in general: Proposed models of selection in time (Reeves & Sperling, 1986; Shih & Sperling, 2002) explain selection and patterns of reports in a linear-system framework. However, our data show massive nonlinearities (namely, delay and diffusion) when two selection episodes interact. In the linear-system framework, these changes to the profile of temporal selection can be described as changes to the attentional-gate impulse response function as a function of SOA. Given that such a dynamic-systems account of vision seems to be generally promising, another key step for future research will be to modify the linear model to account for these temporal nonlinearities.

In sum, our results suggest that two proximal episodes of temporal selection interact nonlinearly such that the second selection episode is rendered (a) less effective (suppressed), (b) systematically delayed, and (c) less precise (diffused). We showed that these three effects are dissociable by their time courses. These findings pose a challenge to theories of the AB. Further, our results provide new insights about the fundamental nature of temporal selection by showing that it varies along at least three independent dimensions: efficacy, latency, and precision.

Acknowledgments—This work was supported by National Eye Institute Grant 13455 to N. Kanwisher. We thank Adam Love and Brian Coffey for assistance running the experiment.

REFERENCES

- Botella, J., Arend, I., & Suero, M. (2004). Illusory conjunctions in the time domain and the resulting time course of the attentional blink. *Spanish Journal of Psychology*, *7*, 63–68.
- Botella, J., Garcia, M.L., & Barriopedro, M. (1992). Intrusion patterns in rapid serial visual presentation tasks with two response dimensions. *Perception & Psychophysics*, *52*, 547–552.
- Bowman, H., & Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory. *Psychological Review*, *114*, 38–70.
- Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.
- Chun, M.M. (1997a). Temporal binding errors are redistributed by the attentional blink. *Perception & Psychophysics*, *59*, 1191–1199.
- Chun, M.M. (1997b). Types and tokens in visual processing: A double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 738–755.

- 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., Kawahara, J., Shahab Ghorashi, S.M., & Enns, J.T. (2005). The attentional blink: Resource depletion or temporary loss of control? *Psychological Research*, *69*, 191–200.
- Kanwisher, N.G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, *27*, 117–143.
- Nieuwenstein, M.R., Chun, M.M., van der Lubbe, R.H., & Hooge, I.T. (2005). Delayed attentional engagement in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 1463–1475.
- Nieuwenstein, M.R., & Potter, M.C. (2006). Temporal limits of selection and memory encoding: A comparison of whole versus partial report in rapid serial visual presentation. *Psychological Science*, *17*, 471–475.
- Olivers, C.N., van der Stigchel, S., & Hulleman, J. (2007). Spreading the sparing: Against a limited-capacity account of the attentional blink. *Psychological Research*, *71*, 126–139.
- 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.
- Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review*, *93*, 180–206.
- Shih, S.I., & Sperling, G. (2002). Measuring and modeling the trajectory of visual spatial attention. *Psychological Review*, *109*, 260–305.
- Visser, T.A., Williams, M.A., Cunnington, R., & Mattingley, J.B. (2005, May). *Activation of visual areas during the attentional blink: An fMRI investigation*. Paper presented at the annual meeting of the Cognitive Neuroscience Society, New York.

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

(RECEIVED 3/2/07; REVISION ACCEPTED 6/9/07)

APPENDIX: EQUATIONS USED

We calculated A , the average probability of report (efficacy of selection: our measure of suppression), as follows:

$$A = \frac{\sum_{i=ks}^{ke} P_i}{n},$$

where n is the total number of items in the selection window (in our case, 7), P_i is the probability (empirical frequency) of reporting an item from serial position i (relative to the target position, 0), and ks and ke are the lower and upper bounds, respectively, of the window used to compute the measure, expressed in serial position of the item relative to the cue position. We used $ks = -3$ and $ke = 3$ (0 is the correct target).

C , the center of mass (latency of selection: our measure of delay), was calculated as

$$C = \frac{\sum_{i=ks}^{ke} P_i * i}{A * n}.$$

Finally, V , the variance of the center of mass (precision of selection: our measure of diffusion), was calculated as

$$V = \frac{\sum_{i=ks}^{ke} P_i * (i - C)^2}{A * n}.$$