Journal of Experimental Psychology:
Human Perception and Performance
55 Vol. 21, No. 2, 275-292

0096-1523/95/$3.00

Repetition Blindness: Perception or Memory Failure?

Clark Fagot and Harold Pashler
University of California, San Diego

Repetition blindness (RB) is the failure to report a repeated item when reporting all of the elements from a rapidly displayed sequence of visually presented items. N. Kanwisher (1987) has characterized RB as a perceptual problem. Experiments 1 and 4 manipulated the order of report of the repeated items. Experiments 2-4 employed tasks other than full report that should still have shown RB according to a perceptual hypothesis, but none was found. In Experiment 5, RB was found when one of the repetitions was spoken and the other was presented visually; this RB could be localized to retrieval processes. RB appears to reflect processes and strategies peculiar to full report from rapid displays rather than any fundamental perceptual limitation and might be the same as the Ranschburg effect found in immediate recall tasks.

When people try to report what items were presented in a rapidly displayed sequence of letters or words, commonly termed a *rapid serial visual presentation (RSVP) display,* they tend to make an interesting error when the display includes an item that is repeated. This error is a failure to report the second occurrence of the repeated letter or word (Kanwisher, 1987). Kanwisher has termed this effect *repetition blindness* (RB), and she suggests that it stems from a difficulty in perceiving the second occurrence. The idea that there may be some specific difficulty in registering multiple instances of a single category has been suggested in other contexts as well (Bjork & Murray, 1977; Humphreys, Besner, & Quinlan, 1988; Mozer, 1989; Pashler & Badgio, 1985,1987). The present article presents data that challenge this interpretation of RB.

Consider Kanwisher's (1986,1987) account of RB, which will be referred to as the *types and tokens* hypothesis, in some detail: For an object to be perceived, two processes must operate successfully. First, the *type* (category) into which the item falls must be recognized. When this happens, a "type node" is activated, representing the presence of a token of this type somewhere in the visual field. The second requirement is *token individuation.* This involves "binding" the type with the appropriate token, where the token might stand for, say, the third element in the RSVP display. Thus, if the word *house* occurs in the third serial position in the RSVP display, then to perceive the item, the type node for *house* must be activated (type recognition) and the token for the third serial position must be bound to this type node. If type activation occurs but token individuation does not, the observer does not become aware of the item as a separate instance (hence the blindness in repetition blindness), although he or she may be aware that something of that type occurred somewhere in the display (Kanwisher, 1987). Kanwisher views both processes as part of perception, because "detection" is claimed to fail when either process fails (Kanwisher, 1987, 1991; Kanwisher & Potter, 1989, 1990).

The types and tokens model provides one potential explanation of the basic RB effect. However, the model makes several claims that have not been explicitly tested. First, the types and tokens model is claimed to be a "model of visual information processing" (Kanwisher, 1987, p. 133), not simply visual information processing during full report from RSVP displays. Thus, it holds that token individuation is required every time the observer sees the stimulus as a separate item, not only for the full-report task normally used in RB experiments but also under all other conditions. This is at the heart of what it means to be perceptual. One important implication is that any task requiring observers to see the repeated items as distinct items should show RB. If this is not the case, there is something seriously wrong with the model.

Another disputable claim concerns the actual locus of the effect. Strictly speaking, what is observed in RB is, of course, not blindness but the failure to report an item. The fact that people cannot report something does not imply that they did not perceive it. If people are presented with a list of 500 words, successively exposed for 1 s each, they will not be able to later report more than a fraction of these words. Such a problem is not due to an inability to identify all 500 words, a fact which could be demonstrated by having them read aloud each of the words as it is presented. Even after doing so, observers still cannot report them all. Obviously,

Clark Fagot and Harold Pashler, Department of Psychology, University of California, San Diego.

This work was supported by Office of Naval Research Grant N00014-91-J-1735, a National Science Foundation Graduate Fellowship, and National Institute of Mental Health Training Grant 2T32-MH14268-18A. We are grateful to John Duncan, Robert Greene, Larry Hochhaus, and Nancy Kanwisher for their comments, and especially to Jim Johnston for his thorough and highly insightful commentary. We are also grateful to Alarm Lopes for technical assistance and to Bill Nusbaum, Toan Vuong, Shannon Wilcox, and Deborah Woo for assistance in testing participants.

Correspondence concerning this article should be addressed to Harold Pashler, Department of Psychology, 0109, University of California, La Jolla, California 92093-0109. Electronic mail may be sent via Internet to hpashler@ucsd.edu.

the problem here is that they cannot store and retrieve this many words in memory. One of the basic discoveries in Sperling's (1960) classic studies was that even "immediate" report critically depends on short-term memory storage and retrieval. It seems reasonable to us to suppose that what is observed in RB might not be a failure, broadly speaking, of "perception" but instead of "memory".[1] Clearly distinguishing what is perceptual and what is memorial is not an easy task. However, if RB were due to a failure of off-line processes (ones occurring after the stimuli have been presented and during retrieval or report) rather than on-line processes (ones occurring as the stimuli are presented), it would seem more appropriate to characterize the problem as involving memory and not perception. Furthermore, if RB were an on-line effect, it still might not be best to characterize the problem as one of perception. In the next section, several potential on-line loci of RB are delineated more specifically.

Possible On-Line Loci of Repetition Blindness

Because it has been claimed that RB is an on-line effect, it is worth considering possible on-line loci in a fairly explicit fashion. In order to make things concrete, suppose an observer sees the RSVP display of letters *F S K V K L,* in which the letter *K* appears twice. RB refers to the observer's tendency to report only the first *K*. Why might the observer be unable to report the second *K*? Some possible explanations follow:

1. Recognition failure: The fact that a *K* was previously recognized makes the system less able to recognize that a *K* is present when the second one occurs.
2. Individuation failure: The fact that the *K* is a repetition does not affect the determination that a *K* is present, but it makes it more difficult to determine the presence of this particular token of a *K* (including attributes specific to the token, such as position in space and time).
3. Storage failure: When the second *K* is presented, the fact it is a repetition has no effect on the observer's immediate awareness of this particular token of a *K,* but it makes it more difficult to store in short-term memory the fact that a new *K* has occurred (necessary for delayed report).[2]

The first question that needs to be addressed is whether the effect is on-line or off-line. We have sketched out more detailed on-line loci to show that even if the effect is on-line, there are more alternatives than the types and tokens model (individuation failure). Kanwisher (1987) has acknowledged that recognition failure and individuation failure are different. We emphasize the further distinction between individuation failure and storage failure, a distinction Kanwisher (1987) argued against.

The Ranschburg Effect

Memory researchers have known for some time that people are less accurate at reporting the second occurrence of a repeated item in full report of a series of items (Crowder, 1968; Crowder & Melton, 1965; Jahnke, 1969). In the memory literature, this effect is called the *Ranschburg effect.* Like RB, the Ranschburg effect is defined as a tendency to drop a repeated item more frequently than a nonrepeated item. Whereas researchers investigating RB use rapid visual presentation (typically 8 items/s) and (often) words that form sentences or letters that form words, researchers investigating the Ranschburg effect have generally used slower presentation (2 items/s) and unstructured material in either the visual or auditory modalities. Because the Ranschburg effect is found with such slow presentations, it has been universally assumed to reflect memory rather than perception, broadly speaking,[3] and furthermore, various experiments have identified it as a memory output rather than input problem (Crowder, 1968; Greene, 1991). Despite the apparent similarity of RB and the Ranschburg effect, Kanwisher (1987) has argued that the two are fundamentally different. Appendix A examines these arguments and concludes that the two effects are not necessarily qualitatively different.

Dependence of RB on Stimulus Input Timing

It has been claimed that "RB for R2 is observed only when … the (filled) ISI [interstimulus interval] is no more than about 500 ms, and when R2 is presented for less than 200 ms" (Kanwisher & Potter, 1990, p. 42), where R2 is the second instance of the repeated item. The claim that RB critically depends on presentation timing has been widely accepted. It would also seem to support a perceptual theory of RB. However, this dependence has not been observed with the full-report methods used in virtually all RB experiments. Instead, it was found in a single experiment in

---

[1] In later sections we will be considering distinctions more fine-grained than "memory" and "perception." We simply use these terms in a commonsense way without claiming that all processes fit cleanly under either term.

[2] The types and tokens model could be couched in terms of storage processes, and, in fact, Kanwisher (1987) has noted this. However, this would be a very specific storage model, in that the storage operation that fails in RB is required every time the observer must correctly perceive an item as a separate token. Other storage models, on the other hand, might hold that the offending storage process occurs under some perceptual conditions and not under others (e.g., when the observer is set for full report but not when he or she is set for report of only a few items). In fact, there is no loss of generality if we classify such a model as an individuation failure model (the only difference between it and other individuation models is that in this model, observers might be aware of the item as a separate instance as it is presented, but they are unable to store this instantaneous awareness).

[3] Ironically, Ranschburg's original experiments (Ranschburg, 1902) were more similar in many ways to Kanwisher's (1987) than more recent "Ranschburg experiments." For example, in one experiment, Ranschburg presented observers with 6 digits simultaneously for *1/3* s. His explanations of the phenomenon, in fact, were also closer to Kanwisher's than to other researchers, emphasizing attentional and perceptual factors, not memory (according to Jahnke, 1969).

which the observer's task was to decide whether an RSVP display contained any repeated item (Kanwisher, 1987, Experiment 1). Kanwisher found that at 250 ms/item, observers were 92% accurate in deciding whether the display contained a repetition. If it is assumed that failures in this task reflect perceptual RB, then it would follow that at 250 ms/item there is very little RB. However, there is no reason to believe that errors in this task reflect problems perceiving repeated items, even at the faster rates of presentation. Rather, the need to compare each item to all the items before it may determine the limiting presentation rate at which repetitions can be reliably detected. Thus, the fact that observers are accurate at detecting repetitions with a presentation rate of 250 ms/item cannot be used to argue that the effect is perceptual, only that if it were a perceptual effect, this effect disappears at presentation rates slower than this. What about RB in the standard full-report task? No systematic examination of RB as a function of presentation rate has been reported. However, the effect has been found with presentation rates as slow as 1 s/item (the Ranschburg effect), albeit with larger memory loads. In summary, stimulus timing factors do not provide evidence for perceptual or on-line RB.

Present Investigations

The following experiments address two issues. The first is the locus of the RB effect. Our main focus is on whether the effect has an on-line or off-line locus. The basic method we use to try to distinguish between on-line and off-line loci is to manipulate the order of retrieval of the critical items (the critical items are the repeated items for stimuli that have repetitions, and the items in the corresponding locations for control stimuli). On-line models predict that RB should affect the second perceived item whereas off-line models predict that the second retrieved item should suffer. Experiments 1 and 4 address this question.

The second issue tests the types and tokens model more directly. As already discussed, this model holds that RB should occur whenever the repeated items must be seen as separate instances or tokens. Thus, in Experiments 2-4 we test this by using tasks that do not require report of the whole RSVP display but that nonetheless require the critical items to be seen individually in order to perform the task correctly.

Experiment 1

In Experiment 1 we asked whether on- or off-line processes are the source of RB. Study participants reported the letters from moving RSVP displays. In these displays, each succeeding letter was presented shifted to the left or right from the position of the preceding one, and the preceding letter was simultaneously masked. Such displays have previously been found to produce RB (Kanwisher & Potter, 1989). One group of participants reported letters from a rightward moving RSVP display and reported the items from left to right; the other group reported letters from a leftward moving RSVP display but still reported the items from left to right. If RB arises during off-line processes, then participants should have dropped the second reported—rather than the second presented—item in both the leftward and rightward moving displays. On-line accounts, on the other hand, predict that the second presented item will be dropped.

Method

Participants. Twenty-four students participated in this experiment for partial fulfillment of a course requirement at the University of California, San Diego.

Apparatus and stimuli. Stimuli were presented on NEC Multisync monitors connected to IBM PC compatible microcomputers. The stimuli were strings of six randomly generated letters from the set *A* through *J*. There were repetition strings with exactly one letter repeated and nonrepetition (control) strings with no letters repeated. The letters were 0.9 cm wide by 2.2 cm tall from a typical viewing distance of 60 cm.

Strings contained all different letters (nonrepetition strings) or exactly one pair of repeated letters (repetition strings). No repeated item occurred in the first or last position or in adjacent positions. Thus, there were three types of repetition stimuli: 2-4 (repeated items in the 2nd and 4th position from the left or order of report), 2-5, and 3-5.

Design. The experiment contained five blocks of 48 trials and one practice block. Twelve participants were in each of the forward presentation (rightward moving RSVP displays) and backward presentation (leftward moving RSVP displays) conditions.

There were 24 repetition and 24 nonrepetition stimuli in each block. Of the repetition stimuli, there were 8 each of the 2-4, 2-5, and 3-5 stimuli (each nonrepetition stimulus served as a control for all three of the repetition stimulus types). Stimulus order was completely randomized.

Procedure. Participants received written instructions describing their task. The instructions informed them that sometimes some letters might occur twice in a display.

Each trial began with the presentation of a white plus sign (+) in the center of the screen for 1,000 ms. Then, 250 ms after the offset of the fixation point, the letters of the stimulus string were presented at the fixation point at a rate of 133 ms/letter (the offset of each letter coinciding with the onset of the next one, except, of course, for the last letter, whose offset coincided with the onset of the final mask). Each frame was offset to the right (forward condition) or left (backward condition) by 1.4 cm, and the previous frame was simultaneously masked by a pound sign (#). (Each mask was removed in the next frame, except for the last mask, which was removed 150 ms after it was presented.)

After this, a report cue of six underline characters (_) was presented to the participant with each underline in a position that had previously been occupied by one of the stimulus letters. In both the forward and backward presentation conditions, participants entered the letter that occupied that position onto each underline character, from left to right. Thus, in the forward condition, observers reported the letters in the same order as they were presented; in the backward condition, they reported the letters in the opposite order. The letters were reported by pressing keys labeled with the letters *A* through *J*. The keys were arranged on the home row in alphabetical order from left to right. A position could be left blank by pressing the space key, and the backspace key could be used to edit entries. To go on to the next trial, a key labeled DONE was pressed. If more than three items were reported in the correct or an adjacent position, then a pleasing ascending sound occurred (otherwise, no sound occurred). The next trial began 250 ms after the DONE key was pressed.

It is possible that the procedure used for report in this experiment might have introduced more memory problems than, for example, vocal report. To guard against this possibility, participants were given a pad of paper and two pencils in case they found it easier to write down their entries first. They were also told that it would be easier if they said their answer to themselves before typing it in. No participant relied on the paper and pencil; all found it easier to either type their answer immediately or say it to themselves before typing it in.

Between blocks, the proportion of letters reported in the correct or adjacent positions was presented to participants, broken down by block number. Participants pressed a key when they were ready to continue.

Results

Repetition blindness was assessed in two ways. First, accuracy was calculated for each serial position in the report, a letter being recorded as correct only if it was reported in the correct location. Repetition blindness, using this measure, corresponds to an effect of repetition on one or both of the critical positions, that is, position 2 and 4 for 2-4 stimuli, 2 and 5 for 2-5 stimuli, and 3 and 5 for 3-5 stimuli. The problem with this analysis is that if participants often report the right letters in the wrong order, any effect will be washed out. Thus, three additional measures were computed. The *both score,* the proportion of trials in which both critical items were included in report, regardless of location in report, was the first. This is a very sensitive measure of RB; any RB should show up as a negative effect of repetition on the Both score. The other two measures used were *C*1 *score,* the proportion of trials in which the first presented critical letter was reported among the first three letters reported (last three for backward report), and *C2 score,* the proportion of trials in which the second presented critical letter was reported among the last three letters reported (first three for backward report). These two measures gave an indication of which instance of the repeated items was being dropped. (Notice, however, that if the order that participants reported letters were random, the C1 and C2 scores would not be sensitive to which item was dropped; it is debatable whether the question of which item is dropped is even well defined in this case.) Analysis for each of these measures was performed on the mean of the scores for 2-4, 2-5, and 3-5 stimuli for repetition trials, and the average of the scores for control trials using 2 and 4, 2 and 5, and 3 and 5 as the critical positions. Thus, each control stimulus served equally as 2-4, 2-5, and 3-5 control stimuli.

Figure 1 shows the proportion of letters reported in the correct location in the forward (Figure 1A) and backward (Figure 1B) presentation conditions, broken down by stimulus type (control or repetition type 2-4, 2-5, or 3-5) and serial position (in the order of report). For both order conditions, report is best at the start of report and worst in the middle positions. In the backward condition, participants report the third and fourth positions correctly at a rate only a little above chance (10%).

The effect of repetition was analyzed for each combination of stimulus type and presentation direction at each serial position. There were only two significant effects: Repetition had a positive effect in the first serial position for 3-5 stimuli in the forward condition, $F(1, 11) = 9.1$, $p < .05$; repetition had a negative effect in the fourth position for 2-4 stimuli in forward report, $F(1, 11) = 6.4$, $p < .05$. Thus, there is little evidence of RB in the strict location

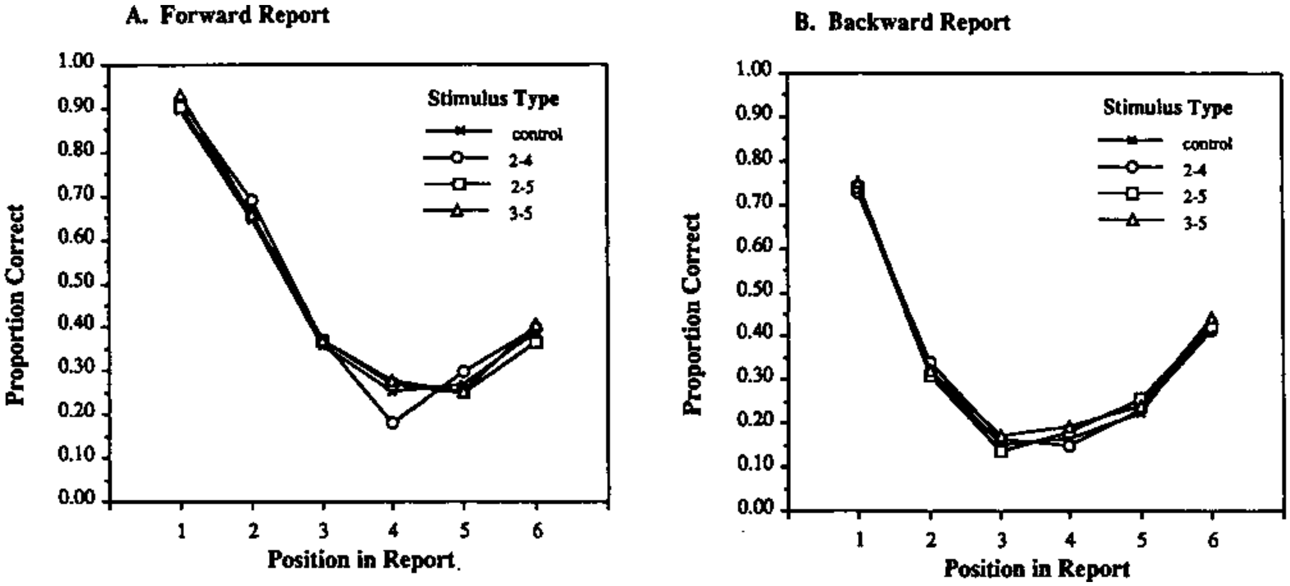

Figure 1. Proportion of letters reported in the correct position for (A) the forward report and (B) the backward report as a function of serial position and types of stimulus for Experiment 1.

reports (significant effect in only 1 of the 12 critical positions). This is probably because participants were not able to accurately report items in the correct location.

Table 1 shows the Both, C1, and C2 scores for each combination of forward and backward report by repetition. There was a 2.8% beneficial effect of repetition in forward report on the C1 score, $F(l, 11) = 7.4$, $p < .05$, an 8.5% negative effect on the C2 score, $F(l, 11) = 12$, $p < .01$, and a 22.9% negative effect on the Both score, $F(l, 11) = 96$, $p < .001$. There was no effect of repetition on the C1 score in the backward condition, $F < 1$, but there was a significant 3.7% benefit of repetition on the C2 score, $F(l, 11) = 5.5$, $p < .05$. Even though no significant harmful effect of repetition was found on the C1 or C2 score, repetition was harmful on the Both score, $F(l, 11) = 43$, $p < .001$ (20.5% effect).

Discussion

RB occurred in both the forward and backward presentation conditions in this experiment. There was a 22.9% harmful effect of repetition on the Both score in the forward presentation condition and a 20.5% effect in the backward presentation condition. The Both score is the standard measure used in previous RB studies: A repetition deficit on this measure plainly indicates trouble reporting both of two repeated items.

The on-line and off-line models make specific predictions for which critical item is dropped. On-line models predict that the second item seen is dropped, whereas off-line models predict that the second retrieved is dropped. In forward report, where both models predict that the second item should be dropped, there was no significant effect of repetition on the C1 score, but there was an 8.5% negative effect of repetition on the C2 score. In backward report the models make different predictions, but there was no harmful effect of repetition on either C1 or C2. This was not due, however, to a lack of RB, since RB was found on the Both score.

Making sense of the details. There is no obvious way to infallibly determine the location of the RB effect. The C1 and C2 measures have the advantage of simplicity but also have some potential problems. These problems might be the reason that no RB is found on the local measures in the backward condition.

Table 1

Proportion of Trials in Which Participants Include the First (C1), Second (C2), or Both Critical Items in Their Report in Experiment 1

Critical item	Presentation condition Forward	 Backward
C1		
Repetition	70.6	39.6
Control	67.8	38.2
C2		
Repetition	42.1	46.0
Control	50.6	42.3
Both		
Repetition	25.3	10.2
Control	48.2	30.7

One problem with the local measures is that they can underestimate the report of nonrepeated critical items. This is because a nonrepeated item found in the wrong half of the report is marked as wrong on both the C1 and C2 scores. For repeated items this problem is avoided because the critical items are identical: an item that migrates in report will never be counted wrong.

A second problem can be seen by considering the following example. Suppose the stimulus is *ABCDBF,* a 2-5 repetition stimulus where *B* is the repeated item. On some trials the participant might report *ADBF,* leaving the last two positions blank. Using the local measures, this case would be recorded as missing the second repetition (in the order of report), even though, from another perspective, it is the first one that is missing. If such a situation arose frequently, it would create a bias toward finding the repetition effect on the second reported critical item.

Both of these problems can be corrected for by computing the local measures in a slightly different way. Here, only those trials in which report includes the one or two items between the critical items in the stimulus are counted. In the case where there are two items between the critical items (2-5 stimuli), we also require that they are consecutive and in order in the report. We call this string of one or two letters the *border string.* (The reason for all this will be clear below). Both critical items are counted correct if they both occur somewhere in report. If only one is reported, then for nonrepeated stimuli, we count the critical item which has the same letter name as the reported item as correct. For repeated stimuli, we count the critical item that lies on the same side of the border string as the reported item as correct. These measures avoid the problem of underestimating the report of critical items in nonrepetition stimuli by always giving credit for reporting a critical item. They also avoid the bias toward finding RB on the second reported critical item because critical items are localized using the border string rather than absolute position for nonrepetition stimuli.

Unfortunately, the modified measures do not satisfactorily determine which critical item was dropped. In forward report there was an 11.7% effect of repetition on C1 and a 21.0% effect on C2. Thus, in forward report, RB is again concentrated on C2. In backward report, however, there was a 10.4% effect on C1 and a 7.6% effect on C2. We are again left unable to determine which critical item was dropped in backward report.

The report matrix in Figure 2 demonstrates one possible reason that the RB effect could not be localized in the backward condition. This figure shows the proportion of trials in which the letter in a given position of the stimulus is found in a given location in report (control trials only), for both forward (Figure 2A) and backward (Figure 2B) conditions. Notice that in the forward report condition, observers are much more likely to report stimulus letter *i* in report position *i* than in positions *i* ± 1 (and further drop-offs are

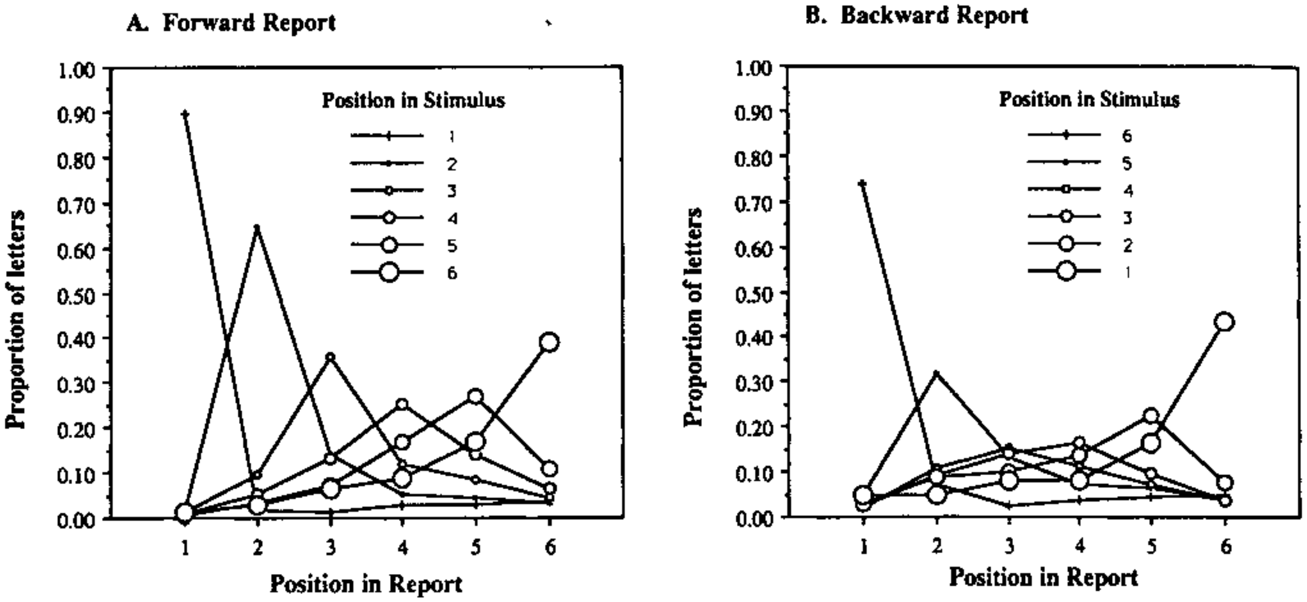

Figure 2. Proportion of letters that occurred in the given stimulus position that were reported in (A) the forward report and (B) the backward report for Experiment 1.

found at positions *i* ± 2, and so on). Thus, participants are doing a reasonably good job at placing items in the correct or almost correct location, when they are able to report them correctly at all. Consequently, the C1 and C2 scores reflect which critical item was dropped. However, the same is not true in the backward condition. Stimulus Letter 3, for example, was reported in Position 4 almost as often as it was reported in Position 3. Thus, it is questionable whether it is logically possible to determine whether C1 or C2 is dropped in backward report.

To assess whether this experiment has the power to distinguish between which critical item was dropped in backward report, we computed the expected C1 and C2 scores (Table 1) and strict location report accuracy (Figure 1), based on a model in which C1 is dropped in RB and on a model in which C2 is dropped. The models took advantage of the fact that on nonrepetition trials, the critical items could easily be distinguished by identity, no matter where they occurred in report. Thus, we used nonrepetition data to model repetition data, removing a critical item on a proportion of trials determined by the size of RB for that observer. It is interesting to note that in order to get close fit to the data, it was necessary to assume that when a critical item was dropped from report, the following items were shifted one position to fill in the gap. In forward report, for both strict location accuracy and C1 and C2 scores, the data were well fit by a model in which C2 was dropped but poorly fit if C1 was dropped. In the backward condition, however, both models fit the data reasonably well. Thus, whereas these analyses support the conclusion that it is C2 that is dropped in forward report, they also suggest that location information is too poor to determine which critical item is dropped in backward report.

All that notwithstanding, report factors dominate perceptual factors in this experiment. If observers simply saw the display and then reported what they saw, we would not expect report variables to have an effect on their performance. However, which critical item is more accurately reported is determined by whether it is the first reported, not whether it is the first seen (Table 1). Similarly, letters early in the report sequence are more accurately reported than letters occurring later, even in backward report (Figure 1). This is important to note because if report factors play a large role in determining performance, they remain a viable candidate for the source of RB.

Summary. This experiment shows that KB is found in both forward and backward report. However, it also demonstrates the difficulty in localizing the effect. Although the data support the prediction that C2 is dropped in forward report, they do not indicate which item was dropped in backward report. However, this experiment does refute the idea that performance on the task mostly reflects how well items in the display were perceptually processed; otherwise, order of report should not have played such a large role, and RB should have been localized to C1 in backward report. A different way of manipulating the order of report of the critical items might yield more decisive results. This is attempted in Experiment 4. In Experiments 2 and 3 we test whether RB really occurs whenever observers must process the critical items as separate tokens.

Experiment 2

If perception of a repeated stimulus is impaired in the way Kanwisher (1987) proposed, then there should be difficulty in detecting repetition per se. Kanwisher (1987) reported v that observers do have difficulty in detecting repetitions. She instructed participants in her study to report whether any items were repeated in an RSVP word stream, under conditions in which they did not know which particular words might be repeated. Performance was quite poor, with many repetitions missed except at the very longest exposure durations.

The hypothesis of perceptual RB predicts that participants should make many errors in her experiment: If the individuation of the repeated item is impaired, then detection of the fact that there is a repetition must also be impaired. Thus, her results are certainly consistent with the hypothesis of perceptual RB. However, such results might still arise even if perception of a repeated stimulus is not impaired. Detecting the presence of (unspecified) repeated items logically requires observers not only to recognize each occurrence of the repeated item but also to remember every preceding item in the list and compare the current item with all the preceding items (Kanwisher, 1987). Thus, item k must be compared to $k - 1$ items. Even if there were no difficulty in perceiving the repeated item, the requirements to retain and compare the previous $k - 1$ items could result in errors. For example, it would not be surprising if, with a list of 100 words, people made many errors even if items were exposed for 5 s each.

In order for repetition detection data to provide strong support for perceptual RB, therefore, one would have to have a way of comparing the error rate in this task with the error rate in a control task that imposes similar retention and comparison requirements, but which does not depend upon perception of a repeated element. Kanwisher (1987) does not present such a control condition, and we are not sure what such a task would be.

Our approach is to hold the memory and comparison demands as constant as possible between repeated and control conditions, while also minimizing these demands overall. At the beginning of each trial, experiment participants were presented with a set of two potential target letters (designated *A* and *B,* although different letters were used on each trial). The participant then saw an RSVP display that contained precisely two target letters (two *A*s, two *B*s, or one of each) in addition to many nontarget letters. On repetition trials, the target letters were two *A*s or two *Bs.* On nonrepetition trials, they were one *A* and one *B.* The types and tokens model only predicts RB when the observer must register the critical items as separate tokens. To insure that our task would require token individuation, we needed a detection response that involved report of some attribute of each token (rather than just its mere presence.) We do this by requiring participants to report the location of the targets in a moving RSVP display (see Experiment 1). Repetition and nonrepetition trials were mixed, so that participants could not anticipate which potential target stimulus would occur next, thus providing every incentive to prepare to detect and localize both *A*s and *Bs* at all times. The logic is simple: If perception of a repeated item is somehow impaired, there should be more failures to detect and localize on repetition trials than on nonrepetition trials.

Method

The methods of Experiment 2 were the same as in Experiment 1 except as noted below.

Participants. Twelve students from the same pool used in Experiment 1 participated in this experiment (these participants were in the third condition of Experiment 1).

Apparatus and stimuli. A new set of two different potential target letters was randomly selected from the set of letters *A-J* at the beginning of each trial. Exactly two occurrences of the potential target letters were present somewhere in the RSVP display. We designated the potential target letters *A* and *B.* There were four types of trials. On *AA* trials, an *A* occurred twice in the RSVP display. On *AB* trials, an A occurred once, followed somewhere by a *B. BA* and *BB* trials followed the same pattern. The two letters drawn from the potential target set that actually occurred will be termed *actual targets* (to be distinguished from potential target letters).

Design. The target items could appear in the second through fifth positions, totaling six possible pairs of target locations (2-3, 2-4, 2-5, 3-4, 3-5, and 4-5). For each of these target locations, there were four *AA*, four *AB,* four *BA,* and four *BB* trials. The 2-3, 3-4, and 4-5 trials were "filler" trials and were not included in analysis.

Procedure. Each trial was initiated with the presentation of the potential target set in the center of the screen (the two letters side by side, separated by one character space, in the same font used in the display) for 1,000 ms. Immediately after the offset of the potential target set, the display was presented like the rightward moving RSVP displays of Experiment 1, beginning with the presentation of the fixation point.

Participants reported the position of the target letters in the RSVP display. Report phase of the trial began with the presentation of six underline characters on the screen, occupying the screen position of each of the six letters of the RSVP display. The positions were reported by pressing the numbers 1-6, in correspondence with the six possible target positions. When a number was pressed, an *X* replaced the corresponding underline character on the screen, to indicate to the participant which location had been selected. As soon as two different responses were made, the computer provided feedback on both responses. A pleasing ascending sound indicated a correct response, and a descending sound indicated an error. The two feedback sounds occurred in sequence, indicating the correctness of the first and second responses, respectively. Each feedback noise lasted 350 ms, with a separation of 500 ms. The next trial began 1.7 s after the second response was made. The proportion of correctly located targets was presented to the observer between blocks.

Results

Table 2 shows the proportion of trials where the first, second, or both targets were located correctly, broken down by trial type (*AA, AB, BA,* and *BB).* A target was considered to be correctly located if either of the participant's two responses corresponded to that location. Only trials in which the targets occurred in locations with at least one position between them were included in the analysis. There was no effect of repetition (*AA* and *BB* vs. *AB* and *BA* trials) on the probability that participants correctly located the first actual target, $F < 1$, the second actual target, $F(1, 11) = 1.0$, $p > .3$, or both actual targets, $F < 1$.

Table 2

Proportion of Trials in Which Targets Are Correctly Located in Experiment 2

	Target		
Trial type	1st	2nd	Both
AA	80.6	59.2	48.6
AB	84.1	65.8	56.1
BA	66.9	62.8	38.9
BB	68.3	60.8	41.7
Repetition	74.4	60.0	45.1
Control	75.6	64.3	47.5

Note. *A* and *B* indicate potential target items.

Discussion

Under the same stimulus conditions in which we found RB in Experiment 1, we also found that participants were essentially as able to detect and localize two occurrences of the same item as they were when the items were different. In particular, the proportion of trials on which both targets were successfully located was not significantly affected by repetition (45.1% and 47.5% for repetition and control, respectively). Although in Experiment 1 it was noted that there is some difficulty in comparing the C1 and C2 scores on control trials to the scores on repetition trials, it should be emphasized that no such problem exists here in comparing control and repetition performance. Thus, because the detection task should require Kanwisher's (1987) putative token individuation process, this is evidence against the types and tokens model.

This experiment demonstrates a situation in which two repeated items must be individuated, yet no RB occurs. It is possible, however, that in this detection task, observers did not bind the name of the target letters to unique tokens as the targets were seen. Instead, it might be that the location was bound to a token, because it is location that must be reported and not the name[4]. This would allow the observers to do the task without RB and so would be an adaptive strategy. Of course, this predicts that if observers must also report the identity of the target letters that appeared in the RSVP display, then RB should appear in either location or identity report (or both), because as both identity and location have to be reported, both identity and location have to be bound to a token. This alternative is explored in Experiment 3.

Experiment 3

This experiment tested whether RB occurs in a target detection task like that used in Experiment 2 when both location and identity of the target items must be reported. We look for RB in location, identity, and the conjunction of location and identity reports. Failure to find RB in all cases would be evidence against perceptual RB.

Method

The methods used were identical to the methods of the detection task in Experiment 2, except as noted.

Participants. Twelve students participated in this experiment for partial fulfillment of a course requirement at the University of California, San Diego.

Procedure. In addition to the procedures used in the detection task in Experiment 1, participants also reported the identity of the target letters that occurred in the RSVP display. They did this by pressing one of four keys corresponding to four alternatives displayed on the screen. If the target set was *A* and *B,* then the four alternatives were always *A-A, A-B, B-A,* and *B-B,* presented from left to right on the computer screen. So, for example, if the first target in the RSVP display was *B* and the second was *A*, the observer would press the third key from the left corresponding to the choice *B-A.* This response was made after the participant made the location response.

Results

The top section of Table 3 shows the proportion of trials where the first target, second target, and both targets were located correctly. As in Experiment 2, there was no effect of repetition in any of the three cases, $F(1, 11) = 3.3$, $.10 > p > .05$; $F < 1$; and $F < 1$, respectively. These results replicate the findings of Experiment 2.

The proportions of trials in which the first, second, and both targets were correctly identified are shown in Table 3. There was no effect of repetition on identification of either the first or second target, $F < 1$ in both cases, nor was there a significant effect on getting both correct, $F(1, 11) = 1.7$, $p > .2$.

The third section of Table 3 shows the proportions of trials in which participants correctly identified and located the first and second targets and the proportion of trials in which both targets were correctly identified and located. There was no effect of repetition on the first or second target, $F < 1$, $F(1, 11) = 1.2$, $p > .3$, respectively. There was also no effect on correctly identifying and locating both targets, $F(1, 11) = 1.1$, $p > .3$.

Inspection of Table 3 shows that in only one case is there a nonnegligible effect of repetition: The effect of repetition on identifying both targets is 7.5%. Though this does not compare with the size of the effect normally found in RB (around 30% and upward) and is not significant, it is large enough to suggest that there might still be some problem in detecting repetitions, even if not enough to account for the whole RB effect. However, the entire advantage for the nonrepetition condition is accounted for by the performance on *AB* trials. In addition, there were many more false alarms for the *AB* response than the other three: 25.1% versus 6.3%, 13.1%, and 7.3% for the *AA, BA,* and *BB* responses, respectively. This suggests a bias toward the *AB* response. If participants tended to guess *AB,* then this could fully explain the slight advantage for the *AB* condition.

To test this hypothesis we did the following analysis based on signal detection theory. For each of the four

---

[4] We thank Nancy Kanwisher for pointing out this possibility.

Table 3
proportion of Trials in Which Targets Are Correctly located or Identified or Both in Experiment 3

	Target		
Trial type	1st	2nd	Both
	Locate		
AA	72.9	60.1	46.2
AB	79.2	63.5	49.3
BA	63.9	58.0	37.8
BB	63.2	61.5	39.9
Repetition	68.1	61.1	43.1
Control	71.5	60.8	43.6
	Identify		
AA	87.5	73.6	64.2
AB	87.2	81.6	77.1
BA	67.0	56.6	52.8
BB	72.2	73.6	50.7
Repetition	79.9	73.6	57.5
Control	77.1	69.1	65.0
Locate and Identify			
AA	67.7	48.6	35.8
AB	71.5	52.8	42.7
BA	45.8	36.5	25.4
BB	49.7	47.2	26.0
Repetition	58.7	47.9	30.9
Control	58.7	44.6	34.0

Note. *A* and *B* indicate the potential target sets.

possible responses *(AA, AB, BA,* and *BB)* we computed d' and the criterion value c using the signal detection model described in Appendix B.[5] This analysis yielded the results found in Table 4. There was a significant effect of response on d', $F(3, 33) = 5.1$, $p < .05$. The effect appears to be almost entirely accounted for by the larger d' for the *AA* condition. There was also an overall effect of response on c, $F(3, 33) = 13.4$, $p < .01$. As expected, there seems to be a bias toward the *AB* response. These results, then, provide no evidence that detection ability is any worse for matching targets.

Discussion

Experiment 3 replicated and extended the findings of Experiment 2. In this experiment, participants were not significantly worse at either locating or identifying repeated targets than they were at doing the same for unrepeated targets. The most straightforward way for participants to have performed this task according to the types and tokens model would have been for them to bind the letter name and location of each target to a separate token.[6] This would have allowed them to become aware of the letters that they saw and where they occurred. This would also, however, have resulted in RB on both tasks, which did not occur. Before accepting that these findings are inconsistent with the types and tokens model, however, it is important to carefully consider how else participants might have performed the task.

Table 4
d' and c for Identification Responses in Experiment 3

Trial type	*d'*	*c*
AA	2.2	.21
AB	1.7	-.48
BA	1.5	.12
BB	1.8	.30
Repetition	2.0	.26
Control	1.6	-.18

Note. Negative values indicate a bias toward that response. c = criterion value.

One other way to perform the task would be to bind some other attribute of the targets to tokens. Say the participants bound the target locations to separate tokens. This would allow them to accurately perform the location task and, moreover, there would be no RB on this task (because locations are not repeated). This can explain the findings of Experiment 2, but it cannot explain the present findings because participants would only be aware of where the targets occurred, not which targets occurred where. Thus, for example, participants would at best perform the task with 50% accuracy on unrepeated trials, because *AB* and *BA* trials would be indistinguishable based on type information alone. This did not happen: 95% confidence interval is 60-70% accurate for unrepeated trials.

Finally, it could be that participants attempted to bind the name and location of each target to separate tokens. When individuation fails, type information is used to determine the correct response. To see how type information could help, consider the case in which the first target was a *B,* and it was individuated. If the participant does not have a record of seeing an *A* (the other member of the target set) anyplace in the display, then he or she might safely assume that it was a *BB* trial, otherwise he or she could assume it was a *BA* trial. Thus, theoretically anyway, it is possible to perform the task perfectly based on only one token and full type information.

The above explanation, it needs to be noted, would force two extensions to the standard types and tokens framework. First, it would require that when token individuation fails because the second target is a repetition, a new token is still created with all the other attributes of the item bound to it. This extension is necessary to explain why location reports do not show RB. Second, it needs to be assumed that probing the type nodes to see if a *B* was not presented (in a sequence of six letters), for example, can be performed as accurately as checking the second target token for its associated letter name (assuming individuation was successful). Otherwise, it would not follow that participants would be

---

[5] A criterion of zero means unbiased responding. Negative criterion is biased toward the response; positive is biased against the response.

[6] It is sometimes convenient to think of tokens as representing different locations or serial positions. An equivalent way to discuss this, done here, is to say that location is a feature that itself is bound to the token along with the letter name.

just as accurate in the repetition condition as the no-repetition condition. So it seems that the types and tokens model cannot account for the present findings without some fairly major (and doubtful) revisions.

Experiment 4

Up to this point we have followed two lines of investigation. First, we tried to manipulate the order in which observers retrieved the critical items from memory (the order manipulation in Experiment 1). Second, we tested whether detection of a repeated item was harder than detection of a nonrepeated item (Experiments 2 and 3). In Experiment 4 we introduce a different manipulation of the order in which the critical items are retrieved from memory and a different test of whether repeated items are simply not "seen" in RB.

Participants in the present experiment were presented with RSVP displays in which one of the letters was red. There were always six letters in the display, and it was always the fourth or fifth letter that was red. Participants had two tasks: One task was to report all of the items in the display; the other was to report which item was red. Sometimes participants reported the red item first, and at other times they made the full report first. If participants actually failed to see repeated items as suggested by perceptual RB, then red item report should have been less accurate when the red item was a repetition. This is analogous to the logic of Experiments 2 and 3.

On-line and off-line loci make different predictions about which of the critical items will be dropped. When full report is made first, both loci predict that C2 will be dropped. An on-line locus predicts this because C2 would be the second seen critical item; an off-line locus predicts this because C2 is the second retrieved critical item. When full report is made after the red report, the second repetition presented will be the first repetition retrieved (i.e., it will be retrieved for the red report) and the first repetition will be the second retrieved. Thus, off-line RB might produce the effect on C1, whereas on-line RB still predicts the effect on C2. But notice that off-line RB is also consistent with the possibility that the RB effect might not reverse (i.e., show up on C1). For one thing, we do not need to assume that retrieval for red report would affect retrieval processes occurring during full report in the same way that the first retrieval of a repetition during full report does. More importantly, however, on-line RB has no obvious way of accounting for a reversal of the RB effect.

In this experiment, participants performed both tasks on each trial and were not told which task they were to do first until after the display was presented. This guaranteed that on-line processing was equivalent for the two tasks. Thus, RB on one task but not on the other would suggest that RB is caused by the difference between the two tasks: the off-line processes. Note, however, that this prediction only applies to the first task performed on each trial, because the response for the second task might in some way be influenced by the first.

A related experiment by Kanwisher (1991, Experiment 5) should be mentioned here. She showed participants in her study RSVP displays of letters that spelled words in which one of the letters was red. The observers task was to make a full report and then circle the letter that was red. She found the standard RB effect on C2. In addition, participants circled the wrong letter significantly more often when the red letter was a second repetition than a nonrepeated letter. The tendency was for participants to circle a letter in their report that was adjacent to the red letter in the stimulus. Kanwisher argued that the second repetition was not seen (even though it was distinct in being the only red letter) and that the color red was illusorily conjoined to an adjacent item (Treisman & Schmidt, 1982). However, one could also explain the data by supposing that participants emphasized the full report task and often forgot what the red letter actually was—even though they saw it—by the time they had finished making the full report. The red item report would then be based on the approximate location of the red item but not its identity. This would explain the repetition effect on the red report task (because second repetitions are dropped in the full report) and also why letters adjacent to the correct letter were circled in its place. In any case, this experiment does not decide the issue.

Method

Participants. Fifteen students participated in this experiment for partial fulfillment of a course requirement at the University of California, San Diego.

Apparatus and stimuli. The apparatus and stimuli were the same as those used in Experiment 1, except that one letter in each stimulus (always the fourth or fifth letter) was colored red and the RSVP display was stationary.

Design. The experiment consisted of four blocks of 48 trials. There was also one block of 24 practice trials. Half of the trials in each block were repetition trials. Of these, there were an equal number of trials using 2-4 stimuli (repetitions in the second and fourth position), 2-5 stimuli, and 3-5 stimuli. Each combination of report order (red report first or full report first) and red item position (4 or 5) occurred twice for each repetition stimulus type in each block, so that the second instance of the repeated item was red half of the time. Repetition trials in which the second critical item was not the red item were filler trials and were not included in the analysis. Finally, each combination of report order and red item position occurred six times per block for control stimuli.

Procedure. Participants received written instructions describing their task. The instructions stated that letters might sometimes occur twice in a given display. The display procedure was exactly as in Experiment 1.

After the stimulus was displayed to the participant, report cues were presented. In the full report first condition, the report cue appeared exactly as in Experiment 1, and participants entered their responses in the same way. After this, a red underline character appeared in the bottom third of the screen (full report entries remained visible in the center of the screen). Participants pressed the key corresponding to the name of the letter they thought was red. They could change their response by pressing a new key corresponding to their new response. Participants pressed the key labeled DONE when ready to go on (even if they had not entered a letter). The report red first condition differed only in that the red item display appeared first, and it appeared in the top third of the display rather than in the bottom third. The red item display remained on the screen as participants made the full report, and as in the other condition, the full report cues appeared in the center of the screen.

Auditory feedback and feedback between blocks occurred as in Experiment 1, except that letters had to be reported in the correct location to be considered correct. The interval of time between the second response on a trial and the beginning of the following trial was 250 ms.

Results

Table 5 shows the results of this experiment broken down by whether full report or report of the red item was made first. Full report performance was analyzed in terms of C1, C2, and Both scores from Experiment 1. Report of the red item was analyzed in terms of the proportion of trials in which the correct letter was reported as the red letter and the proportion of trials in which the letter that followed the red letter in the stimulus string was reported instead. The score for each of these measures was computed in the same manner as in the previous experiment.

Full report first condition. The mean C1, C2, and Both scores for the full report first condition are shown in the left half of Table 5. There was a significant positive effect of repetition on the C1 score, $F(1, 14) = 5.6$, $p < .05$, and a negative effect on the C2 and Both scores, $F(1, 14) = 7.0$, $p < .05$, and $F(1, 14) = 148$, $p < .001$, respectively. These results parallel those in the forward presentation condition of the previous experiment.

Also in the left half of Table 5 are the proportion of trials in which the red item was reported correctly and the proportion of trials in which the item following the red item was reported instead, broken down by whether the stimulus was a repetition or control stimulus. There was no significant effect of repetition in either case, $F < 1$, $F(1, 14) = 1.7$, $p > .2$, respectively.

Table 5
Results of Experiment 4 Broken Down by Whether Full Report or Red Report Was Made First

	Report condition	
Critical item	Full report first	Red report first
	Full report	
C1		
Repetition	47.5	46.7
Control	42.8	40.1
C2		
Repetition	41.9	40.3
Control	52.1	46.9
Both		
Repetition	7.8	5.8
Control	45.9	41.7
	Red report	
Red correct		
Repetition	42.7	55.8
Control	43.6	57.8
Next reported		
Repetition	8.1	11.9
Control	11.0	11.0

Note. C1 = first critical item; C2 = second critical item; Both = both critical items.

Red report first condition. The right half of Table 5 shows the full report data for the red report first condition. The positive effect of repetition on the C1 score was larger than the corresponding (and significant) effect in the full report first condition (6.6% vs. 4.7%). However, it was not significant here, $F(1, 14) = 3.9$, $.05 < p < .10$. The negative effect of repetition on the C2 score was also nonsignificant, $F(1, 14) = 3.1$, $p = .10$, though there was a 6.7% effect (compared with 10.2% in the full report first condition). The effect of repetition was significant on the Both score, $F(1, 14) = 87$, $p < .001$. The data for report of the red item are shown in the right half of Table 5. There were no significant effects of repetition, $F < 1$ for both scores.

Discussion

Participants made full reports of RSVP displays and, in addition, reported which of the items were red. If the RB effect were perceptual, they should have been much less accurate in reporting the red item when it was a repetition (at least when the red report was first). However, no such effect was found for either report condition. This cannot be because the red report measure was not sensitive enough because of ceiling or floor effects: The red items were reported correctly 56.8% and 43.2% of the time for the report red first and full report first conditions, respectively. Any problem seeing the second-presented repeated item severe enough to cause the RB effect in full report (37%) should have shown up. Thus, these results, combined with those of Experiments 2 and 3, argue that the RB effect is not a perceptual effect.

This experiment also provides evidence that the RB effect is specifically due to off-line processes. On-line processing for the two tasks used in the experiment was the same because both tasks were performed on each trial. Nonetheless, the full-report task showed the standard RB effect, but the red report task did not. Thus, the difference must lie in the off-line processes.

If the effect is off-line, why was no reversal of RB onto C1 found when the red report was made first? Recall that the reason the effect might be expected on C1 is that the red report might serve as the first retrieval of the repeated item. However, this prediction depends on the assumption that participants construct their full report around the red item they just reported. But this might not be the case, so the absence of RB on C1 is consistent with an off-line locus.

Can the types and tokens model be reconciled with the data? How might one reconcile the types and tokens model with the fact that there was no repetition effect on the red report task? One might suppose that participants in this study, knowing that extra emphasis is placed on the red item, drop C1 and individuate C2 instead of the reverse. That is, contrary to what we have been assuming up to this point, observers are not bound to miss C2 when it is repeated but instead have a choice between missing C1 or C2. However, this cannot fully explain the present results either. The reason is that when full report is required, it is the second item that is dropped. Thus, the decision of which item to drop could not have been made until the report cue was presented, at which point the decision would already have been made according to the types and tokens model of RB.

Another possibility is that instead of failing to individuate the second repetition, participants merge the tokens for the two repetitions, so that the merged token has some of the attributes from one token and some from the other. Furthermore, in this experiment the position could be inherited from the first instance and the color from the second. This predicts that participants will drop the second repetition in full report, yet have no trouble naming the red item when it is a repetition, which is consistent with the data. However, a previous experiment (Kanwisher, 1991, Experiment 5) makes this model unlikely. In this earlier study, participants viewed RSVP displays of letters that formed words (or pseudowords) in which one letter was red. The participants wrote down all the letters from the display and circled the letter in their report that they thought was the red one. Thus, the stimuli and tasks used in this experiment were very similar to ours, and if participants formed a merged token in our experiment, her participants ought to have done so too. However, participants circled items close to the correct location of the red item, even when they did not report the red item. Were observers forming a merged token, they should have been circling C1 instead. Thus, the merged token theory seems unlikely.

Is a possible guessing bias canceling out a real repetition blindness effect? It has been suggested to us by N. Kanwisher (personal communication, August 27, 1992) that these results might be explained in a manner consistent with perceptual RB if one makes certain assumptions about the strategies that observers use when guessing on the red-report task. Suppose observers make their guesses from the set of items they think they saw on the entire RSVP display. This would certainly inflate the performance of observers in the repeated condition because even when they drop the repeated red item, it is likely that the other repetition will have been seen and thus might be guessed. On the other hand, if the unrepeated item is not seen, it will not be guessed. Can such a guessing bias be hiding a real repetition deficit in detection ability?[7]

In order to test the hypothesis, we applied the standard two-state guessing model. That is, when making a response, observers either know what the red item is or they guess (no partial knowledge). Symbolically,

$$R = S + (1 - 5)\,G, \tag{1}$$

where R is the probability of responding correctly, S is the probability of actually seeing the item correctly, and G is the probability of guessing the item correctly given that it was not seen correctly. Thus, if we can estimate G, we can use it to calculate an estimate of S. Theoretically, S is a measure of the observer's ability to actually detect which item was red, with any guessing biases toward one condition or another partialed out. Though this model is probably an oversimplification, if such a guessing bias exists, then the RB effect should be present in S.

G was estimated in the following way: First, the full report data was used to estimate the chance of correctly guessing which item was red for each individual trial according to the guessing strategy under consideration. This was done by dividing the number of times the actual item that was red appeared in the full report by the number of correctly reported items (if there are no correctly reported items, zero is substituted for the undefined ratio).[8] This number was called the *trial guessing probability.* Theoretically, G is equal to the average trial guessing probability over all trials in which the participants were actually guessing. However, on trials in which the participants were correct, there is no way of knowing whether they were guessing or actually correctly saw the item. So let us view the situation differently. If a participant guesses on 10 trials with a 50% chance of being correct on each of these trials, we would expect the participant—on average—to be correct on 5 of the trials and incorrect on the other 5 trials. Similarly, if we find a total of 5 incorrect trials with a trial guessing probability of 50%, then we can estimate that there were 10 trials, overall, with the same trial guessing probability. Applying this logic generally, we can estimate G, and this estimate can be computed for each participant and for both repetition and nonrepetition conditions separately. Symbolically,

$$G = \frac{\Sigma_i g_i/(1 - g_i)}{\Sigma_i 1/(1 - g_i)}, \tag{2}$$

where g_i is the trial guessing probability and i ranges over all the incorrect trials.[9]

Table 6 shows G and S for both repetition and nonrepetition conditions as a function of report condition. These scores were computed separately for each participant and were then averaged to arrive at the numbers in the table. There was a significant effect of repetition on G, $F(1, 14) = 10.3$, $p < .01$. Although this would indicate that the guessing strategy would lead to a slight bias toward the repetition condition, the effect of repetition on 5 was not significant, $F(1, 14) = 1.5$, $p > .2$, indicating that the bias would be unable to hide anything but a minuscule repetition deficit on detection ability. Certainly, a bias large enough to hide an

---

[7] Of course, if observers always guessed items that were not seen, the bias would go the other way.

[8] The number of correctly reported items is used as the denominator in order to allow for a larger calculated guessing probability (than if the number of items reported, the number of items in the display, or the size of the stimulus set were used) and therefore provides a better chance to see a large bias in favor of repeated items. This was therefore a conservative choice for nonperceptual RB.

[9] This formula is not defined when $g_i = 1$. Thus, when $g_i = 1$, we substituted the denominator plus one for the denominator in the defining formula of g_i.

Table 6
Decomposition of Red Report Task Performance Based on The Guessing Model

	Report condition	
Critical item	Full report first	Red report first
G		
Repetition	29.0	26.6
Control	25.9	22.8
S		
Repetition	18.7	42.2
Control	22.9	46.8

Note. G = the probability of guessing an item correctly, given that it was not seen correctly; *S* = the probability of seeing an item correctly.

effect the size of RB (over 30% on the Both score in this experiment) is inconsistent with the data. Furthermore, it should be emphasized that the small difference in G between repetition and control was computed under assumptions strongly biased toward finding such an effect. In summary, we cannot see how to reconcile our results with perceptual RB.

Experiment 5

In the previous experiments we have found no evidence that the RB effect is a perceptual deficit of the type suggested by Kanwisher (1987). It now appears either that the effect arises from some process occurring on-line other than token individuation (such as a memory storage or perceptual failure that only occurs during full report) or that the effect is caused by events that occur during retrieval from memory (as suggested by Experiment 4). In Experiment 5, the experimenter read aloud three letters to the study participant, who then saw an RSVP display of three letters. In addition, participants concurrently articulated "ba ba ba ..." throughout the trial (from the time just after the experimenter said the first three letters until the participant was cued to begin report). Concurrent articulation was used because without it, it seemed as if participants were able to report the prefix letters too easily, possibly using "echoic" memory. Sometimes letters occurred twice on a trial. When this happened, one of the items was in the auditory prefix and the other was in the RSVP display. Participants reported either the letters in the prefix followed by the letters in the display or just the letters in the display (they did not know which report they had to make until 1 s after the display). If RB were purely a perceptual problem of the type Kanwisher has proposed, no repetition effect should be found here as the time between the two instances of the repeated items is large (over a second) and they are presented in different modalities (the first auditory and the second visual). If RB were caused by off-line processes, on the other hand, then an effect should occur, but only in the condition in which participants report the memory prefix. RB caused by a storage failure predicts that the repeated item should be dropped even if they do not report the prefix.

Method

Participants, Twelve students participated in this experiment for partial fulfillment of a course requirement at the University of California, San Diego.

Apparatus and stimuli. The apparatus was identical to that used in Experiment 1. There were two stimuli on each trial: the *auditory prefix,* a randomly generated string of three letters from the set *A* through *J,* and the *visual stimulus,* a string of seven characters in which the first, third, fifth, and seventh were symbols drawn without replacement from the set @, #, $, %,&,*, and +, and the second, fourth, and sixth were letters from the set *A* through *J.* Repeated trials contained exactly one letter in the auditory prefix that was the same as one of the letters from the visual stimulus. The two repeated letters could occur in any of the nine possible combinations of positions, On unrepeated trials, no letter occurred in both strings, and no letter appeared more than once in either string.

Design. There were six experimental blocks and one practice block of 36 trials each. Half of the trials in each block were repetition trials, with the repetitions occurring in each combination of possible locations equally often. The report tasks were (a) *report all* (report both the auditory prefix and the visual stimulus) and (b) *report visual* (report visual stimulus only). Repetition and report task were counterbalanced. In addition, each combination of repeated letter positions occurred once for each report task within each block. The order of the trials within each block was completely randomized.

Procedure. Experiment participants received written instructions describing their task. They were told that on each trial, the experimenter would say a sequence of three letters and they would see a sequence of letters and symbols. They were also told that on some trials, they would be required to report first the auditory prefix and then the letters in the visual sequence (report all), whereas on other trials they would only have to report the letters from the visual stimulus (report visual). Finally, they were told that they would only know which task they would perform after they had seen the letters.

The sequence of events on each trial was as follows: First, the experimenter told the participant the prefix for that trial. If the participant did not hear the letters, he or she could ask the experimenter to repeat them, but only once. Next, the participant began articulating "ba ba ba ..." at a rate of about $^3/4$ s/ba. After articulation had begun, the participant pressed a key. If the participant pressed the key before articulating, the trial was not included in data analysis and the experimenter told the participant to always begin articulation before pressing the key. Immediately after the key was pressed, a fixation point was presented for 1 s, and 500 ms after the fixation point was removed, the visual stimulus was presented one character at a time. Each character was presented for 100 ms, except for the last symbol, which remained on the screen for 150 ms. One second after the stimulus display, a report cue appeared on the screen (_ _ _|_ _ _ for report all and * * *|_ _ _ for report visual). At this point, the participant, who had continued to say "ba ba ba..." throughout the trial stopped articulating. If the participant had stopped articulating before the report cue appeared, the trial was not included in data analysis and the experimenter warned the participant to articulate until the cue appeared. If the report cue was the visual-only cue, then the participant was to report the visual stimulus outloud. If the report cue was the report-all cue, the participant reported the prefix and then the visual stimulus outloud. In either case, the experimenter copied down the participant's response, and then the next trial began.

Results

The proportion of trials in which the participant reported the first critical item, the second critical item, and both critical items is shown in Table 7 as a function of report cue. There was no effect of repetition on report of the first critical item, $F(1, 11) = 3.2$, $p > .1$ (report all only). There was, however, an overall effect of repetition on report of the second item, $F(1, 11) = 16.6$, $p < .01$, and also an effect of report cue, $F(1, 11) = 21.1$, $p < .01$. The interaction of these two factors was also significant, $F(1, 11) = 10.5$, $p < .01$. Further analysis revealed that the effect of repetition was significant on the report-all task, $F(1, 11) = 24.7$, $p < .01$ but not on the report visual only task, $F(1, 11) = 1.9$, $p > .15$. The pattern of effects found on the proportion of trials in which both critical items were reported was similar. The effects of repetition and report cue were significant, $F(1, 11) = 20.2$, $p < .01$ and $F(1, 9) = 65.2$, $p < .01$, respectively. The interaction was also significant, $F(1, 9) = 29.4$, $p < .01$. Finally, the effect of repetition was found to be significant on the report-all task, $F(1, 11) = 44.8$, $p < .01$, but not on the report visual only task, $F(1, 11) = 1.9$, $p > .15$.

Discussion

In this experiment, participants showed RB effects when they had to report three letters that were spoken to them followed by three letters presented in an RSVP display, even though the occurrences of the repeated items were in different modalities and these were separated by over a second. However, participants showed no difficulty reporting a visual item that matched one of the spoken items in the condition in which they reported only the visual items. Clearly, these effects cannot be due to the sort of perceptual problem Kanwisher (1987) has suggested. This, of course, does not preclude the possibility that there could be an additional perceptual repetition effect when both critical items are visually presented. However, the previous experiments provide no evidence for such an effect. On the other hand, this experiment demonstrates that off-line retrieval processes can produce a robust RB effect: If the present effect were due to on-line processes of any type, then RB would have occurred for the visual report only condition, too, which did not happen.

Table 7
Proportion of Trials in Which Participants Include the First (C1), Second (C2), or Both Critical Items in Their Report in Experiment 5

Critical item	Report condition	
	Report visual only	Report all
C1		
Repetition		87.2
Control		89.5
C2		
Repetition	75.1	58.2
Control	78.0	69.7
Both		
Repetition		50.5
Control		63.8

General Discussion

Let us summarize what the results reveal. First, RB is not a perceptual failure resulting in a lack of awareness of the presence of repeated tokens. If it were, study participants would have had trouble making the detection reports required in Experiments 2 and 3 and the red report in Experiment 4 when the to-be-detected or red item was the second presented repetition. This was not the case. Second, the evidence suggests that RB occurs off-line (during report) rather than on-line (during presentation). In Experiment 4, full report showed RB whereas the red report task did not, even though on-line processing was equivalent for the two tasks. In addition, Experiment 5 clearly demonstrated an off-line RB effect (although Experiment 5 could not test whether an on-line effect exists, too). In summary, then, the most straightforward explanation is that RB is an off-line effect.

The one potential snag in the data is that our attempts to manipulate order of retrieval have not completely displaced RB onto C1. This is not inconsistent with our conclusions, however. The relation of the sought-after reversal to on-line and off-line models is not symmetrical. On-line models demand that the effect is not reversed because the second seen repetition should be affected. Indeed, they predict that the effect should remain on C2. Off-line models, however, make different predictions depending on the specifics of the retrieval processes. For example, when report order is reversed as in Experiment 1, a reversal of RB is predicted if we assume that participants retrieve items in the opposite order that they were seen. On the other hand, if items are retrieved in the order they were seen and reversed before output, no reversal would occur (unless the dropping occurs during output). In summary, off-line models make no firm prediction for the reversal of RB onto C1 because they attribute much of the variance in what gets reported and what gets dropped to the retrieval processes, whose form is not presently well understood.

Potential Retrieval-Based Models

Is there anything specific that we can conclude about how memory retrieval is producing RB? In order to provide the same explanation of both RB and the Ranschburg effect, one might attribute both to the same underlying bias against reporting repeated items. It is worth distinguishing between two types of bias.[10] The first type of bias, a *guessing bias,* would consist of the observer preferring to guess at items that were not seen. A guessing bias that would cause RB is one in which observers guess only items not already re-

---

[10] We thank Nancy Kanwisher for suggesting this way of categorizing bias effects and Jim Johnston for helping to refine our exposition.

ported, thus selectively raising the chance of reporting nonrepeated items. The second type of bias, a *censorship bias,* would make the observer sometimes not report an item when it was actually seen ("say something once, why say it again?"; Byrne, Weymouth, & Frantz, 1977). This type of bias would cause RB if observers often only reported one repetition even though they thought they saw two. By making this distinction, we do not mean to imply that these two types of bias cannot coexist; surely they can.

Kanwisher finds that observers do not make many "intrusion" responses when stimuli are words that form sentences (N. Kanwisher, personal communication, August 27, 1992). This suggests that there is not a large guessing bias in RB with these stimuli. When arbitrary sequences of letters are used as stimuli, however, the situation is not as dear. In Experiment 5 there were very few intrusions (a mean of 0.6 per trial), so a guessing bias does not seem capable of explaining the effect in this case. However, more intrusions were made in the other experiments, so it would be helpful to determine the possible role of a guessing bias in these experiments.

We can put an upper bound on the effect a guessing bias could produce by itself. Let $p(C = i)$, $i = 0$, 1, or 2 be the probability that the participant correctly reports exactly i of the critical items. Thus, $p(C = 2)$ is the same as the Both score used in the analysis of Experiments 1, 3, and 5. Furthermore, let $p(S = i)$ and $p(G = i)$, $i = 0$, 1, or 2 be the probability that the participant actually sees exactly i of the critical items and the probability that the participant correctly guesses i of the critical items, respectively. We will use subscripts of r and nr for repetition and nonrepetition stimuli, respectively. Because we are interested in the maximum effect of a guessing bias, we can make the following assumptions: (a) The participant always guesses items not yet reported (as a tendency to guess already-reported items will reduce the RB effect), (b) There is no censorship bias, (c) There is no perceptual RB, that is, $p(S_r = i) = p(S_{nr} = i) = p(S = i)$. Then we can write

$$p(C_{nr} = 2) = [p(G_{nr} = 2|S = 0) \bullet p(S = 0)]$$

$$+ [p(G_{nr}=1|S = \mathrm{D} \bullet p(S = 1)] + p(S = 2). \quad (3)$$

In this case, $p(C_r = i)$ is an upper bound on $p(S = i)$. (This is where the assumption of no censorship bias comes into play.) Thus, if we can estimate $p(G_{nr} = 2|S = 0)$ and $p(G_{nr} = 1|S = 1)$, then the maximum RB a guessing bias could produce ($\mathrm{MAX_{RB}}$) can be computed as

$$\mathrm{MAX_{RB}} = p(G_{nr} = 2|S = 0) \bullet p(C_r = 0)$$

$$+ p(G_{nr} = 1|S = 1) \bullet p(C_r = 1). \quad (4)$$

In Experiment 4, participants correctly reported 3.9 items, on average, and 1.9 items, on average, that did not occur in the stimulus. Rounding these numbers to 4 and 2, respectively, for convenience, $p(G_{nr} = 2|S = 0) = {}^{1/}{}_{15}$ and $P(G_{nr} = 1|S = 1) = {}^{1}/3$, under the simplifying assumption that on each trial the participant guesses 2 letters from among the 6 not already reported (as there are 10 letters overall). Table 8 shows the proportion of trials in which exactly 0, 1, or 2 critical items were correctly reported for repeated and unrepeated stimuli in Experiment 4 (full report first only). Substituting these values into Equation 4 yields a value of $\mathrm{MAX_{RB}} = .26$, whereas the actual effect was .34. Thus, a guessing bias could account for at most 75% of the effect.[11]

The above analysis suggests that even with arbitrary stimuli, RB is not purely a result of a guessing bias. However, we cannot rule out the possibility that a large portion of the effect is due to a guessing bias. RB might have both guessing and censorship components; with arbitrary stimuli, both of these components are sometimes present, but when sentences are used, only the censorship component occurs. In addition, the Ranschburg effect could similarly be explained by a combination of guessing and censorship bias. It might, therefore, be best to think of RB and the Ranschburg effect as two manifestations of a collection of biases people have that work against reporting repeated items.

The RB effect, of course, does not necessarily have to be accounted for in the same way as the Ranschburg effect. The following alternative account, for example, provides very different explanations for the effects found at slow presentation rates (the Ranschburg effect) and fast ones (the RB effect). First, assume that when presentation rate is slow, items are transferred to articulatory store as they are presented and that there is no particular problem storing repetitions there. However, when retrieving from this store, people tend to drop the second repeated item as a consequence of biases (discussed above). This accounts for the Ranschburg effect. At fast presentation, on the other hand, items cannot be placed in the articulatory store quickly enough to keep up with presentation. Instead, some items are placed in the articulatory store, some are placed in visual short-term memory, and perhaps others are stored in further (perhaps more conceptual) short-term memory systems. Some items make it into two or more stores, whereas others make it into only one or none. Retrieval for the purpose of full report is accomplished, according to this account, by piecing together information from these different stores. If one store contains *AB* and another store contains *CD*, it is clear that *A, B, C,* and *D* were seen (in some order). But if one store contains *AB* and another *BC,* it could be either because *ABC* occurred and the *B* was registered in both stores or because *ABCB* occurred and each *B* was registered in a different store. Because of this, people must either make many false reports of repetitions or miss some repetitions. If the letter were normally opted for (which seems intuitively reasonable), this would account for RB. On the other hand, this model would not predict a difficulty with reporting the

---

[11] Two of the reviewers, John Duncan and Jim Johnston, pointed out that this analysis assumes a high-threshold model, in which nothing is ever falsely perceived. In other applications, a high-threshold model can lead to incorrect conclusions (cf. Pachella, 1974). However, in the present case, misperceptions do not change the conclusion, because in our analysis they are lumped together with "guesses," so the computation would still provide an upper bound on the maximum RB effect a guessing bias could produce.

Table 8
Proportion of Trials in Which Exactly 0, 1, or 2 of the Critical Items Were Reported in Experiment 4

Critical item	Number of critical items reported 0	1	2[a]
Repetition	20.8	72.6	6.5
Control	11.3	48.2	40.6

[a] These numbers differ from the Both score found in Table 5 because the "filler" trials were not included there but are included here.

red item in the red report task of Experiment 4 as long as the participant knew which buffer to probe for the red item.

Notice that this theory has the following properties: (a) The individual stores need not have any trouble storing repetitions, (b) The locus of the effect in this model is off line, because a repeated item is actually seen and stored and is not really "dropped" until retrieval, (c) On the other hand, the reason the effect would occur, on this model, is that storage machinery as a whole is not equipped to store repetitions when its capacity is overloaded (as in full report), (d) This model does not posit perceptual RB, because the problem depends on aspects of the task beyond the question of whether the items must be registered as separate tokens (e.g., whether more than one store is required). Finally, it should be noted that this model would have difficulty accounting for RB when only a very few elements are reported (e.g., Bavelier and Potter, 1992). However, it is possible that a guessing bias produces the effect in this case and also some of the effect when more items are reported.

Conclusions

We were interested at the outset in testing whether a perceptual account like the types and tokens model proposed by Kanwisher (1987) is really plausible. Given our results, such a model now seems very hard to maintain. A more promising explanation for RB involves multiple forms of bias, such as those considered in this article. If this account is correct, then RB may simply be the Ranschburg effect revisited under slightly different conditions. Alternatively, the two effects may originate from distinct but closely related retrieval strategies. In any case, RB seems to reflect the operations and strategies involved in full report from RSVP displays rather than any fundamental and surprising characteristic of on-line perceptual processing.

References

Bavelier, D. (1992). *Phonological repetition blindness.* Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge.

Bavelier, D., & Potter, M. C. (1992). Visual and phonological codes in repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 134-147.

Bjork, E. L., & Murray, J. T. (1977). On the nature of input channels in visual processing. *Psychological Review, 84,* 472-484.

Byrne, D., Weymouth, M., & Frantz, C. (1977). Psychokiller [Recorded by Talking Heads]. On *Talking Heads: 77* [CD]. New York: Sire Records.

Crowder, R. G. (1968). Intraserial repetition effects in immediate memory. *Journal of Verbal Learning and Verbal Behavior, 7,* 446-451.

Crowder, R. G., & Melton, A. W. (1965). The Ranschburg phenomenon: Failures of immediate recall correlated with repetition of elements within a stimulus. *Psychonomic Science, 2,* 295-296.

Greene, R. L. (1991). The Ranschburg effect: The role of guessing strategies. *Memory & Cognition, 19,* 313-317.

Humphreys, G., Besner, D., & Quinlan, P. (1988). Event perception and the word repetition effect. *Journal of Experimental Psychology: General, 117,* 51-67.

Jahnke, J. C. (1969). The Ranschburg effect. *Psychological Review, 76,* 592-605.

Jahnke, J. C. (1972). The effects of intraserial and interserial repetition on recall. *Journal of Verbal Learning and Verbal Behavior, 11,* 706-716.

Kanwisher, N. (1986). *Repetition blindness: Type recognition without token individuation.* Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge.

Kanwisher, N. (1987). Repetition blindness: Type recognition without token individuation. *Cognition, 27,* 117-143.

Kanwisher, N. (1991). Repetition blindness and illusory conjunctions: Errors in binding visual types with visual tokens. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 404-421.

Kanwisher, N., & Potter, M. (1989). Repetition blindness: The effects of stimulus modality and spatial displacement. *Memory & Cognition, 17,* 117-124.

Kanwisher, N., & Potter, M. (1990). Repetition blindness: Levels of processing. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 30-47.

Macmillan, N. A., & Creelman, C. D. (1991). Detection theory: A user's guide. Cambridge, England: Cambridge University Press.

Miller, M. D., & MacKay, D. G. (1994). Repetition deafness: Repeated words in computer-compressed speech are difficult to encode and recall. *Psychological Science, 5,* 47-51.

Mozer, M. C. (1989). Types and tokens in visual letter perception. *Journal of Experimental Psychology: Human Perception and Performance, 15,* 287-303.

Pachella, R. G. (1974). The interpretation of reaction time in information processing research. In B. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition* (pp. 41-82). Hillsdale, NJ: Erlbaum.

Pashler, H., & Badgio, P. C. (1985). Visual attention and stimulus identification. *Journal of Experimental Psychology: Human Perception and Performance, 11,* 105-121.

Pashler, H., & Badgio, P. C. (1987). Attentional issues in the identification of alphanumeric characters. In M. Coltheart (Ed.), *Attention & performance: Vol. 12. The psychology of reading* (pp. 63-81). Hillsdale, NJ: Erlbaum.

Ranschburg, P. (1902). Über Hemmung gleichzeitiger Reizwirkungen. *Zeitschrift für Psychologie, 30,* 39-86.

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

Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology, 14,* 107-141.

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.

Appendix A The Ranschburg Effect Compared With Repetition Blindness

Kanwisher (1987) considered the relationship between the Ranschburg effect and RB in detail and argued that the two effects were substantially different from each other in six ways. We examine these arguments in turn.

The Ranschburg Effect Is Found With Presentation Rates of 2 Items per Second Whereas Repetition Blindness Only Occurs at Rates Over 6 Items per Second

The Ranschburg effect with visual presentation at rates of 2 items/s is actually also an RB effect. Thus, because a Ranschburg effect has been found under these conditions (e.g., Greene, 1991), presentation rate cannot be used to distinguish the two effects. The argument would have some force if it could be shown that with the same set of materials, experiment participants have difficulty reporting repetitions at rates of 2 items/s and 6 items/s but not at intermediate rates. This would suggest that the cause of the effect it slow rates (Ranschburg) and at fast rates (RB) is different. Kanwisher (1986) did not find RB at a rate of 4 items/s, even though she did find it at faster rates. However, the task she used was reporting whether a repeated pair of letters was presented in an RSVP display of six letters. The same task at 2 items/s would almost certainly not produce any RB either. A further problem with using this particular task is that there is no control condition (i.e., detection of two different items), so it is not even clear that this task shows a repetition detection deficit even at fast rates.

The Ranschburg Effect Occurs for Both Visual and Auditory Presentation, Whereas Repetition Blindness Occurs Only for Visual Presentation

In order to present spoken words rapidly, Kanwisher and Potter (1989) used compressed speech. They found no RB effect there, whereas RB occurred with "comparable" visual presentations, even though the overall error rates were approximately matched in the two conditions. It is quite plausible, though, that the errors listeners made with the compressed speech may have arisen mostly because the features necessary for correct recognition of words were degraded rather than because the recognition process had a limited time to operate (as with visual displays terminated with masks; Turvey, 1973). In short, listeners may have operated under limitations of data, rather than limitations of processing time or memory capacity. By contrast, the Ranschburg studies using auditory material used unstructured lists, which were likely to tax memory capacity. Consistent with this view, Miller and MacKay (1994) found RB with rapid auditory presentation when unstructured lists of words are used but not when the words form sentences.

Repetition Blindness Is Greater When Fewer Items Intervene Between the Repeated Pair, but the Ranschburg Effect Is Larger When More Items Intervene

This is not strictly correct. Crowder (1968) undertook what he termed a parametric investigation of the Ranschburg effect. The effect was largest when two items intervened between the repeated pair. When no or one item intervened, only a benefit of repetition was found. This is presumably because at slow presentation rates, observers notice that there are repeated items and code them as such, effectively reducing the memory load. But more importantly, in every case, the Ranschburg effect was smaller when more than two items intervened than when only two items did. Thus, the Ranschburg effect mirrors the RB effect in the relationship between effect size and the number of intervening items, except that when only one or no item intervenes, the Ranschburg effect reverses, showing repetition facilitation, whereas RB shows the usual cost. It is not so surprising that this should occur at slower presentation rates, where more time is available for deliberative processing.

The Ranschburg Effect Is Smaller Than the Repetition Blindness Effect

There is no basis for comparing the two in this way, largely because different measures are used to quantify observers' tendency to drop the second occurrence of the repeated item. Researchers studying repetition blindness have used the proportion of trials in which both repeated items are reported compared to the proportion of control trials in which both nonrepeated control items are reported. Investigators of the Ranschburg effect, on the other hand, have compared the proportion of trials in which the second instance of the repeated item is reported in the correct location to the proportion of trials in which items in the same location from stimuli without repetitions are reported in the correct location.[A1]

Whereas Output Interference Might Explain the Ranschburg Effect, It Cannot Explain Why Repetition Blindness Does Not Occur for Auditory Presentation or for Slow Visual Presentation

Output interference may or may not explain the Ranschburg effect. As to why RB does not occur for auditory or slow visual presentations, the first two sections of this Appendix deal with this.

The Ranschburg Effect Requires Items to Be Reused Between Lists (Jahnke, 1972), Whereas Repetition Blindness Does Not

Jahnke (1972) originally interpreted these finding in terms of an "isolation" effect: When a word occurs for the first time during an experiment (at 500 ms/item) and then occurs 2-3 words later, the observer notices this and so encodes that the word occurred twice. However, when words are repeated between lists, participants often do not notice the repetition (because of proactive interference) and as a result are prone to drop one of them. By this account, observers do not have problems reporting repeated items when items are not reused between lists for the same reason that they do not have problems with repeated items that are adjacent to each other (i.e., they notice that there is a repeated item). At the fast presentation rates used in RB, there may not be enough time for the observer to notice repetitions, so, by this account, whether or not items are repeated between lists should not matter.

[A1] Results from Experiment 1, in fact, show that when analyzed using the same measure, RB experiments yield data comparable to Ranschburg experiments in size.

Appendix B

Signal Detection Analysis for Experiment 3

The goal of the analysis presented here is to compute sensitivity (*d'*) and criterion (*c*) separately for each response in the task of Experiment 3. The most straightforward method would be to apply the standard signal detection model by equating the three incorrect responses. Thus, in the *AB* condition, the *AB* response would be considered a "signal" response and *AA, BA,* and *BB* would all be considered a "no signal" response. This would give us a *d'* and c for each of the four responses, but the scale would be off because we expect fewer false alarms owing to the fact that there are actually three incorrect responses, not just one. The following analysis was performed in order to avoid this scaling problem.

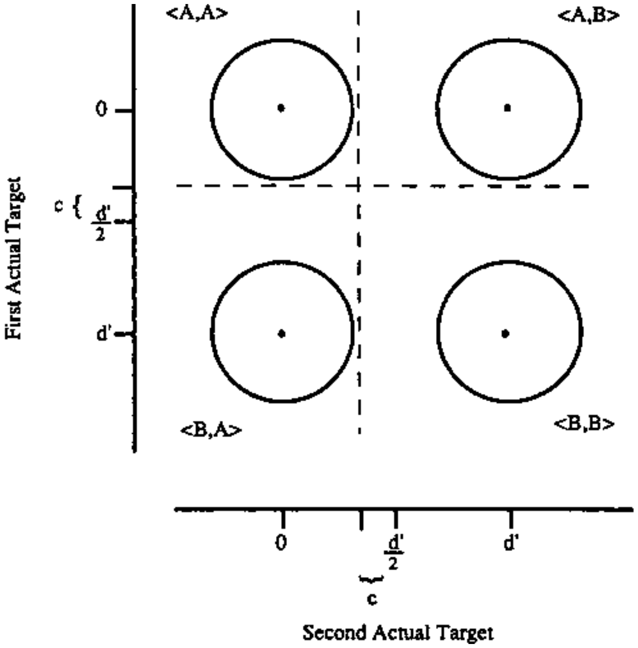

Figure B1. Diagram of Experiment 3 participants' internal representation of the actual targets. Dashed lines = criterion; d' = distance (in standard deviation units) between internal stimulus distributions along one axis; c = distance (in standard deviation units) between criterion line and midpoint between distributions along one axis.

On each trial, the stimulus can be represented in the two-coordinate space shown in Figure B1, in which the ordinate represents the internal representation of the name of the first actual target and the abscissa represents the internal representation of the name of the second actual target. The circles in the figure represent equal-likelihood curves of the four distributions (one distribution for each possible trial type: *AA, AB, BA,* and *BB*). This is a similar stimulus representation to that used by signal detection applications to same-different paradigms (see Macmillan and Creelman, 1991).

For the purposes of deriving values of *d'* and *c* separately for each response, we will assume that the same *d'* separates the distributions vertically and horizontally, and we will represent the criterion with a single real-valued number as shown in Figure B1. Thus,

$$p(\mathrm{HIT}) = \Phi(d'/2 - c) \cdot \Phi(d'/2 - c)$$

$$p(\mathrm{FA}) = [2 \cdot (1 - \Phi(d'/2 + c)) \cdot \Phi(d'/2 - c) + (1 - \Phi(d'/2 + c))^2]/3,$$

where Φ is the standard normal distribution. Using the observed hit and false alarm rates, we can compute *d'* and *c* for each participant. This analysis is carried out separately for each response (*AA, AB, BA,* and *BB*).

Received August 3, 1992
Revision received November 29, 1993
Accepted April 25, 1994