Journal of Experimental Psychology:
Learning, Memory, and Cognition
1999, Vol. 25, No. 4, 923–941

0278-7393/99/$3.00

Contextual Variability and Serial Position Effects in Free Recall

Marc W. Howard and Michael J. Kahana
Brandeis University

In immediate free recall, words recalled successively tend to come from nearby serial positions. M. J. Kahana (1996) documented this effect and showed that this tendency, which the authors refer to as the *lag recency effect,* is well described by a variant of the search of associative memory (SAM) model (J. G. W. Raaijmakers & R. M. Shiffrin, 1980, 1981). In 2 experiments, participants performed immediate, delayed, and continuous distractor free recall under conditions designed to minimize rehearsal. The lag recency effect, previously observed in immediate free recall, was also observed in delayed and continuous distractor free recall. Although two-store memory models, such as SAM, readily account for the end-of-list recency effect in immediate free recall, and its attenuation in delayed free recall, these models fail to account for the long-term recency effect. By means of analytic simulations, the authors show that both the end of list recency effect and the lag recency effect, across all distractor conditions, can be explained by a single-store model in which context, retrieved with each recalled item, serves as a cue for subsequent recalls.

The recency effect refers to the decline in memory performance with the passage of time or the presence of interfering events. Although recency effects in recognition memory are long lived and resistant to interference (e.g., Strong, 1912), recency effects in free and probed recall are short lived and are extremely vulnerable to interference (e.g., Postman & Phillips, 1965). In this article we analyze the recency effect in free recall, focusing on the details of retrieval under various distractor conditions.

In free recall, the recency effect is almost completely eliminated by 15 s of a distractor task (Glanzer & Cunitz, 1966; Postman & Phillips, 1965). The special status of the recency effect in free recall is highlighted by findings that numerous experimental manipulations and participant variables have different effects on recency and prerecency items. For example, list length (Murdock, 1962), interitem similarity (Watkins, Watkins, & Crowder, 1974), incidental learning (Marshall & Werder, 1972), and presentation rate (Murdock, 1962) significantly affect recall of prerecency but not recency items. In contrast, modality of presentation (Murdock & Walker, 1969) and interpolated distractor activity (e.g., Postman & Phillips, 1965) affect recall of recency but not prerecency items.

Two sets of findings—the vulnerability of recency in free recall and functional dissociations between memory for recency and prerecency items—led many to a two-store view of human memory (e.g., Atkinson & Shiffrin, 1968; Waugh & Norman, 1965). This view, termed by Murdock (1967) as the *modal* model, held that incoming information was maintained through rehearsal in a limited-capacity short-term store (STS) and was subsequently transferred to a long-term store (LTS).

Long-Term Recency

In the mid-1970s, the modal model came under attack from numerous directions (see Baddeley, 1986, and Crowder, 1982, for reviews). One of the most significant challenges came from the observation of the so-called long-term recency effect. Long-term recency (henceforth, LTR) refers to the well-documented finding that the recency effect, although eliminated by an end-of-list distractor task, is reinstated when participants perform a distractor task in between each of the list items and at the end of the list (Baddeley & Hitch, 1977; Bjork & Whitten, 1974; Glenberg, Bradley, Kraus, & Renzaglia, 1983; Glenberg et al., 1980; Nairne, Neath, Serra, & Byun, 1997; Neath, 1993; Neath & Crowder, 1990; Thapar & Greene, 1993; Tzeng, 1973). Specifically, for a given duration of the end-of-list distractor period (retention interval; RI), increasing the duration of the within-list distractors (interpresentation interval; IPI) results in increased recency. This increased recency is manifest in an increased probability of recall for the last list item (the level of recall effect; Glenberg et al., 1980) and a steeper slope of the serial position curve at the end of the list. These findings of LTR have been observed on time scales ranging from tenths of seconds (Neath & Crowder, 1996) to weeks (Baddeley & Hitch, 1977; Glenberg et al., 1980). LTR is

Marc W. Howard and Michael J. Kahana, Volen National Center for Complex Systems, Brandeis University.

We acknowledge support from National Institutes of Health Research Grants MH55687 and AG15852. We also acknowledge support from the W. M. Keck Foundation. We thank Donald Laming and James Nairne for sharing their unpublished results with us. We would like to thank an anonymous reviewer for suggesting that we include the modeling section and Ben Murdock, Bob Sekuler, John Lisman, and Jeremy Caplan for their careful reading of this article. We also thank Svetlana Krivoruchenko and Sachin Shah for help with the experimental work and Travis Gebhardt and Dmitri Rezhnikov for assistance with some of the programming.

Correspondence concerning this article should be addressed to Michael J. Kahana, Volen National Center for Complex Systems, M.S. 013, Brandeis University, Waltham, Massachusetts 02454. Electronic mail may be sent to kahana@brandeis.edu.

clearly not the result of continuous maintenance in a limited-capacity STS.[1]

Some researchers have suggested that the recency effect in immediate free recall remains valid evidence for STS but that LTR results from some other process (e.g., Healy & McNamara, 1996; Nairne, 1992; Raaijmakers, 1993). Others, noting the parallel effects of variables, such as semantic similarity (Greene & Crowder, 1984) and word frequency (Greene, 1986a) on continuous-distractor and immediate free recall, have argued that LTR and immediate recency result from the same process and discount STS altogether (e.g., Crowder, 1982; Greene, 1986b). A third position holds that STS, although not providing an explanation of serial position effects in free recall, may remain a useful heuristic in other areas of memory research. Working memory (Baddeley, 1986; Baddeley & Hitch, 1974) has inherited the role of STS in the modal model. Baddeley has taken this last position, concluding that "although a short-term working memory exists, it is *not* responsible for the recency effect" in free recall (italics in the original; Baddeley & Hitch, 1977, p. 647).

Complexity of Free Recall

Analyses of the serial position curve in free recall have provided much of the evidence fueling the debate over two-store models of human memory. Using serial-position-based analyses, investigators have proposed methods to isolate the contributions of long-term and short-term memory to the serial position curve (e.g., Raymond, 1969; Tulving & Patterson, 1968; Watkins, 1974). However, this interpretation of the serial position curve as a straightforward record of the quality of memory is unwarranted. The serial position curve reflects the end product of a rich and dynamic process. Recall probability, a unidimensional measure, fails to capture this process in sufficient detail to constrain theories of free recall. Models of the serial position curve have been based on distinctiveness (Murdock, 1960), spreading activation (Anderson, 1976), forward and backward chaining (Metcalfe & Murdock, 1981), and, of course, short-term and long-term memory (Atkinson & Shiffrin, 1968). The serial position curve alone has failed to distinguish among these widely varied theoretical approaches. The serial position curve, in collapsing over output positions, discards information about sequential dependencies in retrieval. In this article, we demonstrate that these sequential dependencies can distinguish among competing classes of models.

In the experiments reported in this article, two additional measures allow us to examine this process. The probability of first recall measures where in the list participants begin recall. The conditional response probability (CRP) measures how one recall follows another. Taken together, these measures contain more information than the serial position curve that is their result. A theoretical description of each of these measures is necessary for an accurate and complete description of single-trial free recall.

The probability of first recall is a serial position curve for the participants' first response. Recency items, in addition to being more likely to be recalled during the recall period, as revealed by the serial position curve, are more likely to be recalled early in participants' output sequence (this was known at least as early as Deese & Kaufman, 1957). This general tendency is revealed by inspecting the probability of first recall[2] (see Figure 1, serial position curve labeled *1st*).

Kahana (1996) introduced a measure of the tendency for participants to consecutively recall items that shared nearby list positions. This measure, the CRP, gives the probability of recalling item $i + lag$ after recalling item i. Positive lag values indicate forward recalls, whereas negative values indicate backward recalls. Large absolute values indicate remote items, whereas small values indicate items from nearby serial positions (see Figure 1, inset). For example, if the list had contained the subsequence *ABSENCE HOLLOW PUPIL* and a participant recalled *HOLLOW* followed by *PUPIL*, the recall of PUPIL would have a lag of +1. If, instead, the participant recalled *HOLLOW* followed by *ABSENCE*, the recall of *ABSENCE* would have a lag of −1. *ABSENCE* followed by PUPIL would yield a lag of +2. Note that this would be true no matter where in the list the subsequence appeared. Appendix A describes these measures in more detail.

Kahana (1996) found that the CRPs from several large studies of immediate free recall all have the following properties in common:

1. After recalling a given word, the next word recalled

[1] Defenders of the modal model have launched a number of attacks on the empirical interpretation of LTR. One possibility is that LTR is an artifact resulting from participants habituating to the distractor task. In continuous-distractor free recall, as participants become practiced at the distractor task, it is possible that it loses its effectiveness in displacing items from STS (Koppenaal & Glanzer, 1990; Poltrock & MacLeod, 1977). Consistent with this view, Koppenaal and Glanzer found that switching to a new distractor task at the end of the list significantly reduced the LTR effect. If habituation to the distractor is the sole factor producing the LTR effect, then using a different distractor after every list item should also eliminate the LTR effect. However, both Thapar and Greene (1993) and Neath (1993) found LTR under just these conditions. Furthermore, researchers who have examined continuous-distractor free recall under conditions of incidental learning have obtained significant LTR effects (Baddeley & Hitch, 1977; Glenberg et al., 1983; Neath, 1993). Consistent with the view that incidental learning disrupts rehearsal, these researchers found little or no primacy and very low levels of asymptotic recall. Similar findings have been observed when using particularly taxing distractor tasks (e.g., Watkins, Neath, & Sechler, 1989) that would certainly be expected to clear STS of list items.

[2] D. Laming (personal communication, December 16, 1996) analyzed the probability of first recall for the Murdock (1962) and Murdock and Okada (1970) studies. He found a general trend toward a "hump" in the probability of first recall (see Figure 1, serial position curve labeled *1st*). Our own secondary analyses have confirmed that the effect holds for all conditions in the Murdock and Walker (1969) study and for the two fastest conditions of the Roberts (1972) study. This hump is inconsistent with the view that recency simply reflects a monotonic decrease in the strength of list items. It is consistent with a limited-capacity short-term store driving the early stages of recall. Because the last several items are all likely to be in the short-term store, and all items in the store are equally available for recall, this predicts a flattened region at the end of the list in the probability of first recall. The size of the hump would then give a crude estimate of the capacity of this short-term store.

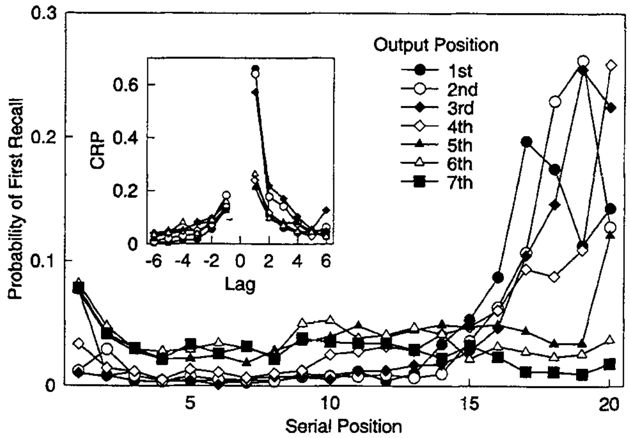

Figure 1. The complexity of free recall. This figure shows the probability of recall and conditional response probability (CRP; inset) partitioned by output position for the 20-word list, 1 item per second condition of Murdock (1962). The shape of the serial position curve is largely determined by the probability of first recall and the conditional response probability. The serial position curve labeled *1st* is the probability of first recall. Note the "hump" at the end of the list.

tends to come from a nearby serial position. We refer to this property as the *lag recency effect.*

2. There is an asymmetry in this advantage—forward recalls are more likely than backward recalls.

3. The proportion of nearby recalls dissipates as recall progresses—there is a tendency to make more nearby recalls early in the output sequence than later on.

The lag recency effect is not simply an artifact, or even simply a correlate, of end-of-list recency. At late output positions, there is no end-of-list recency effect in the probability of recall; nonetheless, the lag recency effect is still apparent in the CRP (see Figure 1, output positions greater than four). Although the serial position curve is flat at these output positions, this does not mean that serial position no longer plays a role. The relevant serial position, however, is relative to the just-recalled item. Although the end-of-list and lag recency effects are empirically distinct, an explanation of both by the same theoretical mechanism is possible. Indeed, both recency effects are explained in immediate free recall by the operation of STS.

The Modal Model and the Complexity of Free Recall

The search of associative memory model (SAM; Raaijmakers & Shiffrin, 1980, 1981), a sophisticated variant of the original Atkinson and Shiffrin (1968) modal model, is the only major memory model that can account for the detailed findings obtained in the free-recall paradigm (Kahana, 1996). In SAM, as in the modal model, STS is responsible for rehearsal and for the transfer of information into LTS. Raaijmakers and Shiffrin have shown how SAM can account for a broad range of benchmark data in free recall, including serial position effects, list length effects, presentation rate effects, growth of interresponse times with output position, part-set cuing, and many other findings. Many of this model's successes depend on the action of STS—both in the formation of associations in LTS and in the determination of what information is available for an immediate test. We refer to this version of SAM as *SAM–FR.*

SAM–FR predicts recency in immediate free recall because the contents of STS are available for recall at the time of test. Because end-of-list distractors displace list items in STS, SAM–FR predicts attenuated recency in delayed free recall. In this case, retrieval is from LTS mediated by fixed list context. SAM–FR makes the same prediction for continuous-distractor free recall. Because the same end-of-list distractor occurs before the free-recall test, SAM–FR fails to produce LTR. Nonetheless, SAM–FR describes immediate free recall in impressive detail. Kahana (1996) demonstrated that a modified version of SAM–FR predicted both the lag recency effect and the change in CRPs with output position in immediate free recall.

In SAM–FR, the existence of the lag recency effect is a consequence of the effect of STS during list presentation on the behavior of LTS during retrieval. Nearby list items tend to share time in STS, thereby strengthening interitem associations in LTS. The lag recency effect is predicted because the just-recalled item contributes to the cue for recall of the next item. Because nearby items are likely to have shared time in STS with the just-recalled item, they are likely to have a stronger interitem association in LTS. This produces the lag recency effect.

The change of the CRP with output position in immediate free recall is explained by an STS component to the CRP at early output positions. Several items from the end of the list (and hence nearby serial positions) are available from STS at the time of test. Items in STS are available for recall and are not subject to the process of retrieval from LTS. In retrieval from LTS, items must compete with all of the other items in the list and are subject to recovery failure. Because the end-of-list items in STS are not subject to this process, SAM–FR predicts a stronger lag recency effect at early output positions.

Explanations of LTR

The view that positional or temporal information drives a competitive retrieval process underlies several attempts at dealing with the phenomenon of LTR. Temporal distinctiveness theory (Glenberg & Swanson, 1986; Murdock, 1960; Nairne et al., 1997; Neath & Crowder, 1990) can be seen as a version of this general approach. These two postulates, a temporally sensitive construct and competitive retrieval, are sufficient to explain LTR. To see this, consider a construct that is sensitive to the time between study of an item and the recall test. We label this hypothetical construct as *trace strength.* In delayed and continuous-distractor free recall, the presence of a filled distractor interval causes all items in the list to have lower trace strength than they did in immediate free recall. In continuous-distractor free recall, the last item in the list will have the same absolute strength as it did in delayed free recall. However, because of the distractor prior to the last item, the other items will have even less strength than they did in delayed free recall. The last item will therefore be subject to less competition from other items, and there will be an increase in the recency effect relative to delayed free recall. This is the LTR effect. Because of the competitive retrieval process, recall probability is a function of the relative rather than the absolute strength of an item.

Variable context, as formulated by Mensink and Raaijmakers (1988, 1989), is a temporally sensitive construct, much like the hypothetical trace strength of the preceding paragraph. Fixed list context is an absolutely necessary component of Raaijmakers and Shiffrin's (1980, 1981) SAM–FR model because there is no other cue to initiate recall from LTS when STS is empty. Studies of list discrimination (e.g., Shiffrin, 1970) illustrate the importance of list context. However, SAM–FR, with all-or-none fixed list context, fails to predict such basic phenomena as proactive interference (Mensink & Raaijmakers, 1988). To deal with these phenomena, Mensink and Raaijmakers postulated that the context cue that participates in retrieval from LTS fluctuates over time. When an item is encoded, it is associated with the currently active subset of contextual elements. The activation the item receives from the test context is determined by the overlap of the item's encoding context and the context at the time of test. This overlap will be maximal at short delays and will decay over longer intervals. Combined with the competitive retrieval structure used in SAM, this should prove sufficient to generate the finding of LTR (as argued by Raaijmakers, 1993), although this has not yet been demonstrated in the literature. Although one would also expect recency in immediate free recall, the ability of this model to produce LTR and an adequate description of immediate recency with the same choice of parameters is an open question.

Context and Long-Term CRPs

A unified theory of recency effects in free recall across time scales must, by definition, maintain that LTR and end-of-list recency in immediate free recall arise from the same mechanism. In such a theory, end-of-list recency must be a function of the relative spacing of the list words and the time of test. This sensitivity to the relative spacing requires that a unified theory of recency effects maintains an equivalence principle between immediate and continuous-distractor free recall (RI = IPI) because the relative spacing is the same.

A general theory of serial position effects must also predict the observed lag recency effect in immediate free recall (Kahana, 1996). Because the relative spacing of the items within the list is the same in immediate, delayed, and continuous-distractor free recall, the equivalence principle means that a unified theory should predict a lag recency effect in all three conditions. We refer to the hypothesized finding of a lag recency effect in continuous-distractor free recall as the *long-term lag recency effect.* Finding such an effect in continuous-distractor free recall would suggest that there is a single theory underlying all recency effects in free recall. Conversely, failure to find a long-term lag recency effect would violate the equivalence principle and would suggest that there is not a unified theory of recency effects in free recall.[3]

How might one construct a unified theory of free recall that predicts a long-term lag recency effect? SAM with contextual variability (Mensink & Raaijmakers, 1988, 1989) has a construct sensitive to the temporal structure of the list and a competitive retrieval mechanism. Depending on the role of context in retrieval, such a model may or may not predict the hypothesized long-term lag recency effect. Consequently, the long-term lag recency effect provides an opportunity to distinguish two large classes of contextual variability models. We refer to these two classes of models as *retrieved-* and *passive-context formulations.* This question of retrieved versus passive context has not explicitly been addressed in the memory modeling literature.[4]

Suppose that when an item is recalled, the context at the time it was encoded is also retrieved. This context will have a greater overlap with the context associated with neighboring items than with the context associated with more remote items. This retrieved context will be a more effective cue for neighboring items. If context is not retrieved, then the context at the time of test will serve as the cue throughout recall. Because the contextual cue will be the same at any given retrieval attempt, regardless of what item is recalled, the retrieved item will have no bearing on which item is retrieved next (in the absence of direct interitem associations). We refer to the former case as the retrieved-context formulation and the latter case as the passive-context formulation of contextual variability SAM.[5] Mensink and

[3] SAM–FR makes clear predictions about the lag-recency effect in immediate, delayed, and continuous-distractor free recall. The fit of SAM–FR to the lag recency effect in immediate free recall has been documented by Kahana (1996). SAM–FR predicts the change in the CRP with output position because of retrieval from STS at early output positions. Because there is no retrieval from STS in delayed free recall, SAM–FR predicts no change in the CRP with output position. In continuous-distractor free recall, as a result of the interitem distractor, list items do not share time in STS. For this reason, SAM–FR predicts that there will be no lag recency effect in this task.

[4] It seems to us that most investigators assume a passive formulation for contextual variability. For instance, Mensink and Raaijmakers (1988) in their treatment of the *A–B/A–C* paradigm used the state of context at the time of test to retrieve both List 1 and List 2 items. They did not use the context associated with the first retrieved item to serve as the cue for the next retrieval attempt. This latter approach would have constituted a retrieved context formulation. In some cases, Mensink and Raaijmakers made use of retrieved context—they used the context stored in an image to drive list discrimination.

[5] The important distinction between retrieved- and passive-context also has implications for distinctiveness-based approaches. Nairne et al. (1997) have explored the possibility that the perturbation process that gives rise to positional uncertainty continues during retrieval. This is analogous to the passive formulation of contextual variability. Positional distinctiveness could be cast into an active formulation by, for instance, calculating the distinctiveness of subsequent recalls from the two sublists formed by breaking the list at the point of the just-recalled item. Similarly, the Glenberg and Swanson (1986) model could be made active by constraining the temporally defined search sets to be centered on the item just recalled.

Table 1
End-of-list and Lag Recency Effects in Free Recall (Before Current Experiments)

Experimental condition	Recency effect	SAM–FR	Passive context	Retrieved context	Data
Immediate	End-of-list	Y	—	—	Y
	Lag	Y	—	—	Y
Delayed	End-of-list	N	N	N	N
	Lag	Y	—	—	—
Continuous distractor	End-of-list	N	Y	Y	Y
	Lag	N	N	Y	—

Note. Search of associative memory—free recall (SAM–FR) gives the predictions of the Raaijmakers and Shiffrin (1980, 1981) version of SAM that relies on a short-term store and fixed list context. Passive context gives the predictions of a version of SAM based exclusively on variable context, without retrieval of context. Retrieved context gives the predictions of a version of SAM based on variable context that is retrieved and used as a cue for subsequent recalls. Entries for lag indicate whether a lag recency effect is observed (i.e., whether the conditional response probability is graded or not). Dashes indicate that the result is not known or the prediction of the model is not obvious.

Raaijmaker's SAM (1988, 1989) should predict LTR either way, but it only predicts a long-term lag recency effect if context is retrieved.

Experiment 1

Table 1 summarizes what is known about recency effects in free recall before the results presented in this article. The CRP in continuous-distractor free recall enables us to distinguish between retrieved- and passive-context models. Passive models predict no long-term lag recency effect. Retrieved-context models predict that there will be a long-term lag recency effect. This is important in determining whether it is necessary to retain STS as a component of a description of immediate free recall. A unified theory of recency effects in free recall should predict a long-term lag recency effect and should not require STS.

The CRP and the probability of first recall require extensive data collected from well-practiced participants. Because the presence of rehearsal is important in how we interpret the lag recency effect, we made an effort to attenuate rehearsal by utilizing a fast presentation rate and by requiring participants to make concreteness judgments on each presented item. In Experiment 1, we investigated the dynamics of retrieval in immediate and delayed free recall only. Because immediate and delayed free recall are well studied, compared with continuous-distractor free recall, this provided us with a baseline for comparison for Experiment 2 under these specific conditions. In Experiment 2, we examined continuous distractor free recall under similar conditions.

Method

Participants. Sixty-three Brandeis undergraduates participated to fulfill a course requirement. All of these participants took part in a single 1-hr session.

Procedure. Participants studied lists of words for a subsequent free-recall test. In an immediate condition, the free-recall test was given immediately after list presentation. In a delayed condition, participants performed an arithmetic distractor task for at least 10 s before recall. Lists were composed of 12 items chosen at random and without replacement from the Toronto Noun Pool (Friendly, Franklin, Hoffman, & Rubin, 1982).

Lists were presented visually at a rate of 1 word per second. During list presentation, participants were required to perform a semantic-orienting task on the presented words. The participants were to press the left control key if they judged the word to be concrete and the right control key if they judged it to be abstract. The presentation rate of the items was not dependent on the concreteness judgments.

In the immediate condition, participants were cued to begin recall immediately after list presentation. Recall was cued with the presentation of three asterisks accompanied by a 500-ms tone. Participants were given 45 s to recall as many items as possible from the list. Vocal responses were recorded for later scoring. A semiautomated speech-parsing algorithm[6] was used to assist with off-line scoring of responses and determination of interresponse times.

In the delayed condition, before free recall, participants were given an arithmetic distractor task that lasted at least 10 s. In this task, participants made true–false judgments on simple arithmetic equations as quickly and as accurately as possible. Each equation remained on the screen until a response was made. Equations were of the form $A + B + C = D$, where A, B, and C were randomly chosen integers from 0 to 9. On half of the trials the equation was true (i.e., $D = A + B + C$), and on the other half of the trials the equation was false ($D = A + B + C \pm 1$). The signal to begin free recall did not begin until the participant finished the problem he or she was working on.

Before the experimental trials, participants were given instructions on how to perform the orienting task. Participants then made concreteness judgments on two practice lists each consisting of 20 words that were not in the Toronto Noun Pool. The first of these practice lists was presented at a self-paced rate. The second list was presented at a 1 word per second rate—the same rate that was used in the actual experiment. After practicing the orienting task, participants were given instructions and practice on the performance of the arithmetic distractor task. Participants were then instructed to perform the orienting task with the words in the experimental lists as before, except that there would be a memory test on the words. Participants were given standard free-recall instructions and were warned that sometimes there would be a math test between the end of the list and the signal to begin recall. They were instructed to respond to the math problems as quickly as possible without sacrificing accuracy. Participants were given 25

[6] Retrieved from D. Utin and M. J. Kahana at http://fechner.ccs.brandeis.edu.

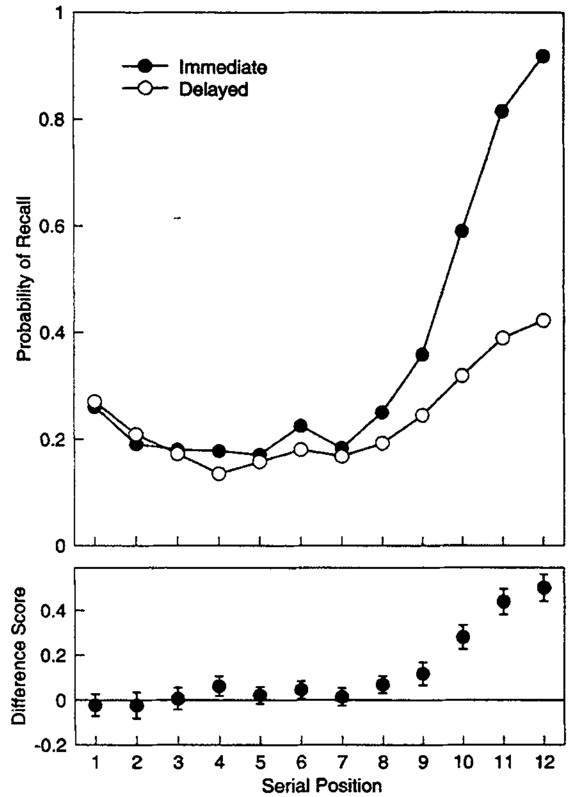

Figure 2. Serial position curves: Experiment 1. Top part of figure shows serial position curves for the immediate and delayed conditions of Experiment 1. In this experiment, participants performed a semantic-orienting task on list items to attenuate rehearsal. In the delayed condition, participants performed 10 s of an arithmetic distractor task before recalling the list items. The differences between the delayed and immediate conditions are plotted at bottom of figure. Error bars reflect 95% confidence intervals for within-subject designs calculated according to the procedure of Loftus and Masson (1994).

lists, with conditions randomized within subject after the first two practice trials. The first practice trial was from the delayed condition, and the second was from the immediate condition. In this way, participants had experience with both conditions before the experimenter left the room. These two trials were treated as practice and removed from further analysis.

Results

Serial position effects, probability of first recall, and CRP curves are reported for both immediate and delayed test conditions. Figure 2 shows serial position curves for all trials in the immediate and delayed conditions. Consistent with previous studies, the 10-s end-of-list arithmetic distractor has virtually no effect on the early list items but dramatically reduces the recency effect. Because the primacy effect is associated with rehearsal, the use of a semantic-orienting task to minimize rehearsal was expected to result in a diminished primacy effect compared with other free-recall studies (as in Marshall & Werder, 1972).

The probability of first recall for the immediate and delayed conditions is shown in Figure 3. There is again a strong recency effect for the immediate condition and a reduced recency effect in the delayed condition. The curve for the immediate condition is clearly positively accelerated. This constitutes a qualitative difference between the immediate free-recall data collected in this experiment and that from previous studies (see Footnote 2 and the serial position curve labeled *1st* in Figure 1). We attribute this qualitative difference to a disruption of rehearsal.

Figure 4 shows CRP curves partitioned by output position for immediate and delayed free recall. There is clearly a substantial effect of output position on the CRP in the immediate condition, but there is no such effect on the delayed CRPs. The change with output position for the immediate condition is consistent with the results reported in Kahana (1996, Figure 3) for other large studies of immediate free recall. The pattern reported by Kahana is that the CRP changes shape for about the number of output positions associated with the recency effect and then stabilizes. Taken together, the stabilization at late output positions in immediate free recall and the static nature of the CRP in delayed free recall suggest that the change in the CRP with output

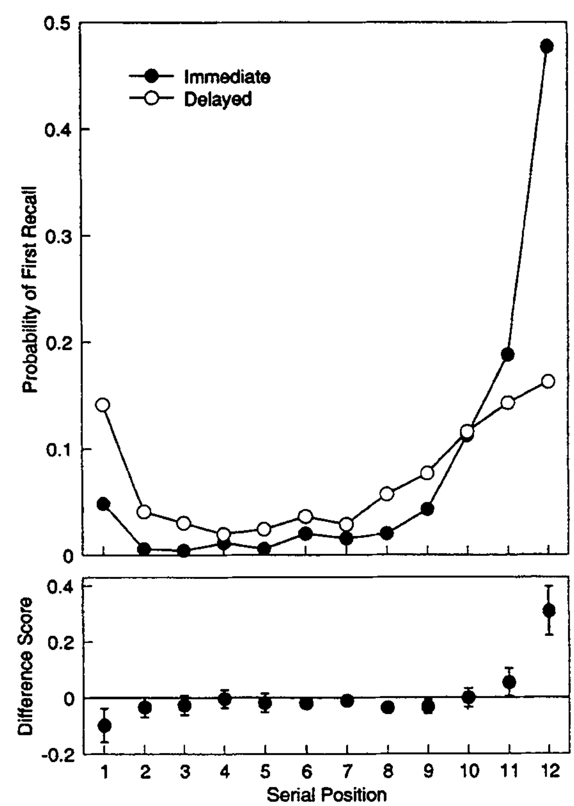

Figure 3. Probability of first recall: Experiment 1. Top part of this figure shows the probability of first recall as a function of serial position for the immediate and delayed conditions of Experiment 1. The differences between the delayed and immediate conditions are plotted at bottom of figure. Error bars reflect 95% confidence intervals for within-subject designs calculated according to the procedure of Loftus and Masson (1994).

position is a function of the absolute rather than relative spacing of study and test. The probability of first recall and serial position curves for the delayed condition show a slight recency effect. The qualitative difference in the CRP with output position in immediate and delayed free recall makes explanations of this residual recency effect on the basis of STS tenuous.[7]

Because the CRP did not change with output position in delayed free recall, the CRP was collapsed over all output positions and recalculated (shown in Figure 5). The average CRP–lag correlation (across subject) for forward lags ranging from 2 to 6 was $-.23$ ($p < .001$), and the average correlation for backward lags ranging from -6 to -2 was .19 ($p < .01$). These analyses confirm that the lag recency effect extends beyond the immediately adjacent list items in this experiment.

Experiment 2

In Experiment 1, we examined immediate and delayed free recall under conditions designed to eliminate rehearsal.

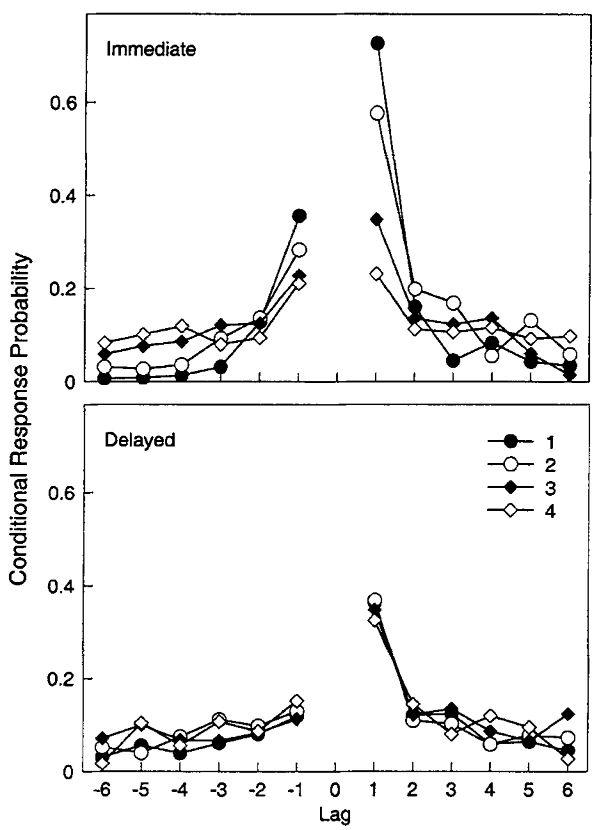

Figure 4. Conditional response probabilities (CRPs): Experiment 1. This figure shows the CRP curves from Experiment 1. Data from the immediate condition are plotted at the top of the figure; data from the delayed condition are plotted at the bottom. Lag is the distance in serial (input) positions between successively recalled items. Here, the CRP curves are partitioned by output position (1–4). At later output positions, there is insufficient data to plot reliable CRP curves.

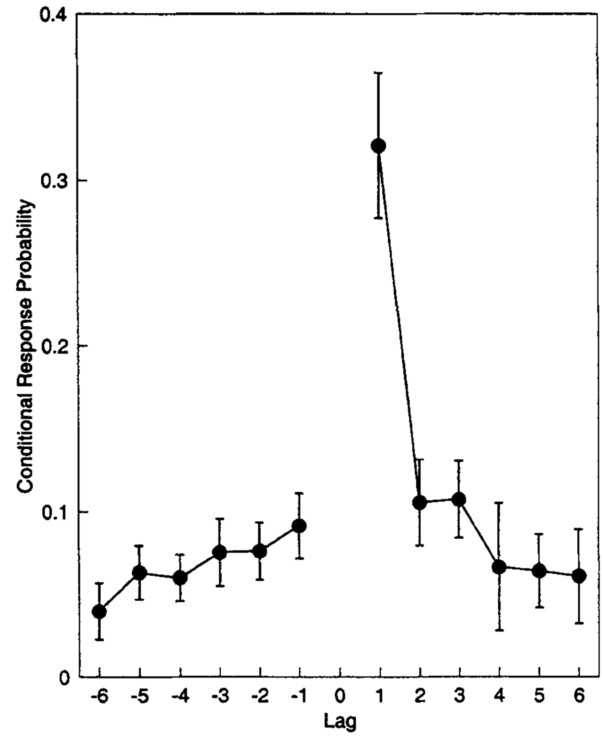

Figure 5. The lag recency effect: Experiment 1, delayed free recall. This figure shows conditional response probability (CRP) curves collapsed over output positions from the delayed condition of Experiment 1. Lag is the distance in serial (input) positions between successively recalled items. The decline in the CRP from Lag 2 to Lag 6 in the forward direction and from Lag -2 to Lag -6 in the backward direction are both statistically reliable (see text for details). Error bars reflect 95% confidence intervals for within-subject designs calculated according to the procedure of Loftus and Masson (1994).

In Experiment 2, we examined continuous-distractor free recall. Continuous-distractor free recall offers a challenge for a two-store account of recency effects. Experiment 2 offered another opportunity to replicate the long-term re-

[7] To examine participants' compliance with our instructions to perform the arithmetic distractor task during the delay, we divided trials into two groups: Trials that were below average on arithmetic performance were assigned to one group, and those that were above average were assigned to the other group. If variability in arithmetic performance was related to surreptitious rehearsal of end-of-list items during the delay, then trials on which participants performed poorly on the arithmetic task might exhibit greater recency in delayed free recall. In both Experiment 1 and Experiment 2, the serial position curves and CRPs did not differ for the two groups of trials. This suggests that variation in arithmetic performance was not related to surreptitious rehearsal during the arithmetic distractor periods. In Experiment 2, in which participants were given multiple sessions of practice, one might expect that as sessions progress, participants allocate less effort to the distractor task in order to recall more words. If this were the case, recall performance should have increased over sessions. Fortunately, we found no significant changes in recall performance from early trials to late trials. This suggests that participants remained compliant with our instructions throughout the experiment.

cency effect. The hypothesized long-term lag recency effect is the critical test needed to distinguish between passive- and retrieved-context contextual variability models. In addition, a long-term lag recency effect is a natural prediction of almost any unitary explanation of free recall across different time scales. In Experiment 2 we followed the procedure of Experiment 1 as closely as possible.

Method

Participants. Forty-two participants were tested in a trial session to determine who would be invited to take part in the full experiment. Sixteen participants took part in the full 10-session experiment. All participants were Brandeis undergraduates who participated for payment. Participants were paid $7.50 for their participation in the trial session, which lasted about 1 hr. Participants were selected to take part in the full experiment on the basis of their performance on the orienting and distractor tasks. The measure used to describe performance on the arithmetic distractor was seconds per raw score. This measure was calculated by taking the total time spent performing arithmetic and dividing that by the difference between the number of problems correct and the number of problems incorrect. Participants who were invited back got a score of no more than 3.35 s per raw score and responded to at least 70% of the orienting tasks. The 16 participants who met these criteria and who were willing to commit to 10 additional sessions were paid $6.00 for each session in the experiment, with an additional bonus of up to $1.50 paid on the basis of their performance on the distractor and orienting tasks. Data from the trial session were excluded from further analysis.

Procedure. There were four conditions in Experiment 2. All four conditions had a filled RI of 16 s between study of the last item and the beginning of the recall period. The conditions varied in the length of the filled distractor period between items within the list (IPI). Condition 0 had no IPI; Condition 1 had an IPI of one problem (about 2.5 s); in Condition 2, IPI = RI/2 = 8 s; and in Condition 3, IPI = RI = 16 s. Words were presented at a rate of 1 word/1.2 s. Participants were given 60 s for free recall. The presentation time and recall periods were increased slightly from the values used in Experiment 1 (1 word per second and 45 s for recall) to ensure reasonable performance levels in the difficult long-IPI conditions. Because of the increased duration of each trial, there were 15 trials in a 1-hr session, rather than 25 as in Experiment 1. The procedure for Experiment 2 followed that of Experiment 1 in all other respects.

Results and Discussion

Serial position effects, probability of first recall, and CRP analyses are reported for the various distractor conditions. For all situations in which we report a difference between the no-IPI condition (Condition 0) and the longest IPI condition (Condition 3), the other two conditions fall in rank order between them. For clarity of viewing, these intermediate conditions are omitted from the figures. For purposes of comparison between experiments, Figure 6 shows the serial position curve for the delayed condition in Experiment 1 (IPI = 0, RI = 10) and the no-IPI condition in Experiment 2 (Condition 0: IPI = 0, RI = 16). As can be seen from the figure, there is neither a trend toward more recency nor more primacy in Experiment 2, although the level of recall is higher. This is consistent with the longer free-recall period and slower presentation rate used in Experiment 2.

Figure 7 shows serial position curves for the no-IPI and longest IPI conditions. Recall of the most recent item in the longest IPI condition was significantly better than it was in the no-IPI condition, $t(15) = -2.6$, $p < .02$. In contrast, recall of prior list items was significantly better in the no-IPI condition for all serial positions ($p < .05$). This demonstrates classic LTR: As the length of the IPI is increased, recency is enhanced.

Another manifestation of LTR can be seen in the effect of the IPI on the probability of first recall (shown in Figure 8). Increased IPI results in a strong tendency to initiate recall at the end of the list. As in Experiment 1, probability of first recall exhibits monotonic recency. If anything, the recency portion of the longest IPI condition is more sharply acceler-

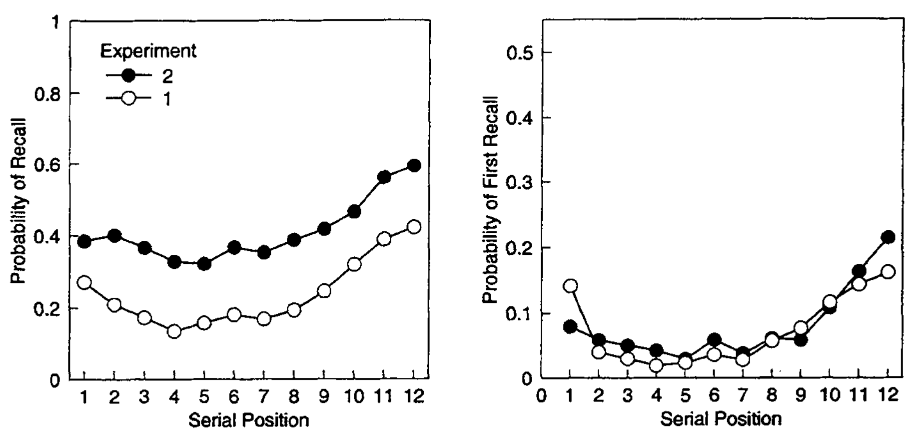

Figure 6. Comparison of the serial position effects in the delayed conditions of Experiments 1 and 2. The right side of this figure shows probability of first recall as a function of serial position for each of the experiments. The left side shows serial position curves for each of the experiments. See text for further details.

ating than the analogous curve for the immediate condition of Experiment 1.

Figure 9 shows the CRP curves for the no-IPI and longest IPI conditions. As in the delayed condition of Experiment 1, the CRP did not change noticeably with output position. Consequently, the curves shown in Figure 9 are collapsed over all output positions. Despite the presence of a full 16 s of distraction between presentation of each item in the longest IPI condition, the CRP in this condition was not substantially different from the CRP in the no-IPI condition. The CRP–lag correlation from the no-IPI condition over Lags 2 to 6 was $-.33$ ($p < .01$). The CRP–lag correlation for backward recalls over the range -6 to 2 was .21 ($p < .05$). As in Experiment 1, we conclude that the lag recency effect was not simply due to an advantage for immediately adjacent items. The same is true of the long-term lag recency effect. For the longest IPI condition, the CRP–lag correlation for forward recalls over the range 2 to 6 was $-.19$ ($p < .05$). For backward recalls over the range from -6 to -2, the correlation was .27 ($p < .05$). This confirms the existence of the hypothesized long-term lag recency effect.

The long-term lag recency effect observed in continuous-distractor free recall raises serious doubts about the dependence of the lag recency effect on co-occupancy in STS. The finding of a lag recency effect in immediate, delayed, and continuous-distractor free recall suggests that perhaps a single memory process gives rise to the lag recency effect in all three conditions.

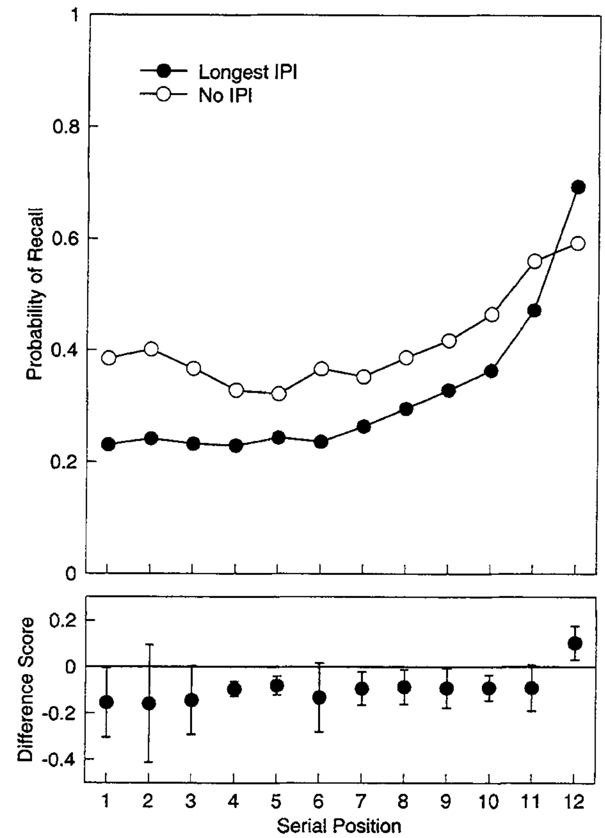

Figure 7. The long-term recency effect. The top of this figure shows the serial position curves for the two extreme conditions of Experiment 2. In both conditions, participants were given a 16-s arithmetic distractor at the end-of-list presentation and before recall. In the no-IPI condition, list items were presented successively at a rate of 1.2 s per item. In the longest IPI condition, there was a 16-s arithmetic distractor task between each list item. These data illustrate what is known as the *long-term recency effect* (e.g., Bjork & Whitten, 1974). The bottom of this figure shows the difference between serial position curves in the two extreme conditions. Error bars reflect 95% confidence intervals for within-subject designs calculated according to the procedure of Loftus and Masson (1994). IPI = interpresentation interval.

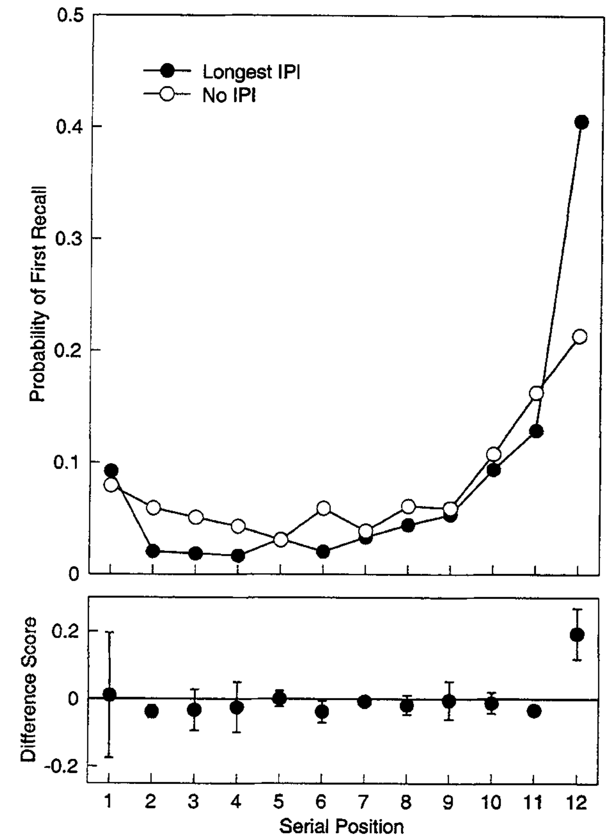

Figure 8. Long-term recency in the probability of first recall. The top of this figure shows the probability of first recall as a function of serial position for the no-IPI and longest IPI conditions of Experiment 2. In both conditions, participants were given a 16-s arithmetic distractor at the end of list presentation and before recall. In the no-IPI condition, list items were presented successively at a rate of 1.2 s per item. In the longest IPI condition, there was a 16-s arithmetic distractor between each list item. The differences between the no-IPI and longest IPI conditions are plotted in the lower part of the figure. Error bars reflect 95% confidence intervals for within-subject designs calculated according to the procedure of Loftus and Masson (1994). IPI = interpresentation interval.

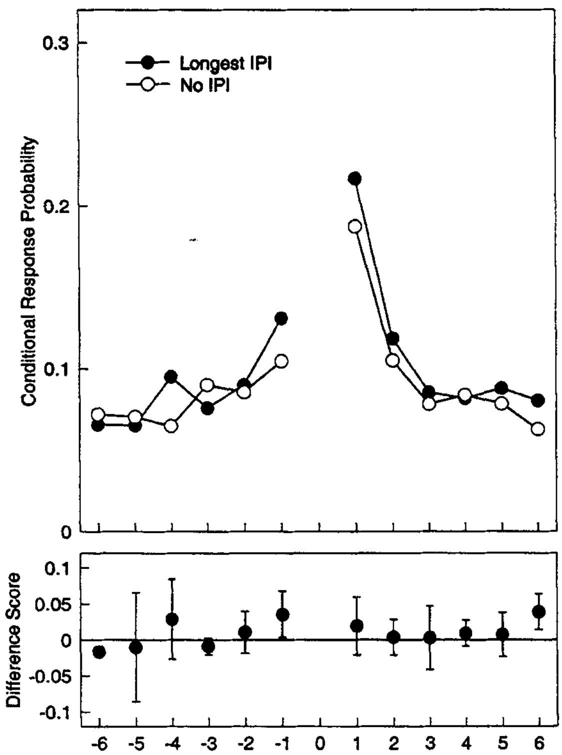

Figure 9. The long-term lag recency effect. The top of this figure shows conditional response probability curves collapsed over output positions from the extreme conditions of Experiment 2. Lag is the distance in serial (input) positions between successively recalled items. The differences between the no-IPI and longest IPI conditions are plotted at the bottom of this figure. Error bars reflect 95% confidence intervals for within-subject designs calculated according to the procedure of Loftus and Masson (1994). IPI = interpresentation interval.

Modeling

Two main variants of the SAM model were examined: a two-store model with fixed list context, SAM–FR (Raaijmakers & Shiffrin, 1980, 1981), and a model with variable context (Mensink & Raaijmakers, 1988, 1989) but no contribution from STS. The contextual variability model had two subvariants: a passive formulation, in which the context at the time of test served as the retrieval cue throughout the recall period, and a retrieved formulation, in which the context of the retrieved item served as the cue for the subsequent recall. These subvariants differed only in their predictions regarding the lag recency effect.

Our goal in modeling serial position effects in free recall was not to determine which model is the "right" one but rather, to evaluate the ability of the mechanisms, STS, passive context, and retrieved context to account for the pattern of results. With detailed modeling, we can address the ability of a single mechanism to account simultaneously for both the end-of-list recency effect and the lag-recency effect across conditions.

Method

A successful model of serial position effects in free recall should account for the end-of-list recency effect and the lag recency effect across conditions. End-of-list recency is concisely described by the probability of first recall. This was especially striking in Experiment 2, in which the lag recency effects from delayed and continuous distractor free recall were highly similar (see Figure 9) and constant with output position. LTR is thus a result of the large end-of-list recency effect in the probability of first recall (see Figure 8). Insofar as a similar curve results in immediate free recall, we can take the probability of first recall as the primary determinant of end-of-list recency in the serial position curve. The immediate condition of Experiment 1 was taken as representative of immediate free recall. The no-IPI condition of Experiment 2 was taken as representative of delayed free recall. The longest IPI condition was taken as representative of continuous-distractor free recall.[8]

The use of the probability of first recall and the CRP simplifies the task of understanding the empirical process of free recall. These measures also lend themselves more readily to analytic solution than does the serial position curve. Rather than a simulation of the entire process of free recall, followed by extracting the measures of interest, the predictions of the models for these measures can be derived explicitly. The relevant equations are detailed in Appendix B.

The models were equated for number of free parameters and process of retrieval. Because we were interested in assessing the advantages of the various structural assumptions across conditions, we attempted no mixing of models. In the variants in which variable context was used, there was no contribution whatsoever from STS, and vice versa. Both SAM–FR and the contextual variability models used the same retrieval process from LTS (described in detail in Raaijmakers & Shiffrin, 1980, 1981; see also Appendix B). The effect of the interitem and end-of-list delays was not allowed to vary across models or conditions but was held at the value 16 s. Each of the models had a subset of its free parameters that did not affect the CRPs. The models were first fit by using the simplex method of Nelder and Mead (1965) to the probability of first recall.[9] The parameter values obtained were then kept fixed while fitting the remainder of the parameters to the CRPs.

[8] Because the CRP changes significantly with output position in immediate free recall, we did not attempt to fit the CRP from the immediate condition of Experiment 1. SAM–FR has already been shown to be consistent with the change in CRP with output position (Kahana, 1996). In their present forms, the models based on contextual variability cannot capture this feature of the data. This is an important barrier to acceptance of these models as a complete explanation of immediate free recall. A retrieved context model might be able to explain the change in the CRP with output position if we included a more complete explanation of recall latencies. The explanation would go something as follows: Retrieval of context takes some finite period of time, whereas recall latency is a function of the absolute, rather than relative, strength of the cue. In immediate free recall, the absolute strength of the context cue is high, leading to recall of several items from the same end-of-list context cue and a steeper CRP at early output positions. Such an endeavor is beyond the limited and focused goals of the present work. The issue of why the lag recency effect changes with output position in immediate free recall is likely a less basic question than the issue of why there is a lag recency effect at all.

[9] The simplex was run for a maximum of 1,000 iterations with a stopping tolerance of 0.001.

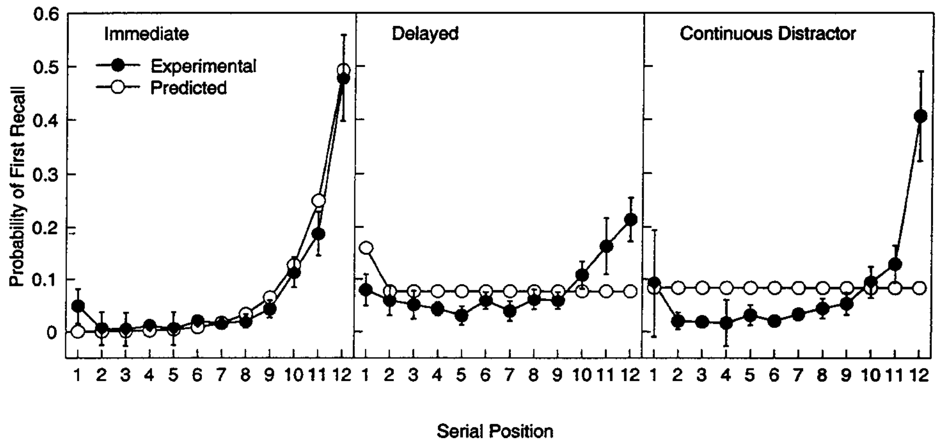

Figure 10. Raaijmakers and Shiffrin (1980, 1981) SAM–FR (search of associative memory—free recall) model: probability of first recall. The model based on the operation of short-term store, with fixed list context, adequately describes the recency effect in immediate free recall and the reduction of recency in delayed free recall, but fails to capture the long-term recency effect. There is a slight primacy effect and negative recency effect in the model's delayed free-recall performance. The best fitting parameter values were $r = 2.03$ for the buffer capacity and $a = 0.14$ for the strength of the item-to-context parameter. For the three conditions, $\chi^2(33) = 442.0$, $p < .001$. The individual contribution to the chi-square from each condition was significant.

Results

Raaijmakers and Shiffrin (1980, 1981) SAM–FR. Figure 10 shows the fit of SAM–FR to the probability of first recall in immediate, delayed, and continuous-distractor free recall. As expected, SAM–FR predicted end-of-list recency in immediate free recall but failed to capture the long-term recency effect. Figure 11 shows the fit of SAM–FR to the CRP curves from delayed and continuous-distractor free recall. SAM–FR adequately described the lag recency effect in delayed free recall. However, it failed to describe the long-term lag recency effect.

Contextual variability models. Figure 12 shows the fit of the contextual variability model to the probability of first recall in the three distractor conditions. It performed better than SAM–FR in describing the qualitative pattern of LTR. There is a distinct recency effect in immediate free recall, a decrement of recency in delayed free recall, and an increase relative to delayed free recall in the continuous-distractor condition.

The CRP distinguishes between passive- and retrieved-context formulations. Figure 13 shows the predictions of the retrieved-context formulation. Figure 14 shows the predictions of the passive-context formulation. As expected, the passive-context model failed to predict the lag recency effect in either delayed or continuous-distractor free recall. The retrieved-context model correctly predicted the existence of the lag recency effect in delayed free recall as well as the long-term lag recency effect.

Discussion

As in prior work, the Raaijmakers and Shiffrin (1980, 1981) SAM–FR model did an adequate job in describing the recency effect in immediate free recall. In delayed free recall, SAM–FR provided an adequate description of the lag recency effect but underpredicted the end-of-list recency effect. In continuous-distractor free recall, as would be expected, SAM–FR failed to predict LTR. Crucially, it failed

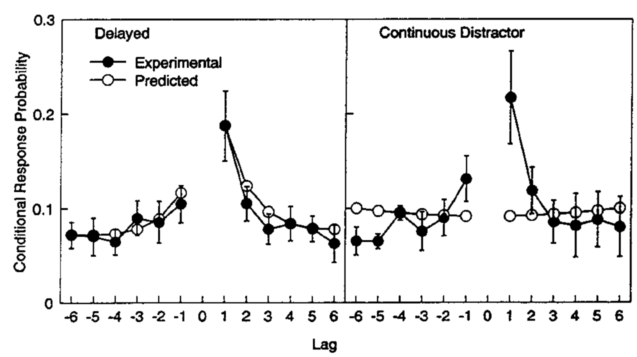

Figure 11. Raaijmakers and Shiffrin (1980, 1981) SAM–FR (search of associative memory—free recall) model: conditional response probability (CRP). The model based exclusively on the operation of short-term store, with fixed list context, describes the shape of the CRP in delayed free recall but fails to predict the long-term CRP effect. The best fitting parameter values were $b_F = 0.27$ for the forward item-to-item association and $b_B = 0.18$ for the backward item-to-item association. The residual strength, d, was found to have a minimal effect, so it was kept fixed at 0.1. The other parameters were kept fixed at the values given in Figure 10. For the five fits together, $\chi^2(55) = 524.7, p < .001$. The contribution to this value from the CRP in delayed free recall was $\chi^2(7) = 14.9, p > .03$. The contribution from the continuous distractor CRP was $\chi^2(7) = 67.5, p < .001$.

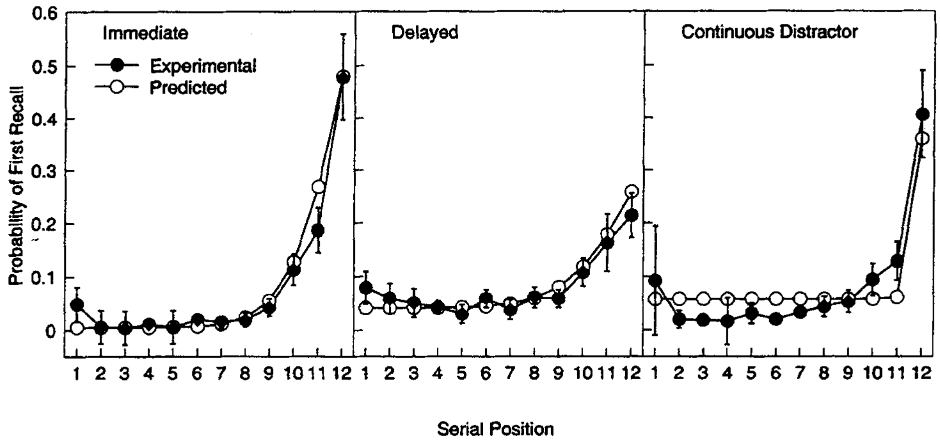

Figure 12. Contextual variability SAM (search of associative memory): probability of first recall. The model with variable context and no short-term store does an adequate job of characterizing the qualitative pattern of recency in probability of first recall in immediate, delayed, and continuous-distractor free recall. The best fitting values of the model parameters were $\beta = 0.096$ and $\gamma = 0.0099$ for the two rate constants. K and a were held fixed ($K = 1$, $a = 0.1$). For the three conditions taken together, $\chi^2(33) = 138.9, p < .001$. The contribution to this value from the immediate condition was $\chi^2(9) = 47.3, p < .001$. The contribution from the delayed condition was $\chi^2(9) = 22.8, p > .005$. The contribution from the continuous-distractor condition was $\chi^2(9) = 68.8, p < .001$.

to predict the hypothesized and confirmed long-term lag recency effect.

The contextual variability models adequately characterize the qualitative pattern of end-of-list recency across conditions. They predict recency in immediate free recall, diminished recency in delayed free recall, and increased recency relative to delayed free recall in continuous-distractor free recall. In short, contextual variability SAM predicts LTR. The passive formulation, however, is incapable of predicting any lag recency effects at all. In contrast, the retrieved formulation predicts a lag recency effect in both delayed and continuous-distractor free recall. The finding of lag recency and long-term lag recency simultaneously with LTR is a basic property of the structural assumption of retrieved context and is not a consequence of the particular parameter values chosen.

Those features of the data that we found to be well described by the SAM–FR (end-of-list recency in immediate free recall and lag recency in delayed free recall) were also surprisingly well described by the retrieved-context model.

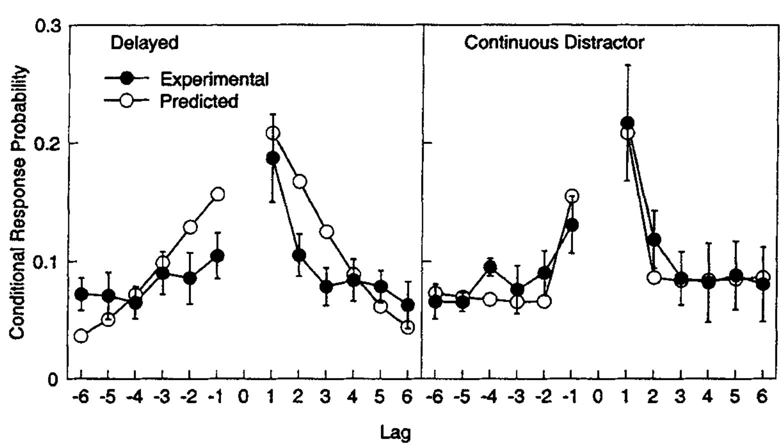

Figure 13. Contextual variability SAM (search of associative memory), retrieved formulation: conditional response probabilities. Retrieved, variable context predicts the existence of the lag recency effect in both delayed and continuous distractor free recall. The best fitting parameter values were $d = 6.98$ for the residual strength and $f = 1.40$ for the contextual asymmetry parameter. The other parameters were as in Figure 12. For the five conditions, $\chi^2(55) = 313.4$, $p < .001$. The individual chi-square analyses were both highly significant.

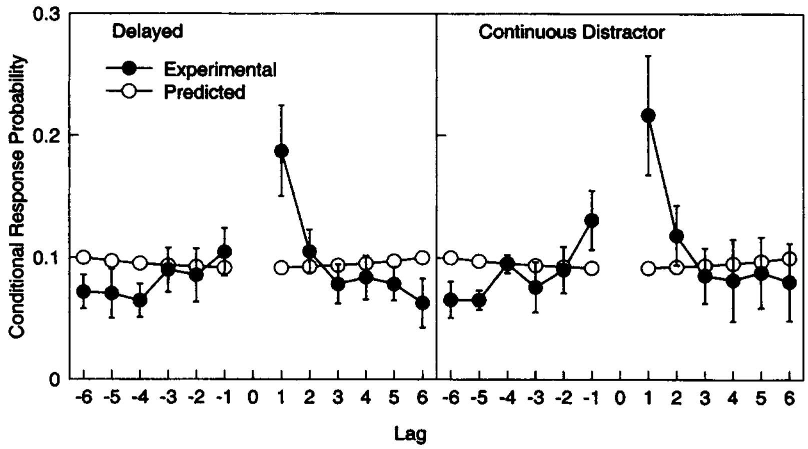

Figure 14. Contextual variability SAM (search of associative memory), passive formulation: conditional response probabilities. The passive formulation of variable context fails to predict the lag recency effect in either delayed or continuous-distractor free recall. None of the parameter values tried produced any lag recency effect at all. The parameter values used in the figure were $d = 8.1$ for the residual strength and $f = 0.80$ for the contextual asymmetry parameter. All other parameters were kept at the values listed for Figure 12. For the five conditions taken together, $\chi^2(55) = 301.0, p < .001$. The individual chi-square analyses were both highly significant.

It qualitatively described the end-of-list recency effect in immediate free recall and captured the qualitative nature of the lag recency effect in delayed free recall. This makes it a good candidate as an explanation of serial position effects in general rather than simply as an explanation of the long-term recency effect.

In delayed free recall, SAM–FR correctly predicted that recency would be attenuated but overestimated this effect. Rather than a failure of the model, might this residual recency reflect a failure of the data? Perhaps the end-of-list distractor was not completely successful in displacing items from STS. If this were true, then LTR could be explained as a result of the last item in the list remaining in STS until the time of test. As discussed in Appendix B, an explanation of LTR as a result of retrieval from STS hinges on there being more than one item available in STS at the time of test in delayed free recall. Because all items in STS are held to be equally available at the time of test, if there is more than one item available, each of those individual items is less likely to be recalled first than it would if it were the only item in STS. An interitem distractor flushes out prior items, leading to increased probability of first recall for the last item and an artifactual increase in recency in continuous-distractor free recall. For this explanation to hold, more than one item must be available from STS in delayed free recall. If this were true, then the CRP should change with output position, reflecting an STS component at early output positions, as in immediate free recall. This prediction is inconsistent with the finding that the CRP was unchanged with output position in delayed free recall (see Figure 4). An artifactual account of LTR by retrieval from STS is therefore inconsistent with our data.

Perhaps LTR does not reflect retrieval from STS but is nonetheless a consequence of the operation of STS during list presentation. If, for instance, we assume that items are displaced more slowly by distractors than by other list items, then items in STS at the start of the distractor period will spend more total time in STS. Items at the end of the list will therefore have a stronger association with fixed list context and will be more likely to be retrieved from LTS. This would clearly predict some recency effect in delayed free recall. This line of reasoning, however, cannot be extended to predict LTR. In continuous-distractor free recall, each of the items in the list is followed by a delay. Because retrieval from LTS is competitive, and all items have a similar benefit, we would not expect this mechanism to produce LTR.

The residual recency observed in delayed free recall is quite consistent with the contextual variability model. The contextual variability model predicts that there should be some (perhaps vanishingly small) recency effect in any list in which all other factors are equated. This is not at all inconsistent with the finding of negative recency in final free recall after an immediate test (Craik, 1970). Explanations of this negative recency hinge on more total study time for prerecency items. If there was little or no rehearsal for any of the items in our experiments, one would not expect to find such an effect. In contrast, experimental manipulations that encourage rehearsal could obscure the small recency effect predicted by contextual variability models such as the one under consideration here.

The notion of temporal distinctiveness (Nairne et al., 1997; Neath & Crowder, 1990) has figured prominently in previous attempts to describe LTR quantitatively. Our arguments in favor of a key role for retrieved variable context as an explanation of serial position effects, and free recall in general, should be seen as complementary to a distinctiveness-based approach. The only constraint that our data clearly place on the temporal distinctiveness models is that

Table 2
End-of-List and Lag Recency Effects in Free Recall (Including Results Presented in This Article)

Experimental condition	Recency effect	SAM–FR	Passive context	Retrieved context	Data
Immediate	End-of-list	Y	Y	Y	Y
	Lag	Y	N	Y	Y
Delayed	End-of-list	N	N	N	N
	Lag	Y	N	Y	Y
Continuous distractor	End-of-list	N	Y	Y	Y
	Lag	N	N	Y	Y

Note. Search of associative memory—free recall (SAM–FR) gives the predictions of the Raaijmakers and Shiffrin (1980, 1981) version of the SAM that relies on a short-term store and fixed list context. Passive context gives the predictions of a version of SAM based exclusively on variable context, without retrieval of context. Retrieved context gives the predictions of a version of SAM based on variable context that is retrieved and used as a cue for subsequent recalls. Entries for lag indicate whether a lag recency effect is observed (i.e., whether the conditional response probability is graded or not).

they be cast in such a way as to explain the lag recency effects, as reflected in the CRP. Retrieved variable context should be seen as one construct that could underlie such a formulation of temporal distinctiveness.

General Discussion

The shape of the serial position curve is largely the result of the probability of first recall and the CRP. The first retrieval is described by the probability of first recall, and subsequent retrievals are described by the CRP. Thus, a single mechanism that accounts for these two functions largely characterizes the retrieval process in free recall.[10]

The recency effect in single-trial free recall is seen clearly in the probability of first recall—participants begin recall at the end of the list. Serial position can be thought of as a lag measured from the end of the list. Probability of first recall is then the special case of a CRP when no prior items have been recalled. The end-of-list recency effect means that recall is high for small values of this end-of-list lag. Analogously, the existence of a graded CRP is very much like the recency effect, only recency in this case is measured relative to the item just recalled rather than relative to the end of the list. To emphasize this equivalence, we have referred to the finding of a graded CRP as the lag recency effect.

We examined immediate, delayed, and continuous-distractor free recall by using the analytic framework of the probability of first recall and the CRP. In immediate free recall, recency was seen in the sharply accelerating probability of first recall (see Figure 3) and lag recency was seen in the CRP (see Figure 4). In delayed free recall, the lag recency effect was intact (see Figure 3), whereas end-of-list recency, as measured by the probability of first recall, was attenuated. The CRP in the continuous-distractor condition was similar to that in delayed free recall (see Figure 9). However, although the probability of first recall in the continuous-distractor condition was significantly different from that in delayed free recall (see Figure 8), it was similar to that in immediate free recall (see Figures 3 and 8). It then follows that the long-term recency effect (i.e., the difference between continuous-distractor free recall and delayed free recall) was entirely a consequence of the enhanced recency in the probability of first recall observed in the continuous-distractor condition. Empirically, the end-of-list recency effect, as measured by the probability of first recall, varied as a function of the relative spacing of items and the time of test.

We examined the ability of two kinds of models to account for our data: the Raaijmakers and Shiffrin (1980, 1981) SAM–FR model with STS and fixed list context and two variants of SAM that are based on Mensink and Raaijmakers's (1988, 1989) contextual variability. According to one of these variants, the passive formulation, context at the time of test is used throughout recall. In the retrieved-context formulation, the context of a studied item is reinstated when that item is recalled. The retrieved-context formulation predicts the qualitative pattern of results in the probability of first recall and the CRP across all three conditions. Retrieved variable context is a good candidate as the temporally sensitive construct that causes serial position effects in free recall.

Table 2 summarizes what we know about serial position effects in free recall after inclusion of the results in this article. Empirically, the findings in this article indicate that the lag recency effect is much more general than what might have been thought previously—it is found in delayed and continuous-distractor free recall as well as in immediate free recall.

Theoretically, the present results demand that Raaijmakers and Shiffrin's (1980, 1981) SAM–FR is at best incomplete—some mechanism other than STS must give rise to end-of-list and lag recency effects in continuous-distractor free recall. This can be seen by noting that the SAM–FR

[10] The other main characteristic of the serial position curve that is not described by these two functions is the level of asymptotic recall. This is described empirically by the growth of interresponse times.

model fails to predict both end-of-list and lag recency effects in continuous-distractor free recall. Both of these effects were found in our study.

The modeling results in this article indicate that the contextual variability SAM model of Mensink and Raaijmakers (1988, 1989) correctly predicted the qualitative pattern of end-of-list recency across all three distractor conditions. A comparison of the predictions for the lag recency effect in continuous-distractor free recall with the data showed that only the retrieved-context model can hold. The existence of a long-term lag recency effect implies that we should reject a purely passive formulation and points to a central role for retrieved variable context in free recall.

This particular model of contextual variability is by no means the only one possible, and there is no guarantee that a similar model cannot do a better job in describing the data quantitatively. Similarly, it is possible that STS operates when items are close together in time (as in immediate free recall and within the list in delayed free recall), whereas contextual variability dominates when events are separated in time. However, parsimony demands that we pursue the possibility that free recall is described by one process—that process being recall mediated by retrieved variable context.

Conclusion

The serial position curve in free recall results from the joint operation of probability of first recall and CRP curves. These measures provide information not derivable from the serial position curve, yet the serial position curve can largely be characterized by these two measures. Experiment 2 exhibited a long-term lag recency effect in continuous-distractor free recall. This and the finding of long-term recency under conditions designed to minimize rehearsal strongly suggest that neither end-of-list recency nor lag recency depend critically on rehearsal. These findings indicate that accounts of free recall based solely on rehearsal and fixed list context (e.g., Raaijmakers & Shiffrin, 1980, 1981) are at best incomplete. An account of free recall that assumes that variable context (based on Mensink & Raaijmakers, 1988, 1989) is used as an initial retrieval cue describes the qualitative pattern of end-of-list recency across delay conditions, including immediate free recall. The finding of a long-term lag recency effect is consistent with variable context that is retrieved and used as a cue for all retrievals. Retrieved variable context is thus a viable explanation for serial position effects in immediate, delayed, and continuous-distractor free recall.

References

Anderson, J. R. (1976). *Language, memory and thought.* Hillsdale, NJ: Erlbaum.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, pp. 89–105). New York: Academic Press.

Baddeley, A. D. (1986). *Working memory.* Oxford, England: Clarendon Press.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 8, pp. 47–90). New York: Academic Press.

Baddeley, A. D., & Hitch, G. J. (1977). Recency reexamined. In S. Dornic (Ed.), *Attention and performance VI* (pp. 647–667). Hillsdale, NJ: Erlbaum.

Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology, 6,* 173–189.

Craik, F. I. M. (1970). The fate of primary memory items in free recall. *Journal of Verbal Learning and Verbal Behavior, 9,* 658–664.

Crowder, R. G. (1982). The demise of short-term memory. *Acta Psychologica, 50,* 291–323.

Deese, J., & Kaufman, R. A. (1957). Serial effects in recall of unorganized and sequentially organized verbal material. *Journal of Experimental Psychology, 54,* 180–187.

Friendly, M., Franklin, P. E., Hoffman, D., & Rubin, D. C. (1982). Norms for the Toronto Word Pool. *Behavior Research Methods and Instrumentation, 14,* 375–399.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review, 91,* 1–67.

Glanzer, M., & Cunitz, A. R. (1966). Two storage mechanisms in free recall. *Journal of Verbal Learning and Verbal Behavior, 5,* 351–360.

Glenberg, A. M., Bradley, M. M., Kraus, T. A., & Renzaglia, G. J. (1983). Studies of the long-term recency effect: Support for a contextually guided retrieval theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 413–418.

Glenberg, A. M., Bradley, M. M., Stevenson, J. A., Kraus, T. A., Tkachuk, M. J., & Gretz, A. L. (1980). A two-process account of long-term serial position effects. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 355–369.

Glenberg, A. M., & Swanson, N. G. (1986). A temporal distinctiveness theory of recency and modality effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 3–15.

Greene, R. L. (1986a). A common basis for recency effects in immediate and delayed recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 413–418.

Greene, R. L. (1986b). Sources of recency effects in free recall. *Psychological Bulletin, 99,* 221–228.

Greene, R. L., & Crowder, R. G. (1984). Effects of semantic similarity on long-term recency. *American Journal of Psychology, 97,* 441–449.

Healy, A. F., & McNamara, D. S. (1996). Verbal learning and memory: Does the modal model still work? *Annual Review of Psychology, 47,* 143–172.

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition, 24,* 103–109.

Koppenaal, L., & Glanzer, M. (1990). An examination of the continuous distractor task and the long-term recency effect. *Memory & Cognition, 18,* 183–195.

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

Marshall, P. H., & Werder, P. R. (1972). The effects of the elimination of rehearsal on primacy and recency. *Journal of Verbal Learning and Verbal Behavior, 11,* 649–653.

Mensink, G.-J. M., & Raaijmakers, J. G. W. (1988). A model for interference and forgetting. *Psychological Review, 95,* 434–455.

Mensink, G.-J. M., & Raaijmakers, J. G. W. (1989). A model for contextual fluctuation. *Journal of Mathematical Psychology, 33,* 172–186.

Metcalfe, J., & Murdock, B. B. (1981). An encoding and retrieval model of single-trial free recall. *Journal of Verbal Learning and Verbal Behavior, 20,* 161–189.

Murdock, B. B. (1960). The distinctiveness of stimuli. *Psychological Review, 67,* 16–31.

Murdock, B. B. (1962). The serial position curve in free recall. *Journal of Experimental Psychology, 64,* 482–488.

Murdock, B. B. (1967). Recent developments in short-term memory. *British Journal of Psychology, 58,* 421–433.

Murdock, B. B., & Okāda, R. (1970). Interresponse times in single-trial free recall. *Journal of Verbal Learning and Verbal Behavior, 86,* 263–267.

Murdock, B. B., & Walker, K. (1969). Modality effects in free recall. *Journal of Verbal Learning and Verbal Behavior, 8,* 665–676.

Nairne, J. S. (1992). *Position-based search of associate memory.* Unpublished manuscript.

Nairne, J. S., Neath, I., Serra, M., & Byun, E. (1997). Positional distinctiveness and the ratio rule in free recall. *Journal of Memory and Language, 37,* 155–166.

Neath, I. (1993). Contextual and distinctive processes and the serial position function. *Journal of Memory and Language, 32,* 820–840.

Neath, I., & Crowder, R. (1990). Schedules of presentation and temporal distinctiveness in human memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 316–327.

Neath, I., & Crowder, R. G. (1996). Distinctiveness and very short-term serial position effects. *Memory, 4,* 225–242.

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal, 7,* 308–313.

Phillips, J. L., Shiffrin, R. J., & Atkinson, R. C. (1967). The effects of list length on short-term memory. *Journal of Verbal Learning and Verbal Behavior, 6,* 303–311.

Poltrock, S. E., & MacLeod, C. M. (1977). Primacy and recency in the continuous distractor paradigm. *Journal of Experimental Psychology: Human Learning and Memory, 3,* 560–571.

Postman, L., & Phillips, L. (1965). Short-term temporal changes in free recall. *Quarterly Journal of Experimental Psychology, 17,* 132–138.

Raaijmakers, J. G. W. (1993). The story of the two-store model of memory: Past criticisms, current status, and future directions. In D. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 467–488). Cambridge, MA: MIT Press.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 14, pp. 207–262). New York: Academic Press.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88,* 93–134.

Raymond, B. J. (1969). Short-term storage and long-term storage in free recall. *Journal of Verbal Learning and Verbal Behavior, 8,* 567–574.

Roberts, W. A. (1972). Free recall of word lists varying in length and rate of presentation: A test of total-time hypotheses. *Journal of Experimental Psychology, 92,* 365–372.

Shiffrin, R. M. (1970). Forgetting: Trace erosion or retrieval failure? *Science, 168,* 1601–1603.

Strong, E. K. J. (1912). The effect of length of series upon recognition memory. *Psychological Review, 19,* 447–462.

Thapar, A., & Greene, R. (1993). Evidence against a short-term-store account of long-term recency effects. *Memory & Cognition, 21,* 329–337.

Tulving, E., & Patterson, R. (1968). Functional units and retrieval processes in free recall. *Journal of Experimental Psychology, 77,* 239–248.

Tzeng, O. J. L. (1973). Positive recency in delayed free recall. *Journal of Verbal Learning and Verbal Behavior, 12,* 436–439.

Watkins, M. J. (1974). The concept and measurement of primary memory. *Psychological Bulletin, 81,* 695–711.

Watkins, M. J., Neath, I., & Sechler, E. S. (1989). Recency effect in recall of a word list when an immediate memory task is performed after each word presentation. *American Journal of Psychology, 102,* 265–270.

Watkins, M. J., Watkins, O. C., & Crowder, R. G. (1974). The modality effect in free and serial recall as a function of phonological similarity. *Journal of Verbal Learning and Verbal Behavior, 13,* 430–447.

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

Appendix A

Empirical Analyses

The serial position curve plots the probability of recalling an item from each serial position without regard to output position. The probability of first recall is a serial position curve calculated for only the very first item that is recalled. The sum of the probability of first recall curve, across serial positions, may be slightly less than 1.0. This is because participants may recall an intrusion as the first item or may fail to make any response in the time allotted for recall.

Kahana (1996) introduced another measure of primary organization in free recall; for two items recalled successively, the CRP measures the tendency for successively recalled items to come from nearby serial positions. The CRP is plotted as a function of lag, where lag is the difference between the serial positions of the successively recalled items. The greatest possible lag in a list of N items is $N - 1$ (recall of the first item in the list followed by recall of the last item in the list); the smallest possible lag is $-(N - 1)$; the figures in this article do not plot all possible lags.

For a given lag, k, CRP(k) is defined as the number of successive recalls of pairs with lag k divided by the maximum number of times that pairs with lag k could have been recalled. Let us refer to the numerator of this expression as $n(k)$ and the denominator as $d(k)$,

$$\mathrm{CRP}(i + k|i) \equiv \frac{n(k)}{d(k)}.$$

If the just-recalled item is the last word in the list, there is no way that the participant could have recalled an item that would lead to a lag of +1. The denominators are only incremented for possible lags given the serial position of the previously recalled word. The CRP provides information that is not contained in the serial position curve or in the probability of first recall. Which lags are possible depends on the serial position of the just-recalled item. Had the CRP been calculated without regard to possible lags, it would have been largely redundant with the serial position curve.

Care should be taken when working with the CRP to avoid collapsing over a variable that has an effect on the CRP. The CRP changes substantially with output position in immediate free recall but not in delayed free recall. In his analyses, Kahana (1996) avoided this problem by omitting the first three output positions. The CRP does not vary substantially with serial position, except insofar as serial position is confounded with output position. An exception to this is recall to the very first serial position. D. R. Laming (personal communication, December 16, 1996) has shown that participants have a strong tendency to make transitions in recall from interior list positions to the very first serial position in the list. These transitions are more frequent than transitions of equivalent lag. Our unpublished secondary analyses have confirmed this result. This effect is analogous to the one-position primacy effect seen in the probability of first recall in this study and others (Laming, personal communication, December 16, 1996). Omitting recalls to the first serial position did not substantially affect our CRP curves.

Appendix B

Modeling

In implementing free recall within the framework of SAM (e.g., Mensink & Raaijmakers, 1988; Raaijmakers & Shiffrin, 1980, 1981), our focus has been on the structural assumptions of the models. Rather than implementing the full 11-parameter version of SAM–FR (e.g., Kahana, 1996) and then adding parameters for contextual variability, we opted for a simplified analytic treatment. As we have argued in the text of this article, the serial position curve largely results from the joint operation of the CRP and probability of first recall. These two measures, in addition to providing additional information not apparent from the serial position curve, express relative, rather than absolute, probabilities of recall. This makes them much easier to treat analytically than the serial position curve would have been.

The Raaijmakers and Shiffrin (1980, 1981) two-store model of free recall (SAM–FR) relies on the operation of the STS to produce end-of-list and lag recency effects (see text for details). In contrast, using contextual variability, as proposed by Mensink and Raaijmakers (1988, 1989), we can generate the end-of-list recency effect in free recall without relying on STS. In our implementation of free recall with contextual variability, we assumed no contribution from STS; rather, context serves as the sole retrieval cue at the start of recall. We go further to contrast two different uses of contextual variability: In the passive-context formulation, context at the time of test serves as the retrieval cue throughout the recall process. In the retrieved-context formulation, end-of-list context is only used as the cue for the first recalled item. Recall of an item then reinstates the context that was associated with it during study. This retrieved context serves as the cue for the next retrieval attempt. In all of the models analyzed here, we used the same sampling and recovery process for recall from LTS.

STS

STS occupancy affects recall in three ways in the Raaijmakers and Shiffrin (1980, 1981) SAM model. First, items in STS at the time of test are recalled initially, leading to an end-of-list recency effect in the probability of first recall. Second, an item i is more likely to be recalled from LTS, with fixed list context as the cue, as a function of the time it has spent in STS over the course of the experiment, t_i. Third, the associative strength between two items i and j is determined by the amount of time the two items spend together in STS, t_{ij}. This associative strength gives rise to the lag recency effect—when item j has just been recalled, it contributes to the cue for the next recall.

In immediate free recall, participants typically do not begin recall at the very end of the list but start a couple of items back and then move forward to the end of the list. This tendency results in a "humped" probability of first recall (see Figure 1, Footnote 2). To explain this, Kahana (1996) found it necessary to use a dropout rule from STS in which older items are more likely to be displaced than newer items (as introduced in Phillips, Shiffrin, & Atkinson, 1967). To explain our observation of a positively accelerated probability of first recall in Experiment 1, we assumed a random dropout rule, with all items in STS equally available at the time of test.

Retrieval From STS

In immediate free recall, the probability that an item i remains in STS at time step j is given by

$$B_{ij} = \begin{cases} (1 - r^{-1})^{j-i} & r < i \le j \\ 1 & i \le j \le r, \end{cases} \tag{B1}$$

where r is the capacity of STS. The matrix, B, is of dimension $[L \times (L + 1)]$, where L is the number of items in the list, and $L + 1$ is the number of items plus the time of test. B is sufficient to generate average values for buffer occupancy at the time of test, the amount of time a given item spends in the buffer t_i, and the joint time two items i and j spend in the buffer, t_{ij}.

This analytic treatment introduces a couple of subtle deviations from simulation studies of SAM. The use of the matrix B treats buffer occupancy of multiple items as independent events. In fact, this is inconsistent with the dependency that would obtain in a simulation—if the capacity of the buffer is two, and if it is known that items x and y are in the buffer, then z is not. This simplification is unlikely to affect the conclusions in any significant way. The use of this analytic treatment makes it formally unnecessary to require r to be an integer. Suppose that the size of the buffer is not constant across subjects or across trials. A continuous-valued r could then be derived from this distribution of (integral) buffer sizes.

Let D denote the length of the distractor intervals in delayed and continuous-distractor free recall. In the continuous-distractor condition,

$$B_{ij} = (1 - r^{-1})^{D(j-i)}; \quad i \leq j. \tag{B2}$$

In delayed free recall, D only appears in the entries at the time of test, with $j = L + 1 \equiv T$, as

$$B_{iT} = (1 - r^{-1})^{L-i+D}. \tag{B3}$$

The probability of first recall from STS of any item i (in all three conditions) is then the probability that the specific item is in STS at the time of test, divided by the sum of the probabilities, if there is on average more than one item at test.

$$P_{\text{STS}}(i) = \begin{cases} \dfrac{B_{iT}}{\sum_j B_{jT}} & \sum_j B_{jT} \geq 1 \\ B_{iT} & \sum_j B_{jT} < 1 \end{cases} \tag{B4}$$

Note that this allows for an STS-based long-term recency effect in the probability of first recall if there is more than one item available from STS in delayed free recall. The probability that the last item in the list is in the buffer at test is the same number, B_{LT}, in both the delayed and continuous-distractor conditions (because j in Equation B2 is equal to $T \equiv L + 1$). The probability that this item will be recalled first goes down with the number of other items available in STS. Adding a distractor before the last item decreases this competition from the other items. This effect depends on $\Sigma_j B_{jT}$ being greater than 1 in delayed free recall and has an upper limit determined by the value of $\Sigma_j(B_{jT}) - 1$ in delayed free recall.

Effect of STS on Retrieval From LTS

The item-to-context and item-to-item strengths in the Raaijmakers and Shiffrin (1980, 1981) model make use of the following simple relationships for t_i and t_{ij}.

$$t_i = \sum_j B_{ij},$$

$$t_{ij} = \sum_k B_{ik} B_{jk}.$$

In the fixed context (STS) model, we set the context-to-item strength $S(I_i, C)$ to be at_i and the item-to-item strength to be

$$S(I_i, I_j) = \begin{cases} bt_{ij} & t_{ij} \geq 1 \\ bt_{ij} + d(1 - t_{ij}) & t_{ij} < 1, \end{cases} \tag{B5}$$

where a, b, and d have the same meaning as in Raaijmakers and Shiffrin (1981). To generate the observed asymmetry in the CRP, we set $b = b_F$; $i > j$ for forward recalls and $b = b_B$, $i < j$ for backward recalls, as in Gillund and Shiffrin (1984) and Kahana (1996).

Context and Variability

For the contextual variability models, we set

$$B_{ij} = \delta_{ij},$$

(where δ is the Kronecker delta function) so that $S(I_i, I_j) = d$ for all $i \neq j$ and $P_{\text{STS}}(i) = 0$ for all i and all conditions (including immediate free recall). Thus, there is no contribution from the operation of the STS. Following Mensink and Raaijmakers (1988, 1989), we calculated the change in contextual overlap for item presentations separated by time τ as

$$A(\tau) = A(0) \exp[-(\beta + \gamma)\tau] + K\left(\frac{\gamma}{\beta + \gamma}\right)\{1 - \exp[-(\beta + \gamma)\tau]\}. \tag{B6}$$

We used this equation to generate a matrix, A, of the overlaps between the context at any two times

$$\tau_{ij} = (1 + \text{IPI})|i - j|$$

$$A_{ij} = A(\tau_{ij}) \tag{B7}$$

$$A_{iT} = A[\text{RI} + (1 + \text{IPI})(L - i)],$$

where τ_{ij} just expresses the time between the presentation of item i and item j (with IPI and RI = 0 or D as appropriate for the condition) and $A(0) = 1$. For the probability of first recall, $S(I_i, C) = A_{iT}$.

For the CRP, we tried two variants. In the passive-context model, we used A_{iT} for all i, j. In the retrieved-context model, we used $S(I_i, C) = A_{ij}$ as the context cue for item i following recall of j, implying that the context of item j has been retrieved. To generate an asymmetry in the contextually mediated CRPs, we multiplied the $S(I_i, C)$ terms in Equation B8 by a new free parameter f if $j < i$, providing an advantage for forward recalls if $f > 1$.

Retrieval From LTS

This section applies to both the Raaijmakers and Shiffrin (1980, 1981) and contextual variability SAM models. In SAM, recall of an item from LTS proceeds in two steps: sampling and recovery. Retrieval of subsequent items proceeds first with context and the contents of STS (in particular, the previously recalled item) as cues. After $L_{\max}$ failed attempts, retrieval proceeds with context only as a cue. If $K_{\max}$ failed attempts are reached, recall stops. Because our measures, the probability of first recall and the CRP, imply that something is recovered, we are interested in the relative probabilities of sampling and recovering different items. The possibility that nothing is recovered is not relevant in calculating these statistics—$K_{\max}$ plays no role. The empirical observation that the CRP does not change with output position led us to assume that the stages of sampling with item and context, followed by context alone, do not play a significant role in describing our data. In this case, $L_{\max}$ does not play a role either. Because the CRP assumes that an item different from the one just recalled is recovered, the self-strength of items does not play a role.

In SAM, the probability of sampling an item from LTS on a given sampling attempt is the (multiplicative) strength of those cues with the item divided by the sum of the strength of those cues to all the other items in the list. When no items have been recalled, context is the only cue. When an item has been recalled, it is used as a supplemental cue for further recalls. We did not model the effect of multiple cues in calculating the CRP. If multiple cues contributed to the CRP, we would expect that the CRP would change shape with output position in delayed free recall as more items have been retrieved. Examination of Figure 4 demonstrates that there is no such effect, indicating that multiple cues do not have much effect on the CRP.

If an item has been recalled, the probability of sampling an item i, given that j has just been recalled, is

$$P_S(i|j) = \frac{S(I_i, C)S(I_i, I_j)}{\sum_k S(I_k, C)S(I_k, I_j)},$$

where j is the item just recalled. Given that item i has just been sampled, its probability of successful recovery is given by

$$P_R(i) = 1 - \exp\left[-S(I_i, C) - S(I_i, I_j)\right].$$

The probability of recalling an item at a given sampling attempt is the joint probability that an item is sampled and is subsequently recovered successfully:

$$P_S(i|j) \times P_R(i) = \frac{S(I_i, C)S(I_i, I_j)\{1 - \exp\left[-S(I_i, C) - S(I_i, I_j)\right]\}}{\sum_k S(I_k, C)S(I_k, I_j)}.$$

Let us neglect the rule that an item may not be recovered with a context cue that has previously failed to recover it. The measures we are interested in give the probability that an item is recalled relative to other items in the list. Because the probability of recalling each item at each sampling attempt has the same sum in the denominator, the relative probability of recall is just this joint probability (normalized appropriately). For the first recall, we have

$$P_{SR}(i) = \frac{S(I_i, C)\{1 - \exp\left[-S(I_i, C)\right]\}}{\sum_j S(I_j, C)\{1 - \exp\left[-S(I_j, C)\right]\}}. \quad \text{(B8)}$$

For the CRP, we have, for each item j in the list,

$$P_{SR}(i|j) = \frac{S(I_i, I_j)S(I_i, C)\{1 - \exp\left[-S(I_i, I_j) - S(I_i, C)\right]\}}{\sum_k S(I_k, I_j)S(I_k, C)\{1 - \exp\left[-S(I_k, I_j) - S(I_k, C)\right]\}}. \quad \text{(B9)}$$

The observed CRP is derived from $P_{SR}(i|j)$. We calculated the CRP as

$$P_{\mathrm{CRP}}(\mathrm{lag}) = \sum_{i:1 \le i+\mathrm{lag} \le L} \frac{P_{SR}(i + \mathrm{lag}|i)}{L(L - |\mathrm{lag}| - 1)}, \quad \text{(B10)}$$

where $(1/L)$ is taken as the (constant) probability that item i has just been recalled and $L - |\mathrm{lag}| - 1$ is the number of times that lag can occur. For instance, a lag of -11 can only occur if the previous word recalled was from Position 12. Again, the consistency of the CRP over output positions in delayed free recall indicates that serial position effects do not play an important role in determining the CRP in delayed or continuous-distractor free recall, justifying the neglect of serial position effects in the CRP. Taking the sum weighted by the probability of first recall of each item or simply calculating the CRP from one interior list position did not change any of the conclusions arrived at in this article.

Received November 28, 1997
Revision received January 13, 1999
Accepted January 18, 1999 ■