Psychological Review
1994, Vol. 101, No. 3, 528-533


0033-295X/94/$3.00

# THEORETICAL NOTES

# Critique of the Retrieval/Deblurring Assumptions of the Theory of Distributed Associative Memory

James S. Nairne and Ian Neath

This note critiques the retrieval assumptions proposed by Lewandowsky and Murdock (1989) in their application of Murdock's theory of distributed associative memory (TODAM) to problems in serial-order recall. Two different methods for "deblurring" the products of retrieval are described, along with simulations appropriate for each. The authors demonstrate that the deblurring assumptions used by Lewandowsky and Murdock (1989) in their simulation version of the model, although appropriate for some aspects of serial-anticipation learning, provide an inadequate account of general serial recall and, in fact, predict several trends that are inconsistent with known data. The authors also briefly consider the analytic version of the model and note some additional problems. Overall, the deblurring assumptions critically affect the performance of TODAM, and it is these assumptions rather than the model per se that determine a sizable proportion of its behavior.

The dynamics of serial-order recall have been of concern to psychologists since before the seminal work of Ebbinghaus (1885/1964). Traditionally, memory researchers have opted for either a chaining view, in which associations between successive items form the basis for sequential recall, or a view that relies on the formation of position-to-item associations (for reviews, see Crowder, 1976; Ebenholtz, 1972; Slamecka, 1985; Young, 1968). In their recent extension of Murdock's (1982, 1983) theory of distributed associative memory (TODAM) to serial-order memory, Lewandowsky and Murdock (1989) proposed a revised chaining model and provided quantitative fits to data for a variety of historically relevant benchmark paradigms.

In the serial-order version of TODAM, events are represented as vectors with individual elements drawn from zero-centered, normally distributed random variables. Successive events are associated through the mathematical operation of convolution, and both the item vectors (e.g., $f_j$) and the convolution products (e.g., $f_j * f_{j-1}$) are stored in a common memory vector, **M**. At the point where item $j$ is encoded (Equation 1 from Lewandowsky & Murdock, 1989):

$$\mathbf{M}_j = \alpha\mathbf{M}_{j-1} + \gamma_j f_j + \omega_j f_j * f_{j-1}. \quad (1)$$

Storage of the item and associative information is weighted by the parameters $\gamma$ and $\omega$, respectively, and the common vector **M** is "forgotten" before each new encoding by a weighting factor $\alpha$. Parameters $\alpha$, $\gamma$, and $\omega$ are real numbers that assume values between 0 and 1.

Order information is recovered by cueing **M** with an item (e.g., $f_{j-1}$) through the mathematical operation of correlation, and the product is a noisy, but potentially recoverable, version of item $j$ (Equation 2a from Lewandowsky & Murdock, 1989):

$$f_{j-1} \# \mathbf{M} = f'_j. \quad (2)$$

The subject is assumed to probe the memory vector successively using either the correlation product ($f'_j$) or the original item ($f_j$) as input for the next retrieval (for details, see Lewandowsky & Murdock, 1989). Before the actual recall, the blurry product of correlation, $f'_j$, needs to be converted (or "deblurred") into a recallable form; the assumptions underlying this deblurring process turn out to be critical determinants of performance and serve as the focus of our commentary.

According to Lewandowsky and Murdock (1989), correct recall of item $f_j$, given $f'_j$, depends on two factors. First, the similarity between $f'_j$ and $f_j$, defined as the magnitude of the dot product between the two vectors, must fall within acceptable limits (defined as tolerance limits; see Lewandowsky & Murdock, 1989, p. 31). Second, $f'_j$ must be more similar to the correct item ($f_j$) than to any other member of the available recall set. If $f'_j$ turns out to be more similar to another list item (e.g., $f_{j+1}$), then $f_j$ is not recalled, and performance is scored as incorrect. For liberal tolerance limits, then, overall performance is largely dependent on the composition of the competitor set.

Consider the extreme case in which the comparison set contains only the correct to-be-recalled item. Assuming that the similarity of $f'_j$ to $f_j$ falls within the acceptable limits, the subject will always interpret $f'_j$ correctly, because there is only one item from which to choose, and performance will be perfect. As the number of possible competitors increases, performance declines and the number of intralist intrusions grows. The charac-

We thank Stephan Lewandowsky for providing the source code for the open-loop implementation of TODAM and for assistance with the simulations. We also thank Stephan Lewandowsky and Richard Shiffrin for helpful criticisms of an earlier draft of this article.

Correspondence concerning this article should be addressed to James S. Nairne or Ian Neath, Department of Psychological Sciences, 1364 Psychological Sciences Building, Purdue University, West Lafayette, Indiana 47907-1364. Electronic mail may be sent to nairne@psych.purdue.edu or neath@psych.purdue.edu.

teristics of the competitor set and how the composition of that set changes during recall of the sequence are therefore important determinants of both the overall level of performance and the resultant shape of the serial-position curve. Note that this conclusion is independent of any specific assumptions about the processes of convolution and correlation and, in fact, applies to a range of simulation and neural-net models. The characteristics of the competitor set will be important anytime it is necessary to interpret a noisy trace through comparisons with a set of standards (see also Hintzman, 1986; Nairne, 1990a; Raaijmakers & Shiffrin, 1981).

In the simulations that follow, we consider two different methods mentioned by Lewandowsky and Murdock (1989) for determining how the composition of the competitor set unfolds throughout recall.[1] We test these methods using the simulation version of the open-loop model because that version of TODAM provides the most complete implementation of the psychological mechanisms proposed by Lewandowsky and Murdock (1989). In a later section, we extend our arguments to cover the assumptions of the analytic version of TODAM, which is designed to compute probability correct directly (and, in fact, does not specify the exact identity of the competitors). We demonstrate that both of the simulation-based deblurring techniques, including the one actually implemented by Lewandowsky and Murdock (1989), produce trends that are inconsistent with known empirical data.

## Deblurring in the Simulation Model

Lewandowsky and Murdock (1989) first mentioned a sampling without replacement scheme in which the size of the deblurring competitor set is reduced systematically after each retrieval attempt:

> The exact functional relation between *m* (the number of competitors) and output position could vary across paradigms and was determined by the parameter N (the Greek nu as distinct from *N*). In particular, the value of *m* was always equal to the number of as yet unrecalled items plus an integer N. Consider a five-item list with N set to zero. When the first item is retrieved, the number of competitors (and therefore *m*) is 4. Similarly, *m* is 3, 2, 1, and 0 for the next four retrievals. Now, consider the same list with N set to 1. The values of *m* across output position are 5, 4, 3, 2, and 1, respectively, which includes the boundary marker in the ensemble of possible output items. Conceptually, the parameter N was designed to reflect differences in the availability of different types of items in serial tasks. *It appeared the simplest and the most sensible implementation of a sampling without replacement scheme, to take into account the fact that subjects almost never commit intrusions of items that have already been recalled* [italics added]. (pp. 33–34)

One interpretation of this excerpt is that when an item is recalled, it is simply removed from the competitor set for the remainder of the trial. Such a scheme would prevent intralist intrusions of already recalled items because those items would no longer be allowable candidates for recall. However, as we shall see, TODAM is unable to produce the benchmark serial-position curves of interest with this scheme. Each time the model samples for recall, beginning with the attempted recall of the first serial position, there is a finite probability that the last item in the list will be sampled and, as a consequence, be removed from the competitor set. Thus, if the subject incorrectly intrudes Item 8 from an eight-item list in an early serial position, there is no possibility that the subject can then correctly recall Item 8 on that trial. Because the likelihood of sampling an item on a trial increases with each sample, end-of-the-list items should be at a disadvantage.

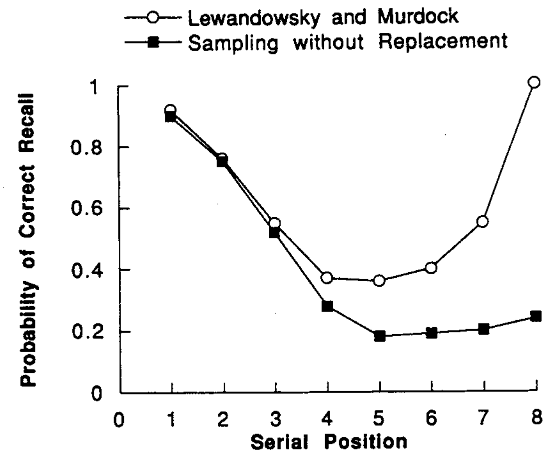


*Figure 1.* Simulation of serial recall of an eight-item list using two different deblurring procedures.

In fact, as the simulation shown in Figure 1 demonstrates, a straightforward sampling without replacement deblurring method (closed squares) produces no evidence of significant recency in the serial-position curve, even with N set to zero. This serial-position curve was produced with a simulation version of the open-loop implementation of TODAM using representative parameter settings from Lewandowsky and Murdock (1989, p. 33; for details about the simulation, see our Appendix). Although we present the data from only one simulation, across a wide variety of simulations, including changes in both list length and the dimensionality of the memory vector, we have been unable to produce serial-position functions of the type described by Lewandowsky and Murdock (1989) using this sampling without replacement technique.

We now consider the second method discussed by Lewandowsky and Murdock (1989). It also uses a sampling without replacement scheme, but it is the item that should have been recalled (i.e., the correct item) that is removed from the competitor set rather than the item that was actually recalled:

[1] There is at least one possible consequence of changing the sampling procedure that warrants mentioning. Once an item has been recalled, it is added back into the memory vector in the following way: "The retrieved item is convolved with the probe, and that convolution plus the recalled item are added to **M**. This implements output interference and also explains how it is that subjects can remember their earlier responses" (Lewandowsky & Murdock, 1989, p. 31). Thus, a change in the sampling procedure could result in a different item being added back into **M** and altering the output interference. However, while the two sampling procedures discussed in this article differ in how they affect the composition of the competitor set, both add the same item back into the memory vector. Our conclusions, then, should not be affected by changes in output interference.

> Taking into account the observation that subjects never repeat an item they have already recalled, competitors were taken from serial positions following the to-be-recalled item. Thus, the retrieved vector ($f'_j$) was compared (by computation of the dot product) with all items following the probe. (p. 35)

For example, for an eight-item list, when the common vector **M** is cued with $f_4$ (or $f'_4$), the correlation product $f'_5$ is compared with a competitor set containing only the intact vectors for presented Items 5, 6, 7, and 8 (assuming N = 0). Regardless of which of these four possible items is recalled, it will be Item 5 that is removed from the competitor set before the next retrieval attempt. Reduction of the competitor set is therefore independent of what was actually recalled; the composition of the competitor set is determined solely by what positions remain to be recalled. As demonstrated in Figure 1, this method produces a bow-shaped serial-position curve with a sharp recency effect (open circles; compare with Figure 7 of Lewandowsky & Murdock, 1989, p. 34). The occurrence of recency with this deblurring scheme is not surprising because end-of-the-list items have few, if any, competitors, and the comparison set will always contain the correct to-be-recalled item.

Although a bow-shaped serial-position curve is produced with the Lewandowsky and Murdock (1989) scheme, there are a number of associated costs. First, there is no plausible psychological mechanism proposed to accomplish the reduction in competitors: How could a system fail to determine the correct item for recall but nonetheless successfully identify that item for removal from the competitor set?

Second, this technique requires the introduction of a new ad hoc parameter, N, as a way of manipulating the size of recency obtained. When N is set to 0 (as in Figure 1, open circles), performance is almost perfect for the last item because the correlation product $f'_n$ has only one possible interpretation: There is only one item left in the competitor set and that is the item that was actually presented last in the list. Errors can occur in this position only if the similarity value between $f'_n$ and $f_n$ fails to fall within the tolerance limits. For this reason, the size of the competitor set is increased, through the parameter N, for certain applications.

Third, because reduction of the competitor set is independent of what has actually been recalled, the Lewandowsky and Murdock (1989) assumptions do not prevent the occurrence of repetitions in the recall output—in which a particular member of the competitor set is chosen repeatedly for recall—despite what the authors suggested in the excerpt. To demonstrate, Figure 2 presents the average number of times each item was recalled per list, broken down by serial position, from the same simulations that produced Figure 1. For the straightforward sampling without replacement scheme described earlier (the closed squares), items are recalled only once per list, by definition. For the Lewandowsky and Murdock (1989) method, repetitions are common—in fact, repetitions occurred on approximately 98% of the simulation trials—and consist primarily of end-of-the-list items. Not only is this last result counter to known data from single-trial serial recall, but it is inconsistent with data from serial-anticipation learning: Intralist intrusions, when they occur, tend to be from middle list items rather than recency items (see Deese & Kresse, 1952; Ebenholtz, 1972).

Finally, because items are competitors in the Lewandowsky and Murdock (1989) scheme only if they follow the current retrieval attempt, it is impossible for the model to produce intralist intrusion distributions that mimic known empirical trends. In single-trial serial recall, in which subjects recall items in order after a single presentation, error distributions typically resemble symmetrical generalization gradients centered on the correct position response. For example, when a subject incorrectly places Item 5 from an eight-item list, that item tends to be placed in an adjacent position (four or six), and response probabilities decrease with increasing distance from the true position response. These error trends are robust and are found regardless of the time scales involved (immediate or delayed recall; see Lee & Estes, 1977; Jahnke, Davis, & Bower, 1989; Nairne, 1990b, 1991, 1992). With the Lewandowsky and Murdock (1989) sampling scheme, because earlier presented items are removed from the competitor set, it is impossible for the subject to intrude an item from early in the list into a later serial position.

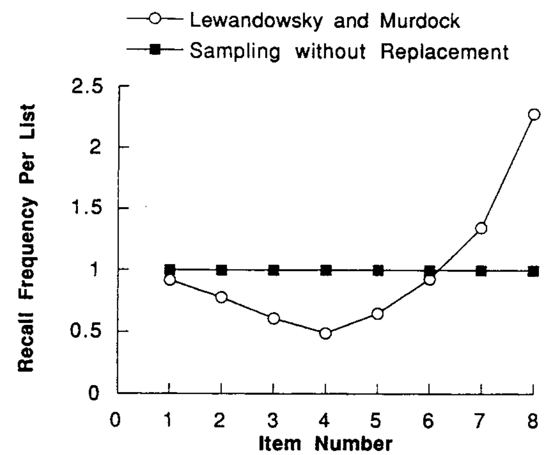


*Figure 2.* Mean number of times each item was recalled on any given list using two different deblurring procedures.

Figure 3 shows the position error gradients for TODAM for each of the items in an eight-item list using the Lewandowsky and Murdock (1989) sampling method. These data are from the same simulation that produced the data represented by open circles in both Figures 1 and 2. Each panel displays the proportion of times a particular item (e.g., the first list item) was recalled in each of the eight possible serial positions. Of most importance is the sudden, abrupt truncation of the gradients. Figure 3 shows quite clearly that the model is incapable of producing the appropriate gradients: Items from the first serial position, for example, can be recalled only in the first serial position and cannot be given as an intrusion response anywhere else in the list. Items from the second serial position can be recalled only in Positions 1 and 2 and so can intrude only in Position 1 and so on. It is only in serial-anticipation learning, in which subjects are given feedback about correct items after each recall, that the Lewandowsky and Murdock method seems reasonable; however, even under these conditions, backward or perseveratory errors still account for 10% to 20% of the total remote intrusions that occur (Bugelski, 1950; Johnson, 1991).

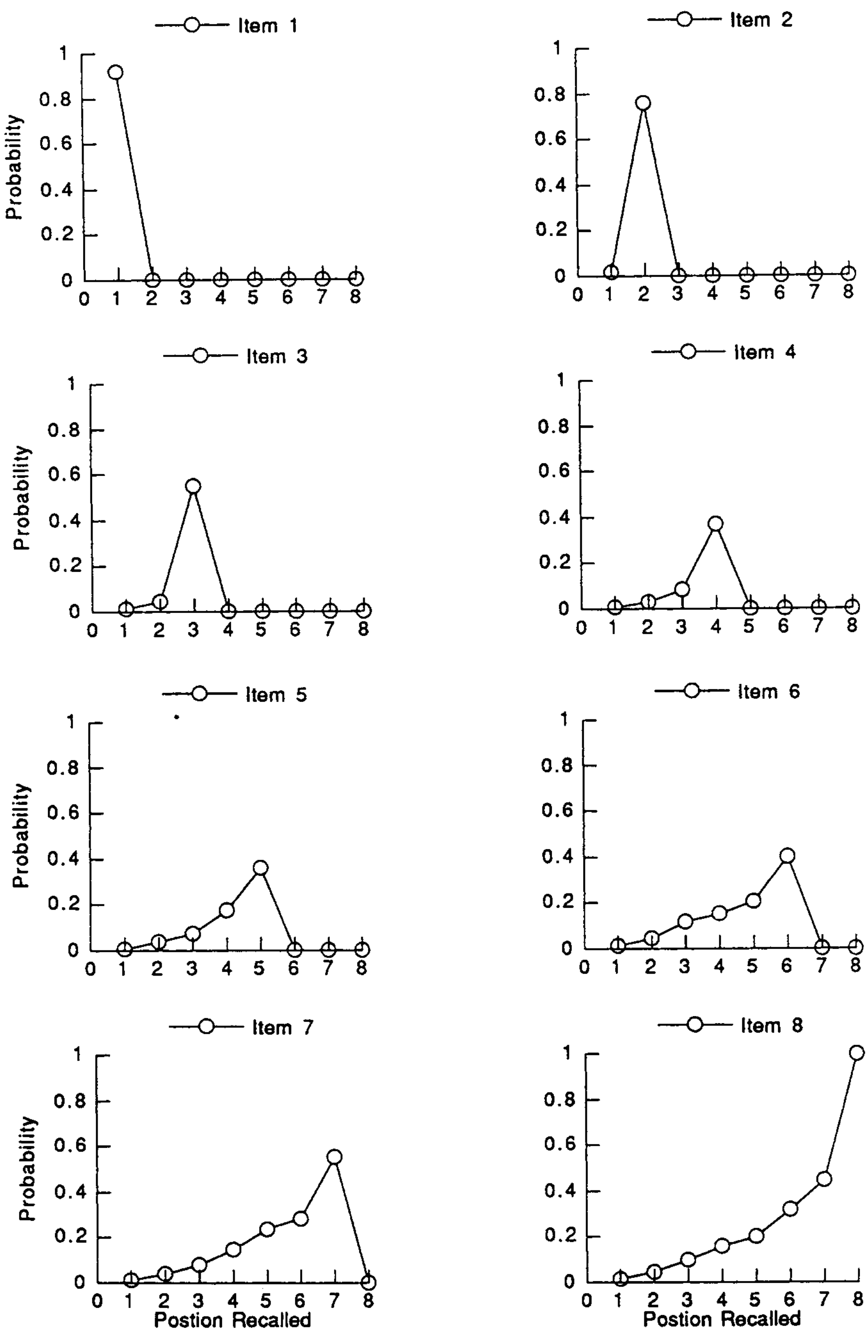


*Figure 3.* Error gradients using only Lewandowsky and Murdock's (1989) deblurring procedure. For example, the graph with the label "Item 3" shows the probability with which the third item was recalled in each serial position.

## Analytic Version of TODAM

Whereas this critique applies primarily to the simulation version of TODAM, the same issues bear critically on the analytic version. In the analytic version, because of the method of numerical integration used (see Equation 3 and Figure 6 of Lewandowsky & Murdock, 1989, p. 31), the exact identity of the competitors is not specified. Rather, response probabilities are based on the average variance of the $m$ competitors. There are additional problems that arise from this conception.

First, recency is again obtained through reduction of the size of the competitor set. There is nothing inherent in the processes of convolution or correlation that produces a bow-shaped serial-position curve: Recall is better at the end of the list because the number of competitors is smaller. Second, by remaining mute about the identity of the competitors, the analytic version cannot produce the appropriate error gradients. The analytic equations reveal only whether the correct item was recalled or not; if the correct item is not recalled, no information is provided about intrusions because the composition of the competitor set is unknown. Therefore, error gradients cannot be modeled and intrusion errors cannot be predicted. If the analytic version is changed somehow to account for the identity of each competitor, then this version should suffer the same problems as the simulation version.

## Summary and Conclusions

The preceding simulations demonstrate that several of the important performance characteristics of the TODAM model, including the shape of the serial-position curve in serial-order recall, depend critically on the deblurring assumptions that define the retrieval process and not simply on the mathematical processes of convolution and correlation. If one adopts a sampling without replacement scheme for determining how the retrieval competitor set unfolds throughout recall, then recency effects are reliably obtained only if "correct" rather than recalled items are systematically removed. Even though Lewandowsky and Murdock (1989) were able to fit a variety of serial-order effects with this scheme, we have shown that it suffers from several serious problems.

First, there is no plausible psychological mechanism that could instantiate the deblurring assumption that subjects might fail to recall an item but still remove it successfully from the competitor set. In fact, even more generally, it is difficult to understand how a composite distributed memory system like TODAM even represents the retrieval competitor set. Is the competitor set also represented in a distributed manner? If so, how can the blurry trace record be compared with each of the items in the set independently? Second, in contrast to the data from single-trial serial recall, the method of deblurring adopted by Lewandowsky and Murdock (1989) leads to intrusions of already recalled items. Third, the intralist intrusions will tend to be items from the end of the list rather than from the middle positions contrary to data from serial-anticipation learning. Fourth, this method of determining the competitor set requires the introduction of an arbitrary parameter, N, to manipulate the size of recency obtained. Fifth, this implementation of TODAM cannot produce intrusion-error gradients that mimic the gradients typically produced for either single-trial serial-recall or serial-anticipation learning. Finally, the analytic version, as currently implemented, cannot produce any error gradients or inform on any other similar intrusion measures.

## References


Bugelski, B. R. (1950). A remote association explanation of the relative difficulty of learning nonsense syllables in a serial list. *Journal of Experimental Psychology, 40,* 336–348.

Crowder, R. G. (1976). *Principles of learning and memory.* Hillsdale, NJ: Erlbaum.

Deese, J., & Kresse, F. H. (1952). An experimental analysis of errors in rote serial learning. *Journal of Experimental Psychology, 43,* 199–202.

Ebbinghaus, H. E. (1885). *Das Gedächtnis: Untersuchungen zur experimentellen Psychologie.* Leipzig: Duncker and Humboldt. Reprinted as Ebbinghaus, H. E. (1964). *Memory: A contribution to experimental psychology.* New York: Dover.

Ebenholtz, S. M. (1972). Serial learning and dimensional organization. In G. Bower (Ed.), *The psychology of learning and motivation* (Vol. 5, pp. 276–314). San Diego, CA: Academic Press.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review, 93,* 411–428.

Jahnke, J. C., Davis, S. T., & Bower, R. E. (1989). Position and order information in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 859–867.

Johnson, G. J. (1991). A distinctiveness model of serial learning. *Psychological Review, 98,* 204–217.

Lee, C. L., & Estes, W. K. (1977). Order and position in primary memory for letter stings. *Journal of Verbal Learning and Verbal Behavior, 16,* 395–418.

Lewandowsky, S., & Murdock, B. B., Jr. (1989). Memory for serial order. *Psychological Review, 96,* 25–57.

Murdock, B. B., Jr. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review, 89,* 609–626.

Murdock, B. B., Jr. (1983). A distributed memory model for serial-order information. *Psychological Review, 90,* 316–338.

Nairne, J. S. (1990a). A feature model of immediate memory. *Memory & Cognition, 18,* 251–269.

Nairne, J. S. (1990b). Similarity and long-term memory for order. *Journal of Memory and Language, 29,* 733–746.

Nairne, J. S. (1991). Positional uncertainty in long-term memory. *Memory & Cognition, 19,* 332–340.

Nairne, J. S. (1992). The loss of positional certainty in long-term memory. *Psychological Science, 3,* 199–202.

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

Slamecka, N. J. (1985). Ebbinghaus: Some associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 414–435.

Young, R. K. (1968). Serial learning. In T. R. Dixon & D. L. Horton (Eds.), *Verbal behavior and general behavior theory* (pp. 122–148). Englewood Cliffs, NJ: Prentice Hall.

## Appendix

Table A1
*Details of the Simulation*

| Parameter | Description | Setting |
|---|---|---|
| $N$ | Dimension of the memory vector | 197 |
| $LL$ | Total list length, including start and end signals | 10 |
| $L$ | No. of items in list *not* including start and end signals | 8 |
| N | Competitor parameter | 0 |
| $\alpha$ | Forgetting parameter | 1.0 |
| $\lambda$ | Rate constant for the exponential decrease of the associative weighting parameter | 0.7 |
| $\omega_0$ | Starting value for the associative weighting parameter | 1.0 |
| $\gamma_0$ | Complement of the associative weighting parameter; determines strength of item information | 0.0 |

Each simulation consisted of 500 trials using an open-loop implementation of TODAM, provided by Stephan Lewandowsky, that was modified to take into account the resemblance between the retrieved vector and the competing items. The parameter settings that resulted in Figures 1 to 3 are listed previously. In all cases, the tolerance level between $a$ and $b$ (see Equation 3 of Lewandowsky & Murdock, 1989) remained constant at 6, as in the original article. The results reported here arose from only two simulations. One simulation produced the curves labeled "Lewandowsky and Murdock" in Figures 1 and 2 and all of the data in Figure 3, and the second simulation produced the curves labeled "Sampling without Replacement" in Figures 1 and 2.

The sole difference between the two simulations was the rule for deciding membership in the competitor set. For the Lewandowsky and Murdock simulation, the "correct" item was removed from the competitor set after each retrieval attempt; when attempting to recall item $i$, regardless of the actual response, item $i$ was removed from the competitor set. For the sampling without replacement simulation, the actual item recalled was removed from the competitor set; when attempting to recall item $i$, the actual response was removed, which may or may not have been item $i$. In both cases, the size of the competitor set decreased as recall progressed, and the recalled item—not necessarily the "correct" item—was added back into the main memory vector, **M**.

The results are not peculiar to the parameters listed here. In general, the results we report are independent of the dimension of the main memory vector ($N$) and the list length ($L$). Whenever a set of parameters, using the Lewandowsky and Murdock sampling procedure, produced a bow-shaped serial-position function, we observed (a) intrusions of already recalled items, (b) intrusions predominantly from end-of-the-list items, as shown in Figure 2, and (c) error gradients similar to those shown in Figure 3. Moreover, the substitution of the straightforward sampling without replacement procedure for the same set of parameters removed the recency effect almost entirely, as shown in Figure 1.

Received April 2, 1992
Revision received September 28, 1992
Accepted December 23, 1992 ■