

UC Office of the President

Reprints

Title

Multiple Reinforcement Effects in Short-Term Memory

Permalink

<https://escholarship.org/uc/item/5zn4r2pt>

Authors

Brelsford, J. W.

Shiffrin, R. M.

Atkinson, R. C.

Publication Date

1968

Peer reviewed

SEP 25 1968

THE UNIVERSITY LIBRARY
UNIVERSITY OF CALIFORNIA, SAN DIEGO
LA JOLLA, CALIFORNIA

Volume 21, Part 1

THE BRITISH JOURNAL OF
MATHEMATICAL
& STATISTICAL
PSYCHOLOGY

EDITED BY R. J. AUDLEY

WITH THE ASSISTANCE OF CYRIL BURT

TAYLOR & FRANCIS LTD
RED LION COURT, FLEET STREET, LONDON, E.C.4

40s. net \$7.45

7251

MULTIPLE REINFORCEMENT EFFECTS IN SHORT-TERM MEMORY

By J. W. BRELSFORD, JR.,¹ R. M. SHIFFRIN and R. C. ATKINSON

Stanford University, Stanford, California

A continuous memorizing situation was studied in which test and study trials alternated throughout an experimental session. The items studied were paired-associates. The interval between study and test for a particular item was randomly determined; and an item was given one, two, three or four reinforcements. A quantitative model is proposed which has two memory stores: a short-term store in which the subject generates a carefully controlled rehearsal scheme of fixed length, and a long-term store in which information is accumulated and lost. A large number of theoretical predictions of the model were verified quantitatively by the data, which confirm results of previous experiments, and support the hypothesis that highly structured rehearsal schemes play a major role in many short-term memory and learning situations.

1. INTRODUCTION

Atkinson & Shiffrin (1968) have proposed a multi-process model of memory that has been employed successfully to describe the effects of a variety of experimental variables. These include: list length, presentation rate and confidence ratings in experiments involving a discrete, independent trial procedure; and mode of rehearsal, size of stimulus set, and judgements of recency in continuous memory experiments. In all of these experiments, items to be recalled were presented for a single period of study with a subsequent test period. In so far as a given item never received more than a single study period (or reinforcement), these experiments may be said to involve memory rather than learning. The present study extends the earlier work into the area of paired-associate learning by giving items varying numbers of reinforcements. An attempt will be made in this paper to explain the results with a model incorporating only a minor extension of the theory used in the previous work.

The task employed in the experiment involves a continuous presentation technique which makes it possible to study the learning process under conditions that are quite uniform and stable throughout the course of an experiment. (This technique is very similar to one employed by Yntema & Mueser (1960, 1962).) In essence the task requires the subject to keep track of the randomly changing responses to eight different stimuli. The eight stimuli are chosen at the start of a session and used throughout that session. A session is begun by presenting for study each of the eight stimuli with associated responses.

¹ Now at Yale University

Following this initial study phase, there is a continuous series of trials, each trial consisting of a test phase followed by a study phase. During the test phase a stimulus is randomly selected from among the set of eight stimuli, and the subject tries to recall the response *last*, or *most recently*, associated with that stimulus. Following the test (and the subject's attempted recall) the study phase of the trial occurs. In the study phase, the same stimulus presented in the preceding test phase is presented, sometimes re-paired with the response that was previously correct, or sometimes paired with a new response; in any case, the subject must study and try to remember the presented pair. In order to distinguish between a particular stimulus and a particular stimulus-response pair, a convention is henceforth adopted that an 'item' will refer to a particular stimulus-response pair. The number of reinforcements for a given item is determined probabilistically but never exceeds four.

The number of trials intervening between the study and test of a given item will be referred to as the 'lag' for that item. Thus, if the test occurs immediately following the study period, the lag is zero. If one trial intervenes (involving test and study on another stimulus), then the lag is one; and so on. Since the stimulus tested is chosen randomly from the set of eight stimuli on each trial, the lag between study and the next test is distributed geometrically with a parameter of 1/8. The task of the subject is simply to remember the current responses assigned to the eight different stimuli. Learning is involved because at the time of testing some of the stimulus-response pairs have had multiple reinforcements, i.e. any given study period may involve the first, second, third or fourth reinforcement of a particular stimulus-response pair. Therefore, the subsequent test of that pair will involve a test after one, two, three or four reinforcements, respectively. The primary dependent variable is the probability of a correct response as a function of lag and the prior number of reinforcements.

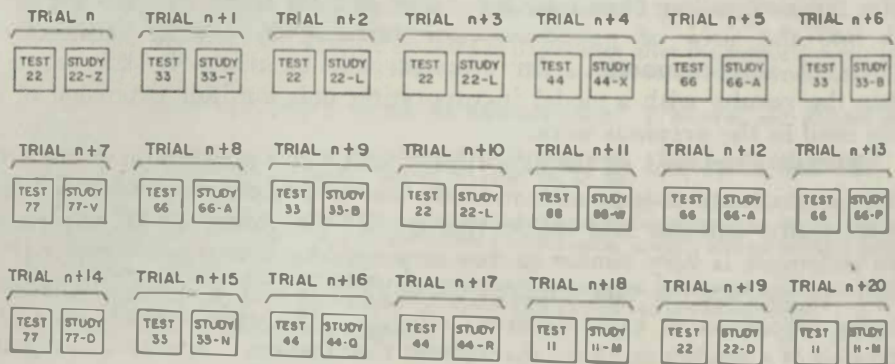


FIGURE 1. A sample sequence of trials.

Before proceeding to the model it may be helpful to illustrate the experimental procedure in some detail. Fig. 1 presents a sample sequence of trials from trial n to trial $n+20$. The stimuli were selected from the set of two digit numbers, and the responses were selected randomly from the alphabet. Various events of the type to be considered later are illustrated in Fig. 1. On trial $n+2$ stimulus 22 is paired with a new response, L , and assigned three reinforcements, the first occurring on trial $n+2$. The second reinforcement occurs on trial $n+3$ after a lag of zero. After a lag of six, the third reinforcement occurs on trial $n+10$. After a lag of eight, stimulus 22 is re-paired with a new response on trial $n+19$. Stimulus 33 is sampled for test on trial $n+6$ and during the study phase is assigned the new response, B , which is to receive two reinforcements, the second on trial $n+9$. Stimulus 44 is tested on trial $n+4$, assigned the new response, X , which is to receive only one reinforcement; thus, when 44 is presented again on trial $n+16$ it is assigned another response which by chance also is to receive only one reinforcement, for on the next trial 44 is studied with response Q . Note also what constitutes a correct response. For example, when stimulus 33 is presented for test on trial $n+6$, the correct response is T ; when presented for test on trials $n+9$ and $n+15$, the correct response is B ; and when presented again, the correct response will be N .

2. MODEL

The theory postulates three memory states: a very short-lived memory system called the sensory register; a temporary memory state called the short-term store (STS); and a more permanent memory storage state called the long-term store (LTS). A consideration of the time intervals and procedure used in the present study leads us to assume that every item is accurately recorded in the sensory register and then at once transferred to STS for active consideration by the subject. In what follows, then, the sensory register plays no part and attention will be restricted to STS and LTS.

Short-term Store

Information entering STS resides there for a short period, perhaps of the order of 15 sec., before being lost completely; this period may be extended indefinitely, however, by processes controlled by the subject, such as rehearsal, coding, and so forth. The memory in the model is imperfect, nevertheless, because only a very limited amount of information may be maintained in STS at any one time. An immediate test of any information that has entered STS (within a few seconds of study, say) will result in correct retrieval of that information. In the present study every item is assumed to enter STS and a test at lag zero occurs within several seconds of study; therefore, the probability of a correct response for any item tested at lag zero is unity. For lags of one or greater the time intervals in the present task are such that items will be lost from STS unless some sort of rehearsal mechanism is invoked to maintain them.

An analysis of subjects' reports and a consideration of similar experiments (Atkinson, Brelsford & Shiffrin, 1967) lead us to propose a very specific, quite orderly, rehearsal scheme called the 'rehearsal buffer' or 'buffer'.

Rehearsal Buffer The buffer is a rehearsal scheme in which a fixed number, r , of homogeneous items are rehearsed at any one time during an experimental session. It is assumed that the series of study items at the start of each experimental session fills the buffer and that the buffer stays filled thereafter. The size of the buffer, r (defined as the number of items held simultaneously), depends upon the nature of the items and thus must be estimated. Once the buffer is filled, each new item that enters causes one of the items currently in the buffer to be lost (i.e. an item currently undergoing rehearsal is eliminated to make room for the new item). It is assumed that a correct response is always given if an item is in the buffer at the time it is tested.

Consideration will now be given to the decision rules by which the subject enters new items into the buffer, and thereby eliminates items currently in the buffer. In order to do this, it is necessary to distinguish theoretically between two kinds of items presented for study: *O*-items and *N*-items. An *O*-item (old-item) is an item presented for study whose stimulus is already a component of an item currently in the buffer. An *N*-item (new-item) is a presented item whose stimulus is not at that time in the buffer.

RULE 1. *Any O-item presented for study is automatically entered into the buffer and replaces there the item with the same stimulus as the presented item.*

There are several reasons for this rule. If the presented item is receiving its second, third or fourth reinforcement, then the item replaced is identical to the one presented and it does not really matter whether we speak of replacement or not, i.e. the state of the buffer remains unchanged in any case. If, on the other hand, the presented item is receiving its first reinforcement, then the item currently in the buffer with the same stimulus will have a different response member. In fact, this response will now be an incorrect response to that stimulus; the subject must therefore change it to the new, correct response, lest he start rehearsing incorrect information.

RULE 2. *When an N-item is presented for study, the subject bases his decisions upon the result of the immediately preceding test phase of the trial: if on the test phase the subject had correctly retrieved from long-term store the item being presented for study, then the item is not entered into the buffer. Otherwise, the item is entered into the buffer with probability α ; if the item enters the buffer, then the item eliminated to make room is chosen randomly (i.e. each item currently in the buffer has probability $1/r$ of being the one eliminated).*

The reasons for this rule are again straightforward. If a correct retrieval of the presented item had just been made from LTS, then the subject has reason to believe he already 'knows' it, and will not attempt to rehearse it further (in this case, of course, the presented item is receiving its second, third or fourth reinforcement). Note that a correct LTS retrieval of the presented

item (and not just a correctly-guessed response or a zero-lag response) is necessary for the subject to prohibit entry into the buffer. If the presented N -item is not retrieved from LTS (which is always the case if the item is being given its first reinforcement) then the subject does not 'know' it, and rehearsal becomes desirable. On the other hand, there undoubtedly is a certain amount of effort involved in rearranging the buffer; in addition, the item which must be removed from the buffer to make room has not yet been tested. In consideration of these conflicting tendencies it seems appropriate to let the entry probability be a parameter, α , to be estimated. Finally, once an item is entered into the buffer, it is necessary to decide which item already there is to be removed. In previous studies using fixed lists, it has proved useful to postulate a tendency for the oldest items in the buffer to be the first eliminated. In all work using the continuous presentation technique of the present study, however, it has been most accurate to postulate a random choice of the item to be removed. Fairly direct evidence supporting this assumption will be presented later in conjunction with Fig. 2.

Long-term Store

LTS is viewed as a memory state in which information accumulates for each item. The term 'information' is not used here in a technical sense. It refers to codes, mnemonics, images or anything else the subject might store that would be retrievable at the time of test. It is assumed that information about an item may enter LTS only during the period that an item resides in the buffer. The status of an item in the buffer is in no way affected by transfer of information to LTS. Whereas recall from the buffer was assumed to be perfect, recall from LTS is not necessarily perfect and usually will not be. At the time of test on an item, a subject gives the correct response if the item is in STS but if the item is not found in STS, the subject searches in LTS. This LTS search is called the *retrieval process*. Two features of the LTS retrieval process must be specified. First, it is assumed that the likelihood of retrieving the correct response for a given item improves as the amount of information stored concerning that item increases. Second, the retrieval of an item gets worse, the longer the information has been stored in LTS. This could result from autonomous decay or active interference from other information being stored in LTS.

It will be specifically assumed in this paper that information is transferred to LTS at a constant rate, θ , during the period that an item resides in the buffer; θ is the transfer rate per trial. Thus, if an item remains in the buffer for exactly j trials, then that item accumulated an amount of information in LTS equal to $j\theta$. Next, it is assumed that each trial following the trial on which an item is knocked out of the buffer causes the information stored in LTS for that item to decrease by a constant proportion τ . Thus, if an item receiving its first reinforcement were knocked out of the buffer at trial j , and i trials intervened between the original study and the test on that item, the amount of information

stored in LTS at the time of test would be $(j\theta)(\tau^{t-j})$. In this experiment an item receiving two, three or four reinforcements may enter and leave the buffer two, three or four times. When the item is in the buffer the θ -process is activated, and when not in the buffer the τ -process takes over.

The probability of a correct retrieval of an item from LTS is now specified. If the amount of information in LTS at the moment of test is zero, then the probability of retrieving the correct response should be zero. As the amount of information increases, the probability of a retrieval should increase toward unity. σ_I is defined as the probability of a correct response for an item which is not in the buffer but has accumulated an amount of information in LTS equal to I at the time of test. Considering the above specifications on the retrieval process,

$$\sigma_I = 1 - (1 - g) \exp(-I), \quad (1)$$

where g is the guessing probability and, in the present experiment, is $1/26$ since there were 26 response alternatives.

As already stated, items that have received two, three or four reinforcements may have been in and out of the buffer at various times and the expression for I becomes fairly complex. Suppose, for example, an item has been in the buffer i_1 trials, then out of the buffer j_1 trials, then in i_2 , out j_2 , in i_3 , out j_3 , and then tested. At that point

$$I = \{[(i_1\theta)\tau^{j_1} + i_2\theta]\tau^{j_2} + i_3\theta\}\tau^{j_3}.$$

As part of the calculation of the probability correct at lag k for a three reinforcement item, one would then have to consider all combinations of i 's and j 's, and compute the conditional probability of each one. It should be clear that this direct method for calculating predictions, although possible, is both formidable and inefficient. Instead, predictions from the model were generated using Monte Carlo methods; the procedure is described in detail later.

At this point it would be appropriate to review the sequence of events and decisions proposed in the model. A trial begins with the presentation of a stimulus for test. The subject first checks STS; if the test was zero lag then the item is present in STS and a correct response is made. Next, or at the same time, the buffer is checked; if the stimulus is in the buffer, then again a correct response is made. If the stimulus is not in the buffer or STS, then LTS is searched, and the probability of retrieving the response correctly is an exponential function of the amount of information then in LTS about the item. If all of these searches fail, then the subject guesses randomly. Next comes the study phase of the trial. The stimulus just tested is now presented again paired with a response for study. In order to determine whether to rehearse this item, the subject refers to the two decision rules given earlier (concerning O - and N -items). It should be noted that the model has four parameters to be estimated: r , the buffer size; α , the buffer entry probability; θ , the transfer parameter; and τ , the long-term decay parameter. The model is identical to that used in previous studies (cf. Atkinson, Brelsford & Shiffrin, 1967) with the sole addition

that an item retrieved from LTS is never entered into the buffer, a state of affairs never arising in experiments where each item receives only one reinforcement.

3. EXPERIMENT

Method

Subjects The subjects were 20 students from Stanford University who received \$2 per experimental session. Each subject participated in at least 10 sessions.

Apparatus The experiment was conducted in the Computer-Based Learning Laboratory at Stanford University. Control functions were performed by computer programs running in a modified PDP-1 computer manufactured by the Digital Equipment Corporation and under the control of a time-sharing system. The subject was seated at a cathode-ray tube display terminal. Stimuli were displayed on the face of the cathode-ray tube (CRT); responses were made on an electric typewriter keyboard located immediately below the lower edge of the CRT. (For a more detailed description of the apparatus see Atkinson, Brelsford & Shiffrin, 1967.)

Stimuli and responses The stimuli were eight two-digit numbers randomly selected for each subject and session from the set of all two-digit numbers between 00 and 99. Once a set of eight stimuli was selected for a given session, it was used throughout the session. Responses were letters of the alphabet, thus fixing the guessing probability of a correct response at $1/26$.

Procedure Every session began with a series of eight study trials; one study trial for each stimulus used in the session. On these trials each of the eight stimuli was paired with a response selected randomly. There were no restrictions on repetition of responses. After the initial study trials the session involved a series of consecutive trials, each consisting of a test phase followed by a study phase. On each trial a stimulus was randomly selected for testing, and the same stimulus was then presented for study. During the study phase of the trial the stimulus was sometimes re-paired with a new response and sometimes left paired with the old response. To be precise, when a particular stimulus-response pair was presented for study the first time, a decision was made as to how many reinforcements (study periods) it would be given; it was given either one, two, three or four reinforcements with probabilities 0.30, 0.20, 0.40 and 0.10, respectively. When a particular stimulus-response pair had received its assigned number of reinforcements, its stimulus was then randomly re-paired with a new response on the next study trial, and this new item was assigned a specific number of reinforcements using the above probability distribution. Reference to Fig. 1 should clarify this procedure. The subject was instructed to respond on the test phase of each trial with the letter that was last studied with the stimulus being tested. Since the stimulus selected for testing was chosen randomly on each trial, the distribution of the lag from study to test was geometric with a parameter of $1/8$.

The temporal arrangement for the eight initial study trials was such that each study trial lasted for 3 sec. with a 3 sec. inter-trial interval. Each trial of the session proper involved the following sequence of events: (1) The word *test* appeared on the upper face of the CRT. Beneath the word *test* a randomly selected member of the stimulus set appeared. This test portion of a trial lasted for 3 sec., during which subjects were told to respond with the last response that had been associated with that stimulus, guessing if necessary. (2) The CRT was blacked out for 2 sec. (3) The word *study* appeared on the upper face of the CRT for 3 sec. Below the word *study* a stimulus-response pair appeared. The stimulus was the same one used in the preceding test portion of the trial. Depending upon the reinforcement schedule, the response was either the one that had previously been correct or a new one. (4) There was a 3 sec. inter-trial interval before the next trial. Thus

a complete trial (test plus study) took 11 sec. A subject was run for 220 such trials during each experimental session. For each subject and session the entire sequence of presentations, the subjects' responses, and response latencies were recorded.

4. RESULTS

In order to eliminate warm-up effects from the data, the first session for each subject and also the first 25 trials of each subsequent session are not included in any of the analyses; otherwise, in what follows, the results are pooled over all sessions and subjects.

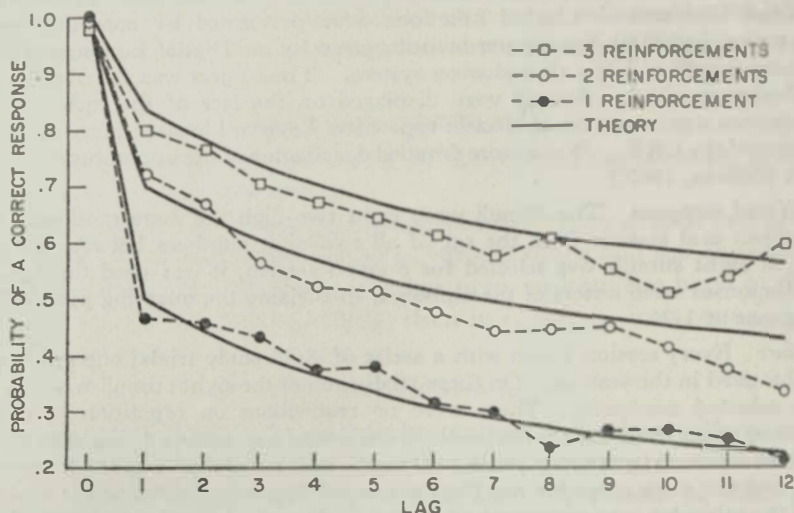



FIGURE 2. Observed and predicted probabilities of a correct response as a function of lag for items tested following their first, second or third reinforcement.

Fig. 2 presents the probability of a correct response as a function of lag for items tested after their first, second and third reinforcements. The number of observations is weighted toward the short lags, and also toward the smaller numbers of reinforcements. The short lags have more observations because the distribution of lags is geometric (with parameter $1/8$). The smaller numbers of reinforcements are weighted because an item receiving k reinforcements provides data for every number of reinforcements less than or equal to k . For example, the one reinforcement lag curve contains not only data from items given just one reinforcement, but also data from the first reinforcement of items given two, three and four reinforcements. No data are presented for the four reinforcement conditions because of the small number of observations. What is graphed is the probability of a correct response to an item that received its j th reinforcement and was then tested after a lag of n trials. The data are presented for values of n ranging from 0 to 15 and for j equal to 1, 2 and 3.

The curves in Fig. 2 exhibit a consistent pattern. The probability correct decreases regularly with lag, starting at a higher value on lag 1 the greater the

number of prior reinforcements. The forms of the 1-reinforcement curve and the probability values are quite similar to those found in a previous study with a similar procedure (Atkinson, Brelsford & Shiffrin, 1967) in which only one reinforcement was given to each item. There is one immediate inference that can be derived from these curves: accepting for the moment the rest of the model, the assumption that items to be lost from the buffer are chosen randomly is reasonably accurate. If there were a sizable tendency for the oldest items in the buffer to be lost first, then the curves would exhibit a pronounced S-shaped effect (Atkinson & Shiffrin, 1968).

For a deeper analysis of this experiment, certain dependencies masked by the curves of Fig. 2 need to be considered. For example, the probability of a correct response to an item that received its second reinforcement and was tested at some later trial will depend on the number of trials that intervened between the first and second reinforcements. To clarify this point consider the following diagram:



Item 22-Z is given its first reinforcement, tested at lag a and given a second reinforcement, and then given a second test at lag b . For a fixed lag b , the probability of a correct response on the second test will depend on lag a . In terms of the model it is easy to see why this is so. The probability correct for an item on the second test will depend upon the amount of information about it in L.T.S. If lag a is extremely short, then there will have been very little time




FIGURE 3. Observed and predicted probabilities of a correct response as a function of both lag a and lag b . (See diagram in text.)

for LTS strength to build up. Conversely, a very long lag a will result in LTS strength building up to a maximum but then decaying once the item has left the buffer. Hence the probability of a correct response on the second test will be maximal at some intermediate value of lag a ; namely, at a lag which will give time for LTS strength to build up, but not so much time that excessive decay will occur. For this reason a plot of probability correct on the second test as a function of the lag between the first and second reinforcement should exhibit an inverted U-shape. Fig. 3 is such a plot. The probability correct on the second test is graphed as a function of lag a . Four curves are shown for different values of lag b . The four curves have not been averaged over all values of lag b because we wish to indicate how the U-shaped effect changes with changes in lag b . Clearly, when lag b is zero, the probability correct is unity and there is no U-shaped effect. Conversely, when lag b is very large, the probability correct will tend toward chance regardless of lag a , and again the U-shaped effect will disappear.

5. THEORETICAL ANALYSIS

Monte Carlo Methods and Parameter Estimates

The evaluation of the goodness-of-fit of the model involves a number of steps. With the present model and experiment, such desirable statistics as maximum-likelihood estimates of the parameters cannot be calculated; indeed, there is no convenient subset of the exact data sequence which will allow us to calculate 'best' estimates of the parameters. Our plan of attack will therefore utilize a representative set of results, namely the data of Figs. 2 and 3, to estimate parameter values; these parameter values will then be used to predict a number of additional results to be presented in the next portion of the paper. This section will be concerned with the method by which the model is fitted to the data of Figs. 2 and 3.

The goodness-of-fit criterion to be used is a type of least-square measure that is calculated in the same manner as chi-square. The criterion will be referred to as minimum chi-square with the understanding that this terminology is based upon the method of calculation and not the theoretical distribution of the measure. (The statistics in Figs. 2 and 3 are by no means independent and our goodness-of-fit measure will not be distributed as chi-square.) Given a goodness-of-fit measure, we now must decide how to generate predictions for a given set of parameter values. As noted earlier, the model is too complex for the usual methods to be employed; instead, Monte Carlo techniques are used. The Monte Carlo procedure involved generating pseudo-data on a computer. These data were generated following precisely the rules specified by the model (and the procedure of the experiment). Therefore, these pseudo-data are an example of how real data would look if the model were precisely correct. Whenever an event occurred in the model that was probabilistic, a random number generator was used to determine the occurrence of that event. The same

analysis routine that was applied to the real data to generate Figs. 2 and 3 was applied to the pseudo-data to generate theoretical predictions. A large sample of pseudo-data was generated so that the theoretical predictions were stable and accurate. Having generated theoretical predictions for a given set of parameter values, the goodness of fit was then measured by the χ^2 criterion:

$$\chi^2(n, r, \alpha, \theta, \tau) = \sum_i \left\{ \frac{1}{N_i \Pr(C_i)} + \frac{1}{N_i - N_i \Pr(C_i)} \right\} \left\{ N_i \Pr(C_i) - O_i \right\}^2, \quad (2)$$

where the sum is taken over all data points in Figs. 2 and 3. For the actual data, the observed number of correct responses for the i th point is denoted by O_i , and N_i is the total number of responses.

For the Monte Carlo data, $\Pr(C_i)$ denotes the probability of a correct response; of course, the value of $\Pr(C_i)$ depends upon the parameter vector $(r, \alpha, \theta, \tau)$. We denote by n the number of simulated trials upon which $\Pr(C_i)$ is based. As usual, the smaller the χ^2 value, the closer are the predicted and observed data.

In order to determine which set of parameter values provides the best fit according to the χ^2 criterion, a fairly exhaustive search of the parameter space was undertaken. (The grid search procedure is similar to that described in Atkinson & Crothers (1964).) Initially the search was carried out using 4,000 different values for the parameter vector $(r, \alpha, \theta, \tau)$, with a χ^2 value being generated for each vector. For that parameter vector yielding the minimum χ^2 , a second, converging, search routine was then begun. These converging searches were continued until the set of parameters generating the minimum χ^2 was determined to two decimal places. One problem that arises when using a Monte Carlo procedure is determining how large a sample of pseudo-data is required to ensure accurate predictions. In this case, an attempt was made to ensure accurate predictions by generating a very large amount of pseudo-data for each set of parameter values: in the initial grid search each set of parameter values were given 200 subjects worth of pseudo-data (each subject consisting of 2200 pseudo-trials). As a check on the stability of the predictions, the pseudo-data were divided in half (100 subjects each) and separate χ^2 's computed for each half [$\chi^2(1)$ and $\chi^2(2)$]. The following inequality was then evaluated:

$$0.95 \leq \frac{\chi^2(1)}{\chi^2(2)} \leq 1.05. \quad (3)$$

If this inequality was satisfied then χ^2 was computed for the entire 200 subjects worth of pseudo-data, and the value was assumed to be appropriate for that point in the parameter space. If the inequality failed, then the number of pseudo-trials was doubled and the split-half test again applied. The number of pseudo-trials was increased in this manner until the inequality was satisfied. Furthermore, as the grid search began to converge on the best set of parameter values, the bounds of inequality (3) were narrowed so that even greater accuracy of predicted values would result. When the final set of parameter values was

obtained, 12,500 Monte Carlo subjects were generated using those parameter values. In all subsequent discussions, the predicted values are based on the output of this final Monte Carlo run, and it is doubtful that they reflect any fluctuations due to sampling error. The best set of parameter values were: $r=3$; $\alpha=0.65$; $\theta=1.25$; $\tau=0.82$.

These parameter values are in reasonable accord with those found in previous experiments using the same model (Atkinson & Shiffrin, 1968). The buffer size of $r=3$ may seem small at first glance but r has been estimated to be 2 or 3 in each of the experiments of the present type. It is not hard to see that a small value of r is to be expected: an r of 3 indicates that the subject is simultaneously rehearsing six numbers and three letters, a fairly difficult task considering interruptions for tests and changes in rehearsal following study periods. The value of τ is fairly low, considering that this decay factor is applied on each trial; one explanation for L'S strength being reduced so quickly would hold that there is a great deal of retroactive interference in this situation. If so, then there should be evidence of this in the data. Evidence along these lines will be presented shortly.

The predictions from the theory are shown as the smooth curves in Figs. 2 and 3. It should be evident that the predicted values are quite close to the observed ones. Note also that the seven curves in the two figures are fitted simultaneously with the same four parameter values. The fact that the spacing of the curves is accurately predicted is particularly interesting.

Some Further Predictions

A number of statistics that were not used in estimating parameters are now examined. These statistics test specific predictions of the model, predictions that were in some cases contrary to the authors' *a priori* intuitions. Most of the machinery of the model, and its most novel features, lie in the rehearsal scheme called the buffer. Fortunately, the very specific assumptions made concerning the working of the buffer lead to clear-cut predictions to be searched for in the data. Consider, for example, the hypotheses concerning *O*- and *M*-items. The model predicts that the kind of items intervening between study and test will influence the probability correct at test. For one thing, the more *N*-items intervening between study and test, the less the probability correct, because only *N*-items can eliminate an item from the buffer. Similarly, the more intervening *O*-items, the greater the probability correct, because *O*-items cannot knock the studied item from the buffer. Although *O*-items and *N*-items are not directly identifiable in the data, the probability of their occurrence can be maximized by selecting appropriate event sequences. Thus in Fig. 4 the 'all-same' and 'all-different' curves are plotted. For the all-same curve, the probability of a correct response is computed as a function of the lag, when all the intervening items between study and test utilize the same stimulus. There are three such curves, depending upon whether the studied item had been given




FIGURE 4. Observed and predicted lag curves for the all-same and all-different conditions.

its first, second or third reinforcement. The model predicts that once any of the intervening items enters the buffer, every succeeding item will be an *O*-item and hence will also enter the buffer (because all these intervening items have the same stimulus). Thus the all-same curve should decrease less (as a function of lag) than the unconditional lag curves presented in Fig. 2. For the all-different curves, on the other hand, the probability of a correct response is computed for instances when every intervening item between study and test utilizes a different stimulus. In this case, the number of intervening *N*-items tends to be maximized, and hence the probability that the studied item will be knocked out of the buffer tends to be maximized. Therefore, the all-different curves should decrease faster than the unconditional lag curves. It will be seen in Fig. 4 that the all-same and all-different curves conform to these predictions. The solid lines in the figure represent predictions from the model using the parameter values estimated in the previous section; the correspondence of data and theory appears to be reasonably close.

Next consider the factors determining the probability that a presented item will enter the buffer. Most important is the probability that the presented item is an *O*-item, since an *O*-item's stimulus is already in the buffer, and thus every *O*-item enters the buffer. Because a high probability of entering the buffer implies a high probability of a correct response at test, it should be possible to manipulate the probability correct by manipulating the probability that a presented item is an *O*-item. In Fig. 5 this has been done. Consider a sequence of consecutive trials all utilizing the same stimulus, but with the last

item in the sequence being given its first reinforcement (thus its response will be different from that of the immediately preceding item). Once any item in this sequence enters the buffer, every following item will do so also; thus the longer the sequence, the greater the probability that the last item in the sequence will enter the buffer. Fig. 5 plots the probability of a correct response as a function





FIGURE 5. Observed and predicted probabilities of a correct response as a function of the number of consecutive preceding items using the same stimulus.

of the length of the sequence of preceding items all utilizing the same stimulus. The curve plotted is *not* a lag curve; the probability correct is pooled for over all lags at which the eventual test occurred. The theoretical predictions are generated from the previously estimated parameter values, and again there is a good correspondence between theory and data. The effect found here is particularly important because it emphasizes the dichotomy between short- and long-term processes. A traditional interference theory would seem to predict just the opposite effect from that found, in that 'proactive interference' should increase as the length of the preceding sequence increases. Nevertheless, indications of typical interference effects will be seen for long lags, where LTS, and not the buffer, is playing a predominant role.

In order to extend and verify the results of Fig. 5, consideration is now given to the effect of the lag preceding an item's presentation for study. To make matters clear, consider the following diagram:



Item 22-Z is studied for the j th time and is then tested at lag a ; on this trial 22 is paired with a new response X , and tested next at lag b . According to the

theory, the shorter lag a , the better performance should be when the item is tested after lag b . This prediction is based on the fact that the more recently a stimulus had appeared, the more likely that it was still in the buffer when the next item using it was presented for study. If the stimulus was in the buffer, then the item using it is an O -item and automatically enters the buffer. In the present analysis, we examine this effect for three conditions, i.e. the preceding item using the stimulus in question could have just received its first, second or third reinforcement. Fig. 6 presents the appropriate data. In terms of the




FIGURE 6. Observed and predicted probabilities of a correct response as a function of lag a . (See diagram in text.)

above diagram, what is plotted is the value of lag a on the abscissa versus the probability of a correct response averaged over all values of lag b on the ordinate. There is a separate curve for $j = 1, 2$ and 3. Note that the results of Fig. 5 are confirmed by those of Fig. 6; again interference theory would appear to predict an effect opposite to that found.

The predicted curves are based upon the previous parameter estimates. The predictions and observations coincide fairly well, but the effect is not as dramatic as one might hope. One problem is that the predicted decrease is not very large. Considerably stronger effects may be expected if each curve is separated into two components: one where the preceding item was correct at test and the other where the preceding item was not correct. In theory the decrease predicted in Fig. 6 is due to a lessened probability of the relevant stimulus being in the buffer as lag a increases. Since an item in the buffer is

always responded to correctly, an analysis made conditional upon correct responses or errors (the centre test in the above diagram) should magnify the effect. To be precise, the decrease will be accentuated for the curve conditional upon correct responses, whereas no decrease at all is predicted for the curve conditional upon correct responses, whereas no decrease at all is predicted for the curve conditional upon errors. If an error is made, the relevant stimulus cannot be in the buffer and hence the new item enters the buffer with probability α , which is independent of lag a . Fig. 7 presents the conditional curves and the

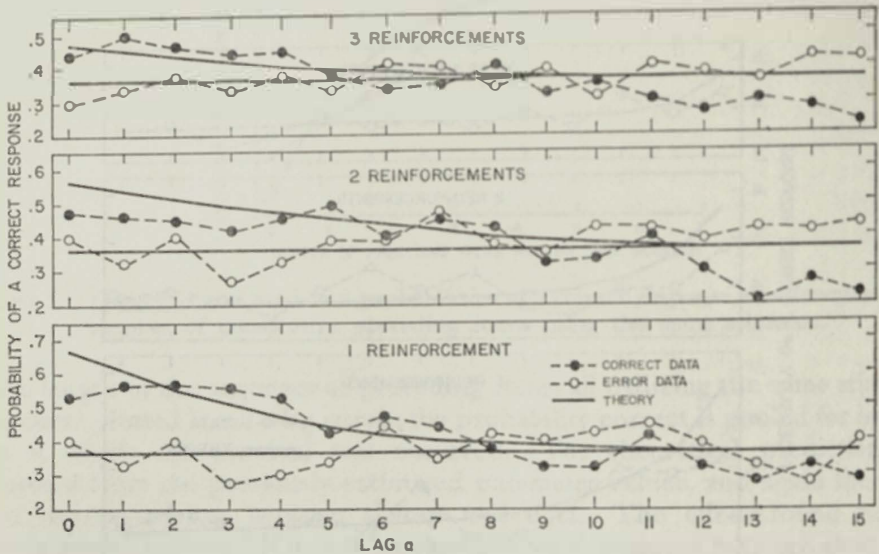



FIGURE 7. Observed and predicted probabilities of a correct response as a function of lag a conditional upon correct and incorrect responses. (See diagram in text.)

predictions. The decreasing effect is fairly evident for the correct curves, whereas the error curves, as predicted, are quite flat over lags. Conceivably one might argue that the effects are due to item selection with correct responses indicating easier stimuli and incorrect responses indicating more difficult ones. However, it is difficult to imagine how item selection could explain the eventual crossing of the correct and error curves found in each of the three diagrams. (Undoubtedly there are some selection effects in the data of Fig. 7, but their magnitude is difficult to determine. Thus, these data should be regarded with some caution.) Indeed, the model does not explain the crossover. The model predicts that the two curves should meet. The model is in error at this point because it has not been extended to include negative transfer effects. Such an extension would not be difficult to implement, however. An item responded to correctly at a long lag probably has a strong LTS trace. This strong trace would then interfere with the LTS trace of the new item which, of course, uses the same stimulus. Thus the data of Fig. 7 appear to imply opposite effects at

short and long lags, with the effect at long lags (when LTS rather than the buffer may be assumed to be the predominating influence) exhibiting the result expected from interference theories.

Figs. 6 and 7 both show that the predicted decrease becomes smaller as the number of reinforcements increases. The fact that the data seem to support this prediction helps confirm certain of the buffer-replacement assumptions. The decreasing effect as reinforcements increase is predicted because the probability of entering the buffer is reduced for an item receiving its third reinforcement; this occurs as a consequence of the assumption that an item retrieved from LTS on the test phase of the trial is not entered into the buffer on the study phase of that trial. Thus as reinforcements increase, the probability of being in the buffer as a result of short lag a is partially counterbalanced.

By and large, the data and predictions to this point may be considered to provide fairly strong support for the details of the model. The feature that has been left out of the model is that of LTS response competition, or negative transfer. The model fails to take account of this effect because it ignores residual information in LTS from previous items using the same stimulus. This lack is most clearly indicated by the occurrence of intrusion errors; particularly, errors that were correct responses on the preceding occurrence of that stimulus. For example, consider the following sequence:



Item 22-Z is studied for the j th time and then tested at lag a . On this trial 22 is paired with a new response X and next tested at lag b . By an intrusion error we mean the occurrence of response Z when 22 is tested at the far right of the diagram. The model predicts that these intrusion errors will be at a chance level ($1/25$) and independent of lag and number of reinforcements. In fact, these predictions fail. Fig. 8 presents the probability of intrusion errors as a function of lag b , where the data are pooled for all values of lag a . Three curves are plotted for $j=1, 2$ and 3. If it is assumed that the probability of giving the previously correct response as an intrusion error is some function of that previous item's current LTS strength, then the pattern of results in Fig. 8 follows naturally. For example, the more reinforcements of the previous item, the greater its LTS strength, and the greater the probability of its response being given as an intrusion error. Thus this failure of the model is not very distressing, rather it was expected. The model could be extended in a number of obvious ways to take account of competing LTS traces without appreciably changing the accuracy of the predictions already considered.




FIGURE 8. Probability of an intrusion error as a function of the prior number of reinforcements.

The major emphasis in this paper has been on rehearsal processes conceptualized in the framework of a buffer mechanism. Because of this emphasis upon short-term processes, the experiment should not be considered as a strong bridge to the usual paired-associate learning situation. However, a number of long-term effects (such as intrusion errors and interference caused by previously learned items on new items using the same stimulus) demonstrate that LTS mechanisms play an important role in this study. It is undoubtedly true that modifications of the theory are required before it can be applied to the typical paired-associate learning experiment. For example, it would be necessary to provide more structure for the workings of LTS. Several possible forms for such structure have been suggested by Atkinson & Shiffrin (1968). These new considerations have yet to be experimentally investigated, but it appears that the theoretical gap between 'memory' and 'paired-associate learning' tasks may be bridged without too much difficulty in the near future.

ACKNOWLEDGEMENTS

This research was supported by the National Aeronautics and Space Administration, Grant No. NGR-05-020-244. The authors wish to thank Susan Shiffrin for collecting the data reported in this paper.

REFERENCES

- ATKINSON, R. C., BRELSFORD, J. W. & SHIFFRIN, R. M. (1967). Multi-process models for memory with applications to a continuous presentation task. *J. math. Psychol.* 4, 277-311.
- ATKINSON, R. C. & CROTHERS, E. J. (1964). A comparison of paired-associate learning models having different acquisition and retention axioms. *J. math. Psychol.* 1, 285-315.
- 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: Advances in Research and Theory*, vol. 2. New York: Academic Press.

Multiple Reinforcement Effects in Short-Term Memory

19

- YNTEMA, D. B. & MUESER, G. A. (1960). Remembering the present state of a number of variables. *J. exp. Psychol.* 60, 18-22.
- YNTEMA, D. B. & MUESER, G. A. (1962). Keeping track of variables that have few or many states. *J. exp. Psychol.* 63, 391-395.