1

A learning model for discrete performance levels

GORDON H. BOWER, *Stanford University*
JOHN THEIOS, *University of Texas*

Recent developments in mathematical models for learning have been concentrated on Markovian interpretations of changes in performance levels (Estes, 1959; Suppes and Atkinson, 1960). During the course of learning, the subject's behavior is conceived to change in discrete steps or stages, and each stage is associated with a different level of response probability. Moreover, the transition of the subject from one stage to the next is assumed to follow certain Markov properties, e.g., the probability that the subject leaves his current state is constant and does not depend upon the prior history.

The simplest example of such a Markov model is the two-state process identified as the one-element pattern model (Bower, 1961; Estes, 1959). In that model, the learning of individual subjects is characterized by a single-step function. The subject begins with an initial probability p of a correct response, and on each trial there is some probability c that an effective learning event occurs, shifting the probability of a correct response to unity. This is a one-step or one-stage model, and it has proved quite useful in integrating data collected in several elementary human learning situations (Bower, 1961, 1962a, 1962b; Bower and Trabasso [this volume, p. 32]; Suppes and Ginsberg, 1963).

It is clear, however, that the simple one-step model is not going to work for a fairly large number of learning situations. It is an easy matter to arrange experimental situations in which the data obviously contradict the expectations

Support for this research was provided by research grant M-3849 to the first author from the National Institutes of Mental Health. Collaboration by the second author was made possible in part through his participation in the Social Science Research Council's Summer Institute on Mathematical Models in Social Science Research, Psychology of Choice and Decision, Stanford University, 1962, and in part through a research leave granted by the University of Texas Research Institute for part of the 1962 summer. Computer time for data analysis was paid for by grant M-6154 from the U.S. Public Health Service.

of a one-step model. To keep matters general for the moment, imagine a learning situation in which initially the probability of the correct or conditioned response is zero, and in which all subjects reach a performance asymptote of unity after sufficient training. Application of the simple one-stage model to this situation leads to the prediction that when the subject makes his first correct response, he has learned; so no errors should follow that first correct response. Although a few experimental situations can be constructed in which this deterministic prediction is approximately correct (Bower, 1962b; Maatsch, 1959), in the majority of ordinary experimental situations the prediction is clearly false—that is, getting out the first correct response is usually not the terminal point of the learning process. Typically, following the first correct response a number of reinforced trials are required before the subject makes his last error and goes into a criterion run of successes.

Faced with such data, we have made the conjecture that perhaps the nonasymptotic behavior following the first success represents a second, intermediate stage in the process of learning. This conjecture derives from earlier theoretical developments by Estes (1959) and by Suppes and Atkinson (1960) with small-element stimulus-sampling models. We asked: Given that the starting and terminal response probabilities are zero and one, respectively, can the data be adequately represented by assuming that the subject goes through a single, intermediate state in which his response probability is some fixed value p between zero and one?

In the cases analyzed to date employing this idea, it has been found consistently that the data can be adequately represented by assuming that a single intermediate performance level occurs between the starting and terminal performance levels. We think that this is an important fact. The fact itself is easily demonstrated and depends in no way upon estimating parameters of learning models or upon complicated statistical manipulations of the data. By simple and direct methods of data analysis, the existence of an intermediate performance state stands out clearly.

The question is how to exploit this fact for the purpose of constructing theoretical models for learning in these situations. Ideally, one would like to have a rational theory which, in these local applications, implies a three-state process. The theory should provide some meaningful interpretation of the intermediate state and the specific level of p observed, and relate the learning parameters (for transitions among the three states) to experimental variables. The specific interpretation of the theory will depend upon whether it is applied to escape and avoidance conditioning in rats, to reversal of paired-associate learning, or to eyelid conditioning. Suitable interpretations of the stimulus-sampling theory (Estes, 1959) are sufficient to provide the necessary theoretical models. Specifically, we require the assumption that the stimulating situation for the subject may be represented by two component stimulus events or patterns. Each stimulus component may or may not be conditioned to the correct response. To admit the possibility that none, one, or both of the patterns may be conditioned leads to the expectation that an individual's response probability will be zero, some intermediate value, or unity over successive trials of the experiment. This general scheme has some flexibility in assumptions at the following points: (*a*) the way in which the components are sampled by the subject on each trial, (*b*) the rule for determining response probability in the presence of a particular sample, and (*c*) the possible changes in conditioning states of the sampled components. Examples of such uassmptions are given in our discussions of data from an avoidance-learning experiment with rats and of data from an eyelid-conditioning study with human subjects. In the following formal analysis, the three-state model is developed simply as a Markov chain and the stimulus-sampling interpretation is suppressed—that is, it may be viewed as a useful empirical model that provides an economic description and summarization of the data. The usefulness of the model does not necessitate a commitment to the stimulus-sampling theory that led us to develop the model.

1. The three-state model

The three states of the Markov model correspond to the three values of response probability, 0, p, and 1. Each subject begins training in state 0, and after a sufficiently large number of training trials he ends up in state 1, where he gives errorless performance. All subjects are assumed to go through the intermediate state p before reaching state 1. The directed graph in Fig. 1 illustrates the Markov chain. The parameter c represents the probability of a transition from state 0 to state p, and θ represents the probability of a transition from state p to state 1.

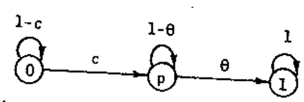

FIG. 1. Markov chain of simple three-state model. The constant c represents the probability of a one-trial transition from state 0 to state p. The constant θ represents the probability of a transition from state p to state 1.

Several comments should be made about this representation. First, it excludes the possibility of a direct transition from the starting state 0 to the terminal state 1. However, this does not mean that all subjects are expected to make at least one error following their first correct response. In fact, a considerable proportion of subjects are expected to go into a criterion run of correct responses starting with their first correct response. If the observed percentage is larger than expected, then one would consider modifying the model to permit one-trial transitions directly from state 0 to state 1. Such a modification would lead to a mixed model consisting of a weighted combination of the one-stage and two-stage models. A second point to notice about the Markov chain diagramed in Fig. 1 is that the number of trials to leave states 0 and p is geometrically distributed as $c(1-c)^{n-1}$ and $\theta(1-\theta)^{n-1}$, respectively. The total number of trials required to reach the terminal absorbing state 1 is the simple convolution of these two geometric distributions, since the waiting times in state 0 and state p are uncorrelated. The distribution of the number of errors in state 0 is the same as that of the number of trials in that state; in state p, the distribution of total errors is the same as for the simple one-element model with parameters p and θ (Bower, 1961).

For carrying out mathematical derivations with this model, a convenient strategy is to separate the correct and incorrect responses made in the intermediate conditioning state, and to imagine that there are two intermediate states corresponding to trials on which a success or an error occurs—that is, to define an intermediate success state S that the subject is in on a trial when he makes a correct response, and an intermediate error state E that the subject is in when he makes an error. The expanded Markov chain is given in Fig. 2, with $q = 1 - p$. We have inserted arbitrary transition probabilities s and ε for learning following a success and an error, respectively, in the intermediate state. If $s = \varepsilon = \theta$, then the model in Fig. 2 is identical to the one in Fig. 1. The model with arbitrary parameters s and ε permits separate assessment of the effectiveness of reinforcing events on S and E trials. In some applications of the model to data, it has been found that one or another of these parameters is zero, so that an effective reinforcing event occurs on S trials or on E trials but not on both. We feel that information of this sort is useful; therefore, s and ε are kept as free parameters, and the data will dictate their values.

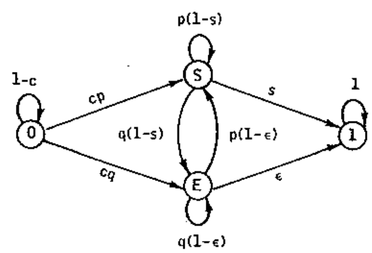

FIG. 2. Expanded Markov chain for the general model.

2. Mathematical analysis of the model

One of the tactical advantages of most Markov learning models is that, compared with corresponding linear models, they permit an extensive number of predictions to be derived in closed form and tested against a single set of data. The present Markov model is no exception to this rule. From the present model one can easily derive the distributions and moments of a variety of random variables representing properties of the subject's response protocol. The derivations assume that the subject is given an infinite number of trials. The condition of infinite response sequences is approximated in practice by training each subject to a stringent criterion of learning.

For most of the derivations, an essential item of information is the probability $w_{i,n}$ that the subject is in state i ($i = 0$, S, E, or 1) at the beginning of

These derivations assume state 0 produces only errors or q=0 trial n, assuming that he begins in state 0 on trial 1. To obtain $w_{0,n}$, note that there is a constant probability c of leaving state 0 on each trial. Hence, $w_{0,n}$ will be $(1 - c)^{n-1}$, representing the probability of $n - 1$ failures to leave state 0. For the probabilities $w_{E,n}$ and $w_{S,n}$, the following recursive expressions can be written:

$$\begin{aligned} w_{E,n+1} &= w_{E,n}(1 - \varepsilon)q + w_{S,n}(1 - s)q + w_{0,n}\, cq, \\ w_{S,n+1} &= w_{E,n}(1 - \varepsilon)p + w_{S,n}(1 - s)p + w_{0,n}\, cp. \end{aligned} \tag{1}$$

Inspection of these two expressions shows them to be the same except that q and p are interchanged. Hence, for all n, it is true that $pw_{E,n} = qw_{S,n}$. This relation may be used to substitute for, say, $w_{S,n}$ in the top equation. If we make this substitution and let $w_{0,n} = (1 - c)^{n-1}$, the top equation then reads

$$w_{E,n+1} = w_{E,n}(1 - \varepsilon)q + w_{E,n}(1 - s)p + cq(1 - c)^{n-1}. \tag{2}$$

Solution of Eq. (2) yields the following result:

$$w_{E,n} = \frac{qc}{\theta - c}\left[(1 - c)^{n-1} - (1 - \theta)^{n-1}\right], \tag{3}$$

where $\theta = ps + q\varepsilon$ is the unconditional probability of leaving the combined set of intermediate states. Also, by the previous relation $w_{S,n} = p/q\, w_{E,n}$, we can easily obtain $w_{S,n}$ from Eq. (3). Because errors occur with probability 1.00 when the subject is in state 0 or state E, the unconditional probability of an error on trial n, q_n, will be the sum of $w_{0,n}$ and $w_{E,n}$.

We now derive the probability that the subject makes k errors before his first success. Let J_0 be a random variable representing the number of errors before the first success. There are two paths through the Markov process that will give $J_0 = k$. The first path is for the subject to make all k errors in state 0 and then, after the kth error, move into state S for his first success. The second path is for the subject to make some intermediate number j $(1 \leq j \leq k - 1)$ errors in state 0, and then move into state E and make $k - j$ errors. Following the $(k - j)$th error in state E, a success occurs on the next trial either in state S or in state 1. Suppose we let $\alpha = (1 - \varepsilon)q$ represent the probability of an error following an error in the intermediate state. The distribution of J_0 can be written as follows:

$$\Pr\{J_0 = k\} = (1 - c)^{k-1}cp + \sum_{j=1}^{k-1} (1 - c)^{j-1}cq(\alpha)^{k-j-1}(1 - \alpha). \tag{4}$$

The second term is the sum over all j of the probabilities of making j errors in state 0 and $k - j$ errors in state E before the first success. Performing the summation in Eq. (4), we obtain the final result:

$$\Pr\{J_0 = k\} = cp(1 - c)^{k-1} + \frac{cq(1 - \alpha)}{\alpha - (1 - c)}\left[\alpha^{k-1} - (1 - c)^{k-1}\right]. \tag{5}$$

The mean value of J_0 is

$$E(J_0) = \frac{1}{c} + \frac{q}{p + q\varepsilon}. \tag{6}$$

Equation (6) is intuitively sensible. The mean number of errors before the first success is equal to the mean trials to leave state 0 plus the mean number of errors in the intermediate state before a success occurs.

We have defined J_0 as the number of errors before the first success. In general, we define J_k as the number of errors intervening between the kth and $(k + 1)$th successes. We wish to find the probability that J_k takes on some arbitrary value $i = 0, 1, 2, \cdots$. A useful method for finding the distribution of J_k is first to find the probability (call it g_k) that the kth success occurs in state S rather than in terminal state 1. If the kth success occurs in state 1, then J_k will equal 0 since no further errors occur once state 1 is entered. We begin by finding g_1, the probability that the first success occurs in state S. It is

$$\begin{aligned} g_1 &= p + q(1 - \varepsilon)p + p[q(1 - \varepsilon)]^2 + \cdots \\ &= p \sum_{i=0}^{\infty} \alpha^i = \frac{p}{1 - \alpha}. \end{aligned} \tag{7}$$

A recursion may be written relating g_k to g_{k-1}. The recursion is

$$g_k = g_{k-1}(1 - s) \cdot p \sum_{i=0}^{\infty} \alpha^i = g_{k-1}(1 - s)g_1. \tag{8}$$

The logic of this recursion is straightforward. In order for the kth success to occur in state S, it must be the case that the $(k - 1)$th success occurred in state S, which has probability g_{k-1}; that conditioning failed to occur on the trial of the $(k - 1)$th success, which has probability $1 - s$; and that the next success occurred in state S, which has probability g_1. Equation (8) is a linear difference equation and can be solved to yield

$$g_k = \frac{1}{1 - s}[g_1(1 - s)]^k \qquad (k \geq 1). \tag{9}$$

As noted before, the value of g_k enters importantly into the distribution of J_k, the number of errors between the kth and $(k + 1)$th successes. The distribution of J_k may be written as follows:

$$\Pr\{J_k = i\} = \begin{cases} 1 - q(1 - s)g_k & \text{for } i = 0, \\ g_k(1 - s)q(1 - \alpha)\alpha^{i-1} & \text{for } i \geq 1. \end{cases} \tag{10}$$

The mean value of J_k is

$$E(J_k) = \frac{q}{p} g_{k+1}. \tag{11}$$

The value of g_k decreases exponentially with k; hence, the expected number of errors between the kth and $(k + 1)$th successes also decreases exponentially with k. By taking appropriate sums of the $E(J_k)$'s, one can predict the mean number of errors before the mth success.

We now derive the distributions of the trial of the last error and the total number of errors. For each problem, useful summary probabilities are f and b, which are defined as the probabilities that no more errors will be made following a particular trial in state S or state E, respectively. The probability of no more errors starting in state S is

$$f = s + (1 - s)ps + [(1 - s)p]^2 s + \cdots \tag{12}$$

$$= s \sum_{i=0}^{\infty} [p(1 - s)]^i = \frac{s}{1 - p(1 - s)}.$$

The probability of no more errors following a response made in state E is

$$b = \varepsilon + (1 - \varepsilon)pf. \tag{13}$$

Using these summary probabilities f and b, we can write the probability distribution of the trial of the last error (call it n') as

$$\Pr\{n' = k\} = w_{0,k}pcf + w_{E,k}b. \tag{14}$$

The first term is the probability that the final error occurs in state 0; the second term is the probability that it occurs in state E. The mean trial of the last error will be

$$E(n') = \frac{1}{c} + \frac{q}{\theta(q + ps)}. \tag{15}$$

To obtain the distribution of total errors, T, we use the relation $T = t_0 + t_1$, where t_0 and t_1 are the numbers of errors made in states 0 and E, respectively. The distribution of T is obtained by convoluting the distributions of t_0 and t_1, which are

$$\begin{aligned} \Pr\{t_0 = k\} &= (1 - c)^{k-1}c && \text{for } k \geq 1, \\ \Pr\{t_1 = j\} &= \begin{cases} pf & \text{for } j = 0, \\ (1 - pf)b(1 - b)^{j-1} & \text{for } j \geq 1. \end{cases} \end{aligned} \tag{16}$$

The convolution is carried out in two parts to handle the two pieces of the t_1 distribution—that is,

$$\Pr\{T = k\} = \Pr\{t_0 = k\}\Pr\{t_1 = 0\} + \sum_{i=1}^{k-1} \Pr\{t_0 = i\}\Pr\{t_1 = k - i\}. \tag{17}$$

Substitution of the t_0 and t_1 distribution functions into Eq. (17) yields the following summated result:

$$\Pr\{T = k\} = cpf(1 - c)^{k-1} + \frac{cb(1 - pf)}{c - b}[(1 - b)^{k-1} - (1 - c)^{k-1}]. \tag{18}$$

The expectation of T is given by

$$E(T) = \frac{1}{c} + \frac{q}{\theta}. \tag{19}$$

Sequential statistics. Let x_n be the response indicator variable taking the values 1 or 0 accordingly as an error or success occurs on trial n. We wish to find the expected number of times j-tuples of errors occur before learning is completed. Define $u_{j,n}$ as follows:

$$u_{j,n} = x_n x_{n+1} \cdots x_{n+j-1} .$$

The variable $u_{j,n}$ indexes the occurrence of j consecutive errors running from trial n through trial $n + j - 1$, since it equals 1 if all the indicator variables x_n through x_{n+j-1} equal 1, and otherwise equals 0. The expectation of $u_{j,n}$ is simply the probability that only errors occur on trials n through $n + j - 1$. There are two possible starting points that give an error on trial n. The possibilities are (a) to start in state E on trial n and stay in state E for $j - 1$ trials; and (b) to start in state 0 on trial n and either stay there for $j - 1$ trials or move into state E after i trials and stay in state E for the remaining $j - 1 - i$ trials. Thus, the probability that $u_{j,n} = 1$, and therefore the expectation of $u_{j,n}$, may be written as follows:

$$E(u_{j,n}) = w_{E,n}\alpha^{j-1} + w_{0,n}\left[(1 - c)^{j-1} + \sum_{i=1}^{j-1} (1 - c)^{i-1} cq\alpha^{j-1-i}\right]. \tag{20}$$

To find the average number of j-tuples of errors, u_j, during the course of learning, Eq. (20) is summed over all trials. The result is

$$u_j = \sum_{n=1}^{\infty} E(u_{j,n}) = \frac{q\alpha^{j-1}}{\theta} + \frac{(1 - c)^{j-1}}{c} + \frac{q\,[\alpha^{j-1} - (1 - c)^{j-1}]}{\alpha - (1 - c)} . \tag{21}$$

Once the u_j are known, predictions of runs of errors are obtained easily. The number of runs of exactly j errors, r_j, is equal to $u_j - 2u_{j+1} + u_{j+2}$, and the total error runs of any length, $R = \Sigma_{j=1}^{\infty} r_j$, is equal to $u_1 - u_2$ (Bush, 1959).

Finally, results are shown for two other sequential statistics that we have found useful for estimating parameters. Derivations are easy and are not detailed. The results are as follows:

$$c_k = E\left[\sum_{n=1}^{\infty} x_n \cdot x_{n+k}\right] = \frac{1}{c}[w_{0,k+1} + w_{E,k+1}] + \frac{q\alpha}{\theta}(1 - \theta)^{k-1} , \tag{22}$$

$$d_k = E\left[\sum_{n=1}^{\infty} (1 - x_n)x_{n+k}\right] = \frac{pq(1 - s)}{\theta}(1 - \theta)^{k-1} . \tag{23}$$

The statistic c_k is the average number of times a pair of errors occurs k trials apart, without regard to the intervening responses. The statistic d_k is the average number of times a success is followed by an error k trials later. These two statistics summarize sequential information akin to that involved in an autocorrelation with lag k.

3. Parameter estimates

In the model there are four parameters, p, c, ε, and s, which, in general, must be estimated from the data. A ready estimate of p is the observed proportion of correct responses over all trials between the first success and the last failure.[1] The estimates of the learning parameters c, ε, and s may be obtained by taking appropriate functions of the observed values of mean T, mean J_0, and mean d_1. Letting $\bar{T}$, $\bar{J}_0$, and $\bar{d}_1$ represent the observed means, we set them equal to the corresponding theoretical expressions, viz.,

$$\begin{aligned} \bar{T} &= \frac{1}{c} + \frac{q}{ps + q\varepsilon}, \\ \bar{J}_0 &= \frac{1}{c} + \frac{q}{p + q\varepsilon}, \\ \bar{d}_1 &= \frac{pq(1-s)}{ps + q\varepsilon}. \end{aligned} \tag{24}$$

By algebraic manipulation of these quantities, it is seen that

$$\frac{\bar{d}_1}{\bar{T} - \bar{J}_0} = p + q\varepsilon. \tag{25}$$

If p has been estimated, then Eq. (25) gives a method for estimating ε. Once ε and p are estimated, then c may be estimated from $\bar{J}_0$ by the relation

$$\hat{c} = \frac{1}{\bar{J}_0 - [q/(p + q\varepsilon)]} = \frac{\bar{d}_1}{\bar{d}_1\bar{J}_0 - q(\bar{T} - \bar{J}_0)}. \tag{26}$$

After p, ε, and c have been estimated, the parameter s may be estimated by solving the expression for $\bar{T}$ or for $\bar{d}_1$. For example, using the expression for $\bar{T}$, the estimate of s is

$$\hat{s} = \frac{q}{p}\left[\frac{1}{\bar{T} - (1/c)} - \varepsilon\right]. \tag{27}$$

There is little doubt that Eqs. (25), (26), and (27) are not the best estimators for the parameters. Their statistical properties are unknown. They are surely biased since each has a theoretically infinite expectation (i.e., each involves the reciprocal of an observed term that conceivably could be zero). These estimators have been used in practice simply because we have not been able to construct any better estimation schemes. The doubtful quality of the estimates should be kept in mind when we later compare the model's predictions with data. It is reasonable to expect that the model would make more accurate predictions if the parameters were estimated more accurately.

---

[1] As this book went to press, we learned from Dr. J. Greeno that this intuitively sensible estimate of p is biased. He showed that the percentage of successes between the first success and the last failure is $p(1 - s)(1 - \theta)$, which equals p only if $s = \varepsilon$. This information arrived too late for us to modify predicted quantities throughout the paper. However, for most of our cases, the bias in p thus introduced appears to be small, and removing the bias would not appreciably change the predictions.

4. Experimental applications

Avoidance learning. Theios (1961, 1963) has used this model to analyze results from an experiment on avoidance learning in rats. Since the analyses are given in detail in Theios' reports, specific results are not duplicated here. The initial critical analyses center upon the responses between the first *CR* (conditioned response) and last non-*CR*. Call these the "intermediate responses" or "intermediate trials." The assumption of a single intermediate state presumes that the intermediate responses may be represented as a stationary, independent series of Bernoulli observations. Three different tests for stationarity of the intermediate responses yielded nonsignificant χ^2 values for the null hypothesis of stationarity. The observed proportion of *CR*'s in the intermediate trials was just slightly over .50. The test for statistical independence of intermediate responses also yielded a nonsignificant χ^2 value. A related test for stationarity and independence comes from examining the distribution of *CR*'s in blocks of four intermediate trials. The observed frequencies were fit quite well by the theoretical binomial distribution with probability .50 of a *CR*. In summary, all these tests on the intermediate responses gave results consistent with the assumption of a single intermediate performance level, represented by a stationary Bernoulli-trials process. For predicting quantitative details of the data, Theios used the simple two-element pattern model of stimulus-sampling theory (Estes, 1959; Suppes and Atkinson, 1960). In the simple model, one of two stimulus patterns is sampled on a given trial, and the response is determined by the state of conditioning of the sampled pattern. With probability c the sampled pattern becomes conditioned to the reinforced response if it is not already so conditioned. These assumptions dictate the following restrictions on the parameters: $c = \varepsilon$, $s = 0$, and, if the two stimulus elements have equal sampling probabilities, $p = .50$. With the estimates $c = \varepsilon = .427$, $s = 0$, and $p = .50$, the model was found to provide a respectable quantitative account of the avoidance data.

Eyelid conditioning. The model has also been used to analyze some data on human eyelid conditioning that have been supplied to us by Dr. I. Gormezano. Though far from perfect, the predictions from the model are promising enough to encourage more effort in this direction of model testing. Approximately 140 college students were run for 70 trials in classical conditioning with a conditioned-stimulus–unconditioned-stimulus (*CS–US*) interval of 500 ms and a relatively strong airpuff as the *US* (200 mm. of mercury). Procedural details may be obtained by reference to Group C in a paper by Moore and Gormezano (1961). Of this pool of 140 subjects, a subset of 114 subjects met a criterion of 10 consecutive *CR*'s at some point in their 70 trials. The preliminary analysis involves only those trials preceding this criterion run for these 114 subjects. We make the idealizing assumption that by the time a subject completes the criterion run of 10 consecutive *CR*'s, he is in absorbing state 1. This assumption can be considered as only an approximation to the actual state of affairs.

The first critical data concern the stationarity of the series of responses between the first *CR* and last failure to make a *CR*. The relevant results are shown in Fig. 3, which is a plot of the average proportion of *CR*'s in blocks of three intermediate trials. The successive estimates appear to be reasonably stationary. The only estimate that is far out of line is in block 12, which was based on results from only seven subjects. The over-all average percentage of *CR*'s during the intermediate trials was 69.6 per cent. The points in Fig. 3 were tested for stationarity by a χ^2 test (Suppes and Ginsberg, 1961). The test yielded a value of $\chi^2 = 14.18$; with 16 degrees of freedom, the probability of such deviations on the null hypothesis of stationarity is greater than .50. About half of this χ^2 value is contributed by the deviant point at the twelfth trial-block.

The χ^2 test does not take into account possible trends of the p-estimates over trials. Since incremental theories predict an upward trend where the three-state model expects no trend, a trend test would seem appropriate. A test suggested to us by Dr. Terry Allen consists of calculating the rank-order correlation coefficient between intermediate trials and the rank of the p-estimates. If the p-estimates show a perfect upward trend over intermediate trials, then this rank-order correlation will be $+1.00$; if there is no trend at all, then the correlation will be zero.

We adopt the procedure, recommended by McNemar (1960, p. 210), of

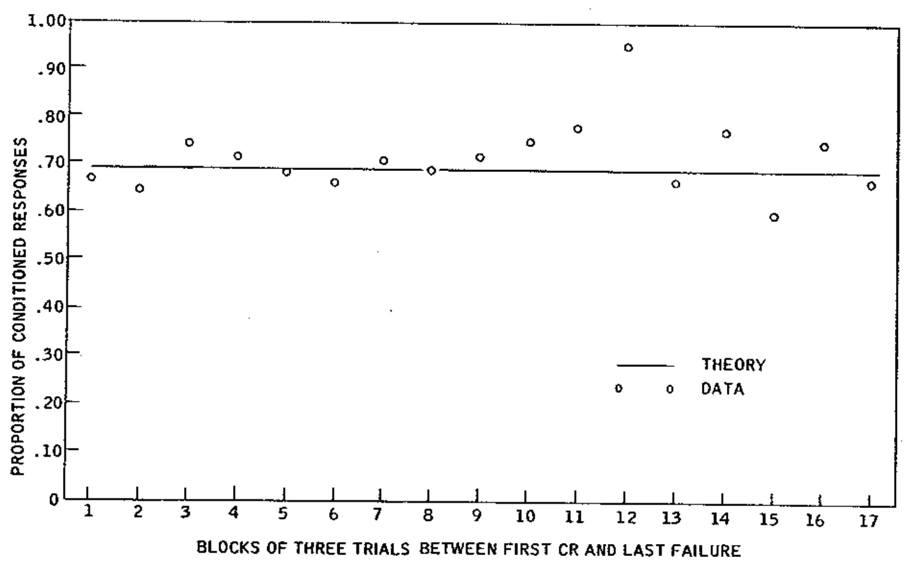

FIG. 3. Proportion of conditioned eyelid responses (*CR*) in blocks of three trials between the first *CR* and the last non-*CR*.

calculating the rank-order correlation, ρ, when the number of paired ranks, N, is 10 or more. The null hypothesis of the three-state model is that the true ρ is zero. The test statistic for an observed sample ρ_s is then

$$t = \rho_s \sqrt{\frac{N-2}{1-\rho_s^2}},$$

which is distributed as Student's t with $N - 2$ degrees of freedom. When N is less than 10, Kendall's tau is used to measure the association between trials and rank of the p-estimates. Probability values on the null hypothesis are then determined by reference to tables in Siegel's book (1956). To return to the p-estimates in Fig. 3: the rank-order correlation is +.17. This value is not significantly different from zero ($t = .71$, df $= 15$, $P > .40$).

A separate test for the statistical independence of the intermediate responses was carried out. Considering all intermediate responses, including the first CR and last failure, we tabulated the frequencies of the four possible pairs of responses on trials n and $n + 1$ over all intermediate trials. The transition frequencies pooled over subjects are shown in Table 1. The χ^2 test on this 2 × 2 table (see Suppes and Ginsberg, 1961) yielded $\chi^2 = 1.65$.

TABLE 1

TRANSITION FREQUENCIES FOR INTERMEDIATE RESPONSES IN THE EYELID-CONDITIONING EXPERIMENT

		Trial $n+1$: *CR*	Trial $n+1$: Failure
Trial n	*CR*	506	252
	Failure	190	113

$\chi^2 = 1.65$, df $= 1$, $P > .15$

With one degree of freedom, the probability is greater than .15 on the null hypothesis of statistical independence. Because of the large number of observations involved (1,061), this test had considerable power to reject the null hypothesis if it were false. This is a conservative test for independence since the χ^2 value could be large if the p-values vary across the individuals that are pooled for the test. To summarize briefly, then, the sequence of intermediate responses forms a stationary, independent Bernoulli series of observations.

For additional analyses, the learning parameters must be estimated. From the intermediate responses, the estimate $\hat{p} = .70$ is obtained for the probability of a CR in the intermediate state. The estimators given in Eqs. (25)–(27) are used along with the observed average values $\bar{T} = 6.43$, $\bar{J}_0 = 3.20$, and $\bar{d}_1 = 2.21$ to obtain estimates of c, s, and ε. The estimates are $\hat{\varepsilon} = 0$, $\hat{c} = .361$, and $\hat{s} = .117$. The fact that $\varepsilon = 0$ means that the effective learning

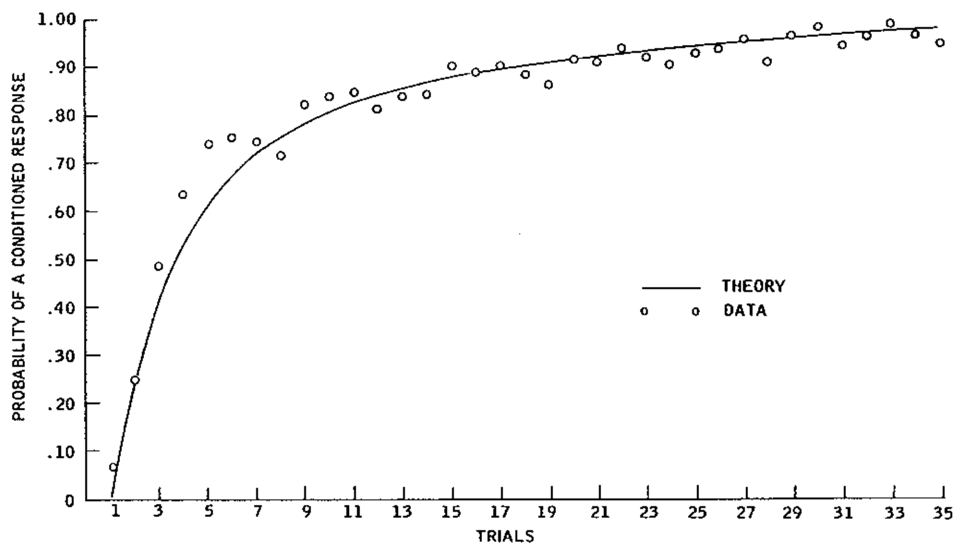

FIG. 4. Predicted and obtained average probability of a conditioned eyeblink response (CR) as a function of trials.

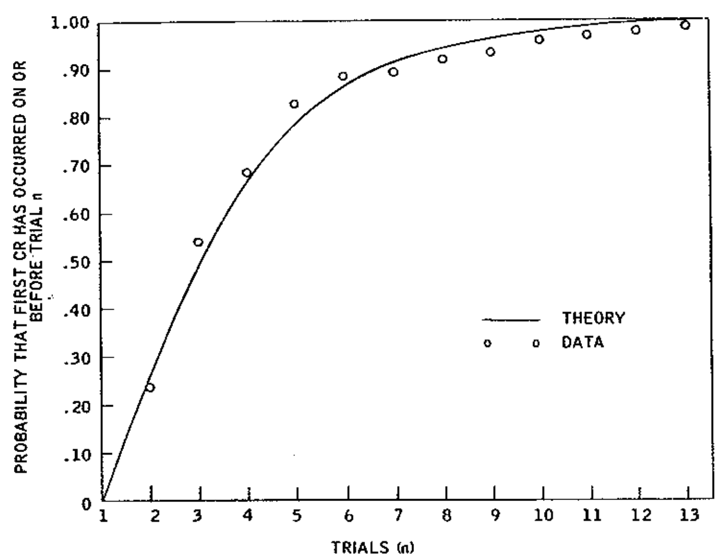

FIG. 5. Predicted and obtained probability that the first CR has occurred by trial n. These curves result from cumulating the probability distribution of the number of failures before the first CR, J_0.

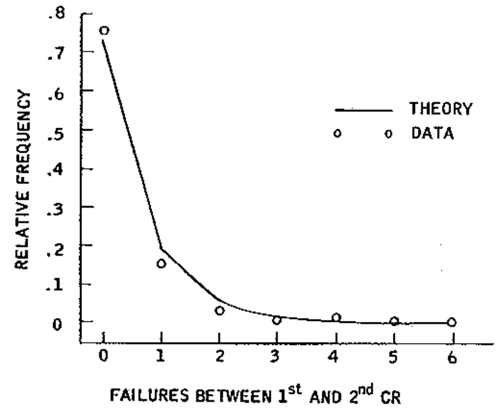

FIG. 6. Predicted and obtained distribution of failures between the first *CR* and the second *CR* in the eyelid-conditioning study, $P(J_1 = i)$.

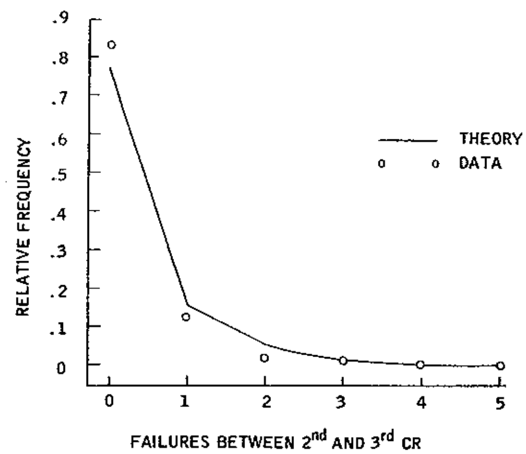

FIG. 7. Predicted and obtained distribution of failures between the second *CR* and the third *CR* in the eyelid-conditioning study, $P(J_2 = i)$.

event in the intermediate state occurs only on *CR* trials; non-*CR* trials are not effective in advancing the subject into the terminal absorbing state.

Using these estimates, a variety of predictions can be made about these data. In Fig. 4 are shown the predicted and observed mean learning curves. Because $\hat{c}$ is relatively large (.361), the subjects quickly move out of the 0 state into the $p = .70$ state. However, the exit from the intermediate to the terminal state is slow, occurring with probability $\theta = ps = .082$ on each trial. Hence, the mean learning curve rises sharply at first, but then drifts only gradually up towards unity.

Figure 5 shows the predicted and obtained probabilities that the first *CR* has occurred by trial n. This is the cumulative distribution of J_0, the number of failures before the first *CR*. The goodness of fit was evaluated by the Kolmogorov-Smirnov one-sample test (Siegel, 1956). The largest discrepancy between observed and predicted values (.05) occurs on trial 3. With an N of 114, the probability of a discrepancy this large is more than .20 on the null hypothesis. Figures 6 and 7 show the predicted and obtained distributions of J_1 and J_2, respectively. Recall that J_k is the number of failures between the kth and $(k + 1)$th *CR*'s. In neither case does the one-sample test detect a significant discrepancy between predicted and observed values.

Finally, Table 2 displays a number of point predictions derived from the model. An "error" may be interpreted as a failure to make a *CR*. An asterisk beside a number indicates that it was used for parameter estimation, and hence the predicted value is forced to equal the observed value at that point. In general, the predictions come reasonably close to the observed statistics. Predictions of sequential statistics, in the lower half of the table, are especially accurate. Predictions are poorest on the standard deviations of total errors and trial of last error. For both statistics, the model predicts less variance than the data exhibit. In fact, in both cases, the observed standard deviations are slightly larger than the corresponding means, whereas the model always predicts standard deviations less than the mean. Thus, the variance

TABLE 2

PREDICTIONS AND DATA FOR THE EYELID-CONDITIONING EXPERIMENT

Statistic	Observed	Predicted
Total errors, $E(T)$	6.43	6.43*
$\sigma(T)$	6.50	4.90
Errors before first *CR*, $E(J_0)$	3.20	3.20*
$\sigma(J_0)$	2.74	2.40
Errors between first and second *CR*'s, $E(J_1)$	.40	.38
$\sigma(J_1)$	.98	.82
Errors between second and third *CR*'s, $E(J_2)$	.27	.33
$\sigma(J_2)$	.69	.77
Trial of last error, $E(n')$	13.28	12.32
$\sigma(n')$	14.72	10.71
Runs of errors, R	3.14	3.26
Runs of 1 error, r_1	1.79	1.88
Runs of 2 errors, r_2	.68	.64
Runs of 3 errors, r_3	.24	.33
Runs of 4 errors, r_4	.21	.17
Error-error pairs, c_k		
1 trial apart, c_1	3.29	3.17
2 trials apart, c_2	2.68	2.60
3 trials apart, c_3	2.36	2.20
CR-error pairs, d_k		
1 trial apart, d_1	2.21	2.21*
2 trials apart, d_2	2.02	2.08
3 trials apart, d_3	1.91	1.91

* Number used for parameter estimation.

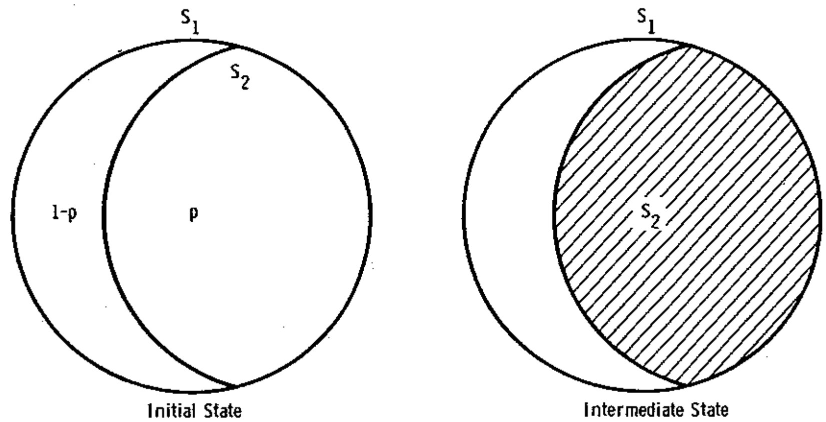

FIG. 8. Venn diagrams of stimulus sets for classical conditioning situations.

predictions could not be brought into line with the data by readjusting the parameter estimates. The only plausible explanation for this excessive variance is that the parameters vary across subjects, whereas, for this application, it has been assumed that all subjects have the same parameter values.

The fact that $\hat{\varepsilon} = 0$ for these data agrees with a theoretical interpretation of classical conditioning proposed by W. K. Estes (personal communication, 1960). Within the stimulus-sampling framework, it is assumed that the onset of the conditioned stimulus gives rise to a set of stimulus elements, S_1. A subset $S_2 \subseteq S_1$ of these elements is present at the time the unconditioned stimulus occurs, evoking the response to be conditioned. The sets S_1 and S_2 may be represented by Venn diagrams as in Fig. 8. Initially none of the elements of S_1 is conditioned to the response, and the response is evoked only in the presence of the subset S_2 of elements (i.e., when the *US* is presented). With probability c the elements of S_2 become conditioned to the response on a single trial, leading to the intermediate state. The conditioning of the S_2 subset is illustrated in Fig. 8 by shading in the S_2 set. After the S_2 subset is conditioned, presentation of the *CS* leads to a conditioned response with probability p, the measure of the subset S_2 within the over-all set S_1. In order for the unique elements, $S_1 - S_2$, to become conditioned, the response must occur when they are present. Since the elements in $S_1 - S_2$ are not present at the time the unconditioned stimulus–unconditioned response (*US–UR*) occurs, the subject cannot leave the intermediate state on a non-*CR* trial (i.e., $\varepsilon = 0$). The only opportunity for conditioning of the set $S_1 - S_2$ arises when an anticipatory *CR* occurs to the set S_1 (and, necessarily, to the set $S_1 - S_2$). Let s represent the probability that an anticipatory *CR* results in effective conditioning of the elements of $S_1 - S_2$. The probability that the subject goes from the intermediate to the terminal state on a particular trial is *ps*. Thus, it can be seen that there is a direct correspondence between this abstract theory of stimulus sampling and conditioning and the three-state empirical model that has been applied to these data obtained from eyelid conditioning.

In evaluating the model's account of these data on eyelid conditioning, the reader should bear in mind the explicit limitations imposed by the selection of data. A number of problems of interpretation stem from the fact that a subject's protocol is examined only until he reaches the trial on which he meets a criterion of 10 consecutive *CR*'s. Subjects who do not meet this criterion within 70 trials are excluded from the analysis. The interpretation problem centers about the question of whether subjects in fact have asymptotes of unity in this situation. Subjects might actually have asymptotes of, say, .90, but if subjects are run long enough, they would eventually meet the criterion of 10 consecutive *CR*'s. However, it can be said in defense of our unity-asymptote approximation that if subjects attained criterion by "chance" only, then the average trials to criterion should have been much larger than that predicted by the model.

Response shift in paired-associate learning. The model has also been applied to data on response shift in paired-associate learning following an initial learning series.[2] Data have been obtained under several conditions of number of available responses and type of training (correction or noncorrection). We present first the details of one of these sets of data; then we summarize briefly the salient results on intermediate responses for three of the other sets of data.

In the initial experiment, 25 university students learned a 12-item paired-associate list. The stimuli were the printed names of colors, such as "red," "blue," etc. The responses were the numerals 1 to 4, and each numeral was assigned as the correct response to a randomly selected set of three stimuli for each subject. The subject was required to make a response to every stimulus on every trial. After his response, the experimenter informed the subject of the correct response: just two seconds later the next stimulus item was pre sented. A trial consisted of one presentation of each of the 12 items in random order.

When the subject attained a criterion of 10 consecutive correct responses to a particular stimulus item, on the next trial the correct response for that stimulus was changed to one of the other three numerals, chosen at random. The subjects were not informed that the correct answers would be changed, and the shifts occurred independently for different stimuli in the list of 12 items. The experiment continued until the last-shifted item had been responded to correctly for 10 consecutive trials following its shift.

[2] The following paired-associate response-shift experiments were conducted by John Theios at the University of Texas with the aid of Miss Mary Gunter and Mr. Edwin Davison.

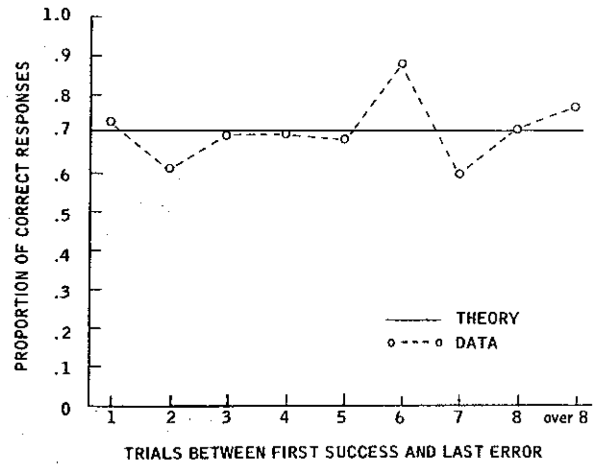

FIG. 9. Proportions of correct responses in trials between the first correct response and last error for the response-shift experiment.

We will be concerned with only the data following the response shift for individual items. The unit of analysis will be the trial-sequence of responses a subject gives to a single item, starting with the trial on which the response shift first occurs for that item.

The critical data on stationarity of the intermediate responses are shown in Fig. 9. The mean proportion of correct responses during the intermediate trials is .708. The χ^2 test for stationarity applied to the data in Fig. 9 yields a $\chi^2 = 10.11$; with 8 degrees of freedom, $P > .20$ on the null hypothesis. The major contributors to the χ^2 value obtained are the points at trials 2 and 6, which are off the constancy line. To test for an upward trend to the estimates of success probabilities, the association between ordinal trials and rank of estimate was calculated by Kendall's tau. The value of tau was .11, which is not significantly different from zero ($P > .66$). Another test for stationarity divided the intermediate responses for each subject-item into a first half and a second half. This division necessarily restricts the analysis to subject-items having at least two trials between the first success and last failure. Pooling over subject-items, the mean proportion of successes in the first half was .691 and in the second half was .718. These estimates do not differ significantly. Finally, all sequences entering into the above analysis (i.e., those sequences with at least two intermediate trials) were scored $+$, $-$, or 0, according to whether the number of successes in the second half was larger, smaller, or equal, respectively, to the number of successes in the first half for that subject-item sequence. Of the 65 sequences entering this analysis, 19 were scored as $+$, 19 scored as $-$, and 27 scored as 0. These results are consistent with the assumption of a stationary series of intermediate responses.

In Table 3 are given the transition frequencies for testing the statistical

TABLE 3
TRANSITION FREQUENCIES FOR INTERMEDIATE RESPONSES IN THE RESPONSE-SHIFT EXPERIMENT

		Trial $n+1$ Success	Trial $n+1$ Error
Trial n	Success	227	90
	Error	67	32

$\chi^2 = .58$, df $= 1$, $P > .30$

independence of the intermediate responses. The χ^2 value for the frequencies in Table 3 is .58; with one df, $P > .30$ on the null hypothesis of statistical independence. Another test of independence and stationarity is shown in Table 4. This records the obtained and predicted frequencies of the number of successes in blocks of four intermediate trials. The theoretical binomial distribution with $p = .708$ gives a satisfactory fit to the observed frequencies ($\chi^2 = 1.97$, df $= 2$, $P > .30$).

Estimates of the model parameters were $\hat{p} = .708$, $\hat{c} = .690$, $\hat{\varepsilon} = .329$, and $\hat{s} = .220$. Using these parameters, we tested predictions of other details. In Fig. 10 are shown the predicted and obtained distributions of errors before the first success. Figure 11 gives the probability of no errors following trial n; this is the cumulative distribution of the trial of the last error. In neither Fig. 10 nor Fig. 11 are we able to reject the fit of the model to the data, according to the Kolmogorov-Smirnov one-sample test.

TABLE 4
NUMBER OF SUCCESSES IN BLOCKS OF FOUR BINOMIAL TRIALS BETWEEN THE FIRST SUCCESS AND THE LAST ERROR IN THE RESPONSE-SHIFT EXPERIMENT

Number of Successes	Obtained Frequency	Predicted Binomial Frequency
0	0	.5
1	5	5.3
2	18	19.3
3	37	31.1
4	15	18.8
Total	75	75.0

$\chi^2 = 1.97$, df $= 2$, $P > .30$

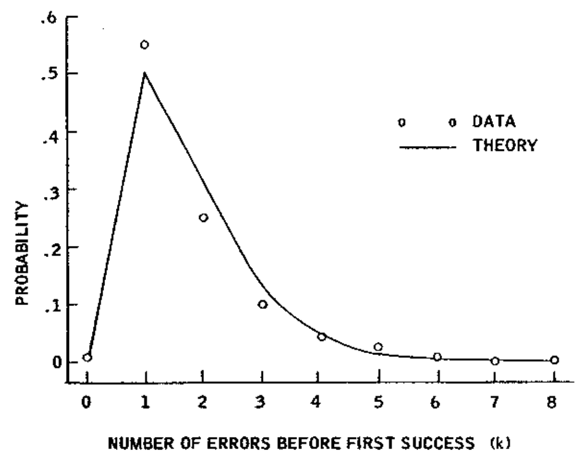

FIG. 10. Predicted and obtained probability distribution of the number of errors before the first success in the response-shift experiment, $P(J_0 = k)$.

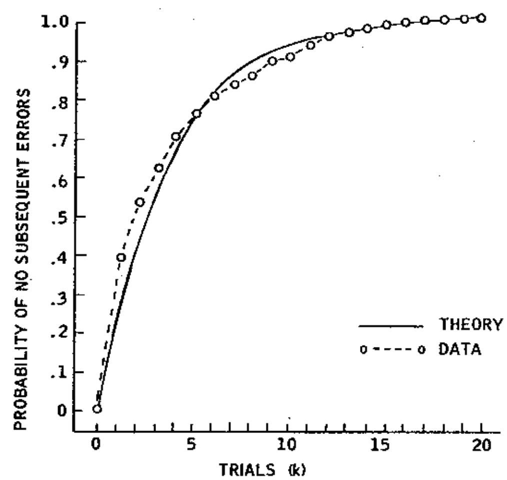

FIG. 11. Predicted and obtained probability of no errors following trial n of the response-shift experiment, $P(n' \leq k)$.

A list of point predictions from the model is given in Table 5. Inspection of Table 5 reveals that the model gives a satisfactory account of these data. It may also be noted that, as in the previous case, the model predicts less variability than is displayed by the data.

In order to provide further evidence concerning the existence of an intermediate state, results from some other paired-associate experiments are presented briefly. In all cases, the stimuli were the 12 color names and the responses were the first few numerals. The data in each case come from a

TABLE 5

COMPARISON OF RESPONSE-SHIFT DATA WITH PREDICTIONS

Statistic	Obtained	Predicted
Total errors, $E(T)$	2.60	2.60*
$\sigma(T)$	1.98	1.50
Errors before first success, $E(J_0)$	1.81	1.81*
$\sigma(J_0)$	1.29	1.02
Trial of last error, $E(n')$	4.00	4.03
$\sigma(n')$	4.05	3.15
Runs of errors, R	1.60	1.64
Runs of 1 error, r_1	1.02	1.01
2 errors, r_2	.35	.42
3 errors, r_3	.12	.15
4 errors, r_4	.05	.05
5 errors, r_5	.03	.02
Alternations of successes and failures	2.21	2.28
Error-error pairs, c_k		
1 trial apart, c_1	1.00	.97
2 trials apart, c_2	.65	.62
3 trials apart, c_3	.57	.43
4 trials apart, c_4	.41	.30
Errors between kth and $(k + 1)$th successes, $E(J_k)$		
$k = 1$	.22	.25
$k = 2$	.18	.17
$k = 3$	.09	.12
$k = 4$	.08	.08
Success-error pairs, d_k		
1 trial apart, d_1	.64	.64*
2 trials apart, d_2	.44	.48
3 trials apart, d_3	.33	.36
4 trials apart, d_4	.28	.27

* Number used for parameter estimation.

reversal-shift series that follows initial learning to a given criterion. In one case, there were two response alternatives, and subjects were trained under noncorrection conditions—that is, following the subject's response, the experimenter simply said "Correct" or "Wrong." The results on responses between the first success and the last error are shown in Fig. 12. These estimates exhibit stationarity at a mean p value of .66. The χ^2 test for stationarity yielded the value $\chi^2 = 6.89$ with 9 df ($P > .50$). The test for a systematic trend in the probability estimates yielded a rank-order correlation of .005 between ordinal rank of the trial and rank of the estimate.

Another condition studied involved three responses with noncorrection

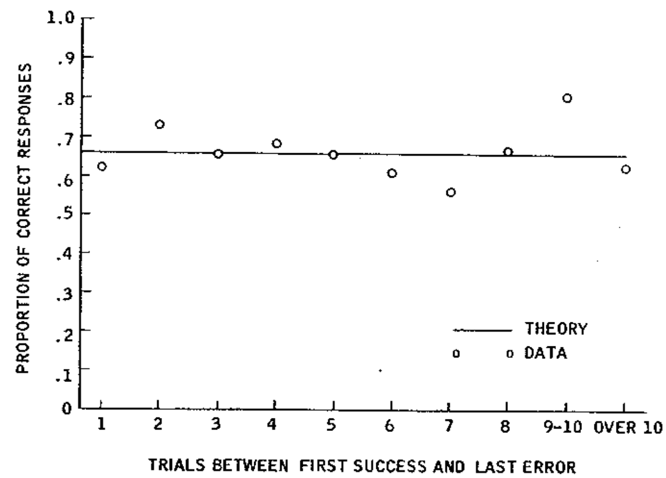

FIG. 12. Proportion of correct responses on trials between the first correct response and the last error in the two-response reversal experiment.

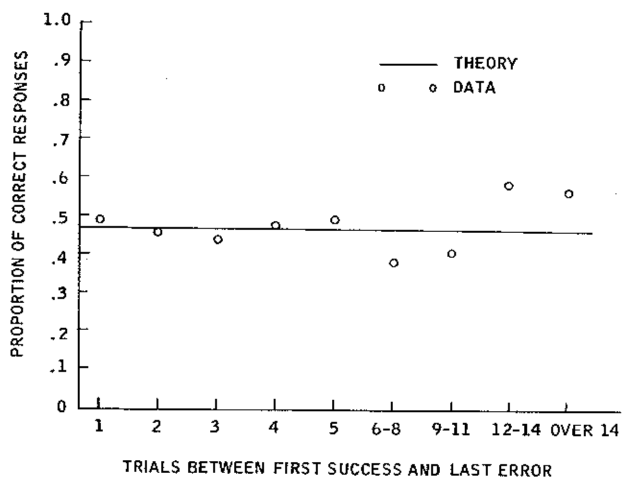

FIG. 13. Proportion of correct responses on trials between the first correct response and the last error in the three-response reversal experiment.

training. The intermediate responses during reversal learning are shown in Fig. 13. These results exhibit stationarity at the mean $p = .47$. The χ^2 test for stationarity yielded a value of $\chi^2 = 5.61$; with df $= 8$, $P > .50$ on the null hypothesis. Kendall's tau between trial rank and probability rank was .11. The corresponding normal deviate is .43; so $P > .66$ on the null hypothesis of zero correlation.

A final condition studied involved four responses and noncorrection training. The graph of probability of a correct response during intermediate trials is shown in Fig. 14. The mean probability of a correct response during intermediate trials was .55. The stationarity test gave $\chi^2 = 3.25$ with 6 df ($P > .70$). Kendall's tau is $-.24$ for the association between trials and rank of probability estimates, indicating a moderate downward trend in success probabilities. With only seven cases, the tau does not differ significantly from zero ($P > .58$).

Let us summarize briefly these results on response-shift following initial paired-associate learning. In the four cases we have examined, the responses between the first correct response and last error were tested for stationarity. To increase the over-all power of the test of the same null hypothesis, the independent χ^2 values for stationarity for the four conditions of the experiment may be added. This gives a total χ^2 value of 25.85 with a total of 31 degrees of freedom. The total value is not significant ($.70 > P > .50$). Because of the power of this combined χ^2 test, it would seem reasonable to accept the null hypothesis. Responses between the first correct response and last error can be represented by a single performance level (i.e., within a particular condition, p does not change over intermediate trials). The specific levels (p's) of intermediate performance vary across experimental conditions, but that fact is immaterial to the general point that each set of data may be represented by only three values of response probability.

Escape learning. As a final illustration of the model, we present an analysis of data obtained from a two-choice shock-escape learning situation employing rats as subjects. We are indebted to Dr. H. W. Coppock for

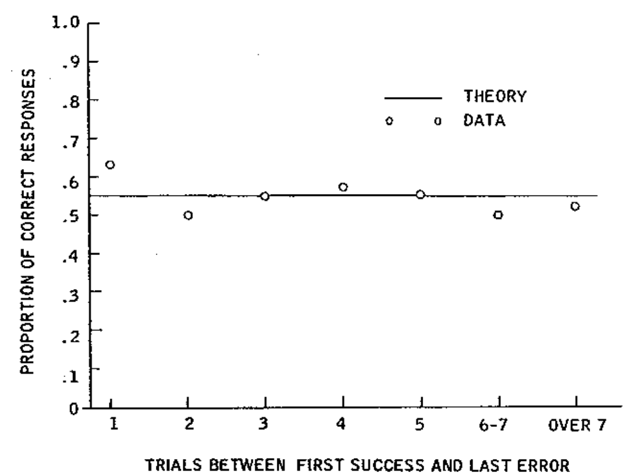

FIG. 14. Proportion of correct responses in trials between the first correct response and the last error in the four-response reversal experiment (noncorrection procedure).

making these data available to us. In the experimental situation, the rat was strapped down to a board and completely restrained except for being able to move its head to the right or left. A photocell arrangement was set up to pulse counters whenever the rat moved its head 50 degrees of arc right or left away from the center, median head position. These right or left head-turns of 50 degrees are identified with the two response classes in the model. After the rat became accustomed to this apparatus, it was subjected to painful electric shocks delivered to its tail. The correct response (right or left) for a given rat was determined by its choice on the first shock-trial; specifically, the correct choice was defined as the choice opposite the first one made by the rat on the first trial. This procedure trained each subject against its initial response bias and ensured an "error" on trial 1. A correction training method was used—that is, the shock on each trial continued until the rat finally made the correct choice, at which time the shock was terminated. Subsequent trials (shocks) occurred at 1-minute intervals. A necessary condition for starting a trial was that the rat have its head in the center position. The performance measure analyzed is the first response (correct or incorrect) on each trial. The analysis is carried out for the responses before a criterion run of 10 consecutive correct responses and includes 276 rats (of a total 304) that met this criterion. For further details of procedure, the reader may refer to an experimental report by Coppock and Freund (1962).

For purposes of analysis in terms of the model, we consider that at the point of the first error on trial 1 the subject is in state 0, and at the end of the criterion run of 10 correct the subject is in state 1. The critical information concerns the stationarity of responses between the first correct response and the last error for each subject. The estimates of success probabilities over successive intermediate trials are shown in Fig. 15. The estimates prove to be stationary with a mean p-value of .75. The χ^2 test for stationarity yields a value of 5.38; with 8 df, $P > .70$ on the null hypothesis of stationarity. A second analysis partitioned the intermediate trials of each subject into quartiles for plotting a Vincent curve (Suppes and Ginsberg, 1961). This restricts the analysis to those subjects having at least four trials between their first success and last failure and ensures that each point in the curve is based on the same number of observations. The quartile Vincent curve of intermediate responses was stationary ($\chi^2 = 2.46$, df $= 3$, $P > .30$); when the first two and last two quartiles were pooled, the percentages of successes were identical in the two halves. To test for an upward trend to the estimates in Fig. 15, Kendall's tau was $+.33$ for the association between trials and rank of success probability. With nine pairs of ranks, a tau this large has a probability greater than .22 on the null hypothesis of zero correlation. A final test for stationarity and independence was the prediction of the distribution of number of successes in blocks of four intermediate trials. The results are shown in Table 6. The observed frequencies are adequately described by the binomial distribution ($\chi^2 = .72$, df $= 2$, $P > .50$).

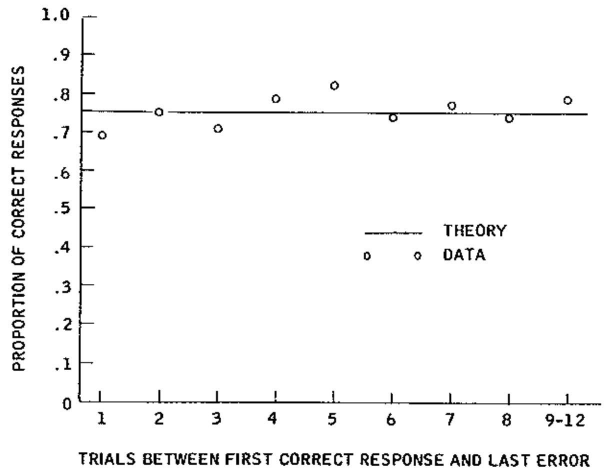

FIG. 15. Responses over intermediate trials for rats learning an escape habit. (Data from Coppock.)

For predicting quantitative details, the parameters were estimated to be $\hat{p} = .75$, $\hat{c} = .60$, $\hat{s} = .30$, and $\hat{\varepsilon} = 0$. Again, as in the eyelid-conditioning experiment discussed earlier, the effective learning event in the intermediate state occurs only on success trials. The predictions of the model for these data will be displayed in a series of graphs. Figure 16 is a graph of the average probability of an error over successive practice trials. Reflecting the values of $c = .60$ and $\theta = ps = .225$, the curve drops rapidly at first but changes its slope after trial 3. In Fig. 17 are shown the predicted and obtained distributions of J_0, the errors before the first success; in Fig. 18, the distribution of total errors; and in Fig. 19, the distribution of the trial of the last error,

TABLE 6

NUMBER OF ERRORS IN BLOCKS OF FOUR INTERMEDIATE TRIALS

Number of Errors	Obtained Frequency	Predicted Frequency
0	36	33.56
1	39	43.01
2	22	20.67
3	4	4.41
4	1	.35
Total	102	102.00

$\chi^2 = .72$, df $= 2$, $P > .50$

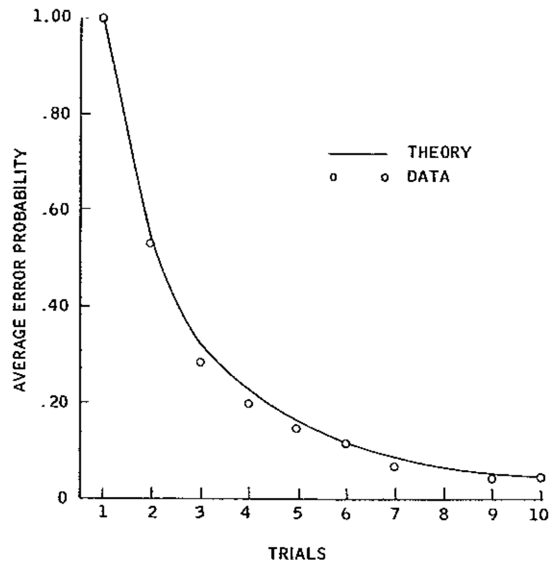

FIG. 16. Average probability of an error over successive practice trials. (Data from Coppock.)

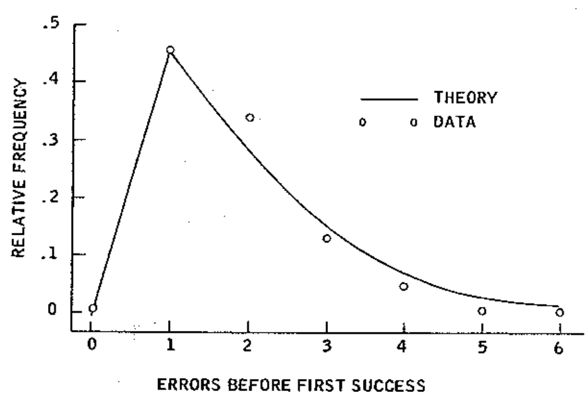

FIG. 17. Probability distribution of number of errors before the first correct response (J_0). (Data from Coppock.)

preceding the criterion run of 10 consecutive successes. In each case, the model gives a respectable approximation to the observed distributions; conventional statistical tests for goodness of fit reveal no significant discrepancies between predicted and obtained values. In summary, then, the three-state model gives a good quantitative account of these data on escape learning in rats.

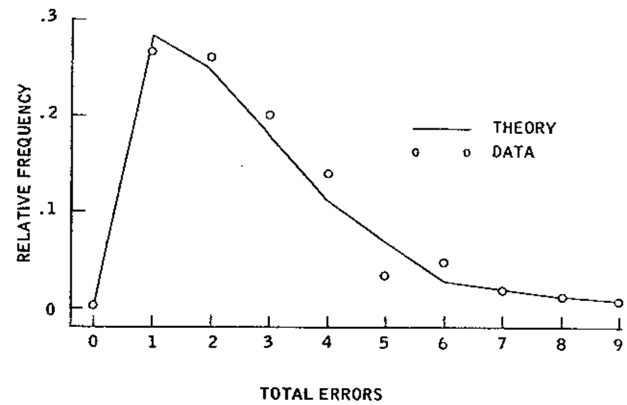

FIG. 18. Probability distribution of total errors (T) per subject. (Data from Coppock.)

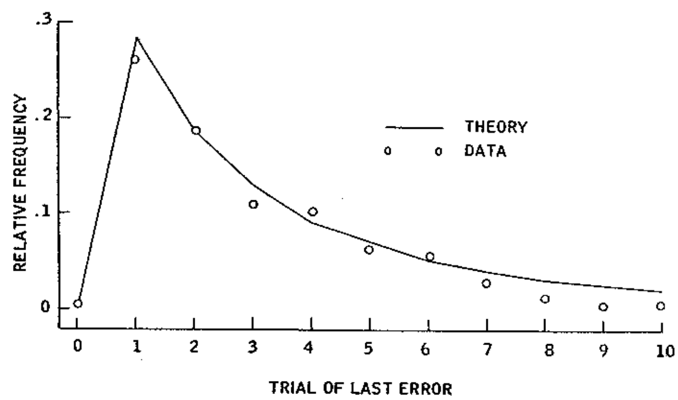

FIG. 19. Probability distribution of the trial of the last error (n'). (Data from Coppock.)

5. Discussion

We have proposed that performance in the experimental situations reviewed can be represented by three discrete performance levels, 0, p, and 1. In these terms, learning consists of two all-or-nothing transitions from lower to higher levels of response probability. This notion seems to have originated with Estes (1959), who also introduced the technique of representing learning by discrete Markov processes, where each state of the process corresponds to a performance level. It was because of Estes' prior theoretical work that we were led to examine our data for evidence of an intermediate performance level. In truth, we have been astonished by the consistency with which such evidence has appeared throughout the range of data examined.

The reader will note that the evidence comes from experimental situations in which initially the probability of a correct response is zero and asymptotically it is unity. Such zero-to-one situations possess an important advantage for our methods of data analysis. The arrangement enables one to identify responses between the first success and last failure as occurring in the intermediate state. The importance of this identification can be understood if one imagines trying to test decisively the notion of a single intermediate state for learning situations in which the initial response probability is greater than zero, or the asymptote is less than unity, or both. In such cases, the evidence would have to be of a more indirect nature, e.g., predicting quantitative details of a variety of statistics. Although we have several verbal learning experiments of this sort to which the three-state model (with initial p greater than zero) has been successfully applied, they are not presented, because their evidence is indirect on the question of primary interest to this paper (cf. Norman, this volume, p. 173).

Our purpose in this paper was the empirical one of reporting evidence regarding an intermediate performance state in several experimental situations. The fact has been emphasized at the expense of specific theoretical interpretations. This approach was adopted because acceptance of the fact entails no particular theoretical commitments. However, we would be remiss not to point out that data showing an intermediate performance level can be interpreted within the framework of stimulus-sampling theory. One specific interpretation in these terms was given on pp. 16–17. Facts about intermediate performance levels could also be interpreted in terms of the multi-stage models of Restle (this volume, p. 116). In constructing and testing the three-state model, we have suppressed the stimulus-sampling rationale and have presented it simply as a descriptive model about learning.

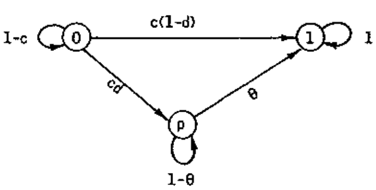

FIG. 20. Alternative model tested and rejected.

The learning model presented exploits the notion of an intermediate state in an obvious way. Certain general Markovian properties were imposed regarding transition probabilities among the states, and the resulting model provided a fairly adequate description of the data on which it was tested. The specific form of the model is not entirely arbitrary since we have been able to reject various plausible alternative three-state models. For example, one plausible alternative model we have considered permits a direct, one-trial transition from the starting state to the terminal absorbing state. This alternative is diagramed in Fig. 20. Here it is assumed that with probability $1 - d$, the subject skips the intermediate p-state, going directly to state 1. This model was applied to data from the avoidance-conditioning, eyelid-conditioning, and response-shift experiments. For the avoidance and response-shift data, the estimate of d was slightly greater than one. Within the restricted unit interval, the best estimate was $d = 1$. For the eyelid-conditioning experiment, the estimate of θ was zero, making all limit expressions inappropriate (e.g., total errors, trial of last error). A modified estimation technique, setting $\theta = 0$, gave the estimate $\hat{d} = 1.68$, which is outside the allowable bounds on probabilities. Thus, even though we allow the possibility of a direct transition from state 0 to state 1, these data indicate that all the subjects go through the intermediate p-state.

The alternative classes of learning models that could be considered for these data are the continuous or incremental theories such as the linear-operator models. Although extensive comparisons have not been undertaken, it seems evident that nearly all continuous models would be rejected by these data. In particular, from continuous models one would expect performance to improve monotonically over trials between the first success and the last error. Such upward trends simply failed to materialize in any of the studies. Our tests for such trends were the χ^2 and the rank-order correlation between intermediate trials and response probabilities. In none of the seven cases was this correlation significantly different from zero, a result in line with the stationarity assumption.

In order to give the continuous model a fair test, one must consider the possible effects on intermediate responses of individual differences in learning rates within the model. The possibility that our plots of forward stationarity curves capitalize on selection of slow learners at the later points on the curve, since fast learners are dropped out early, must be considered. Conceivably, the differential selection of slow learners could balance out the expected increase, giving rise to an approximately flat curve of average percentage of successes over the intermediate trials.

In an attempt to forestall this type of objection, we have run some Monte Carlos for a particular continuous learning model, viz., the single-operator linear model (e.g., Bush and Mosteller, 1955). Similar results would probably be obtained from any continuous model. Let $q_{i,n}$ represent the probability of an error on trial n by subject i, and let α_i (between 0 and 1) be the learning rate for subject i. Then the assumption of the single-operator model is that $q_{i,n} = \alpha_i^{n-1}$. We have carried out 50 Monte Carlo runs; each run will be called a "stat-subject." The distribution of α_i in the 50 runs was 4 stat-subjects at an α of .90, 12 at .85, 18 at .80, 12 at .75, and 4 at .70. This is a symmetric, unimodal distribution with a mean α of .80, a value appropriate to the range of data we have been reporting. Each stat-subject was run until a criterion of 10 consecutive correct responses was achieved. The crucial analyses center upon the responses between the first correct response and the last error.

Figure 21 shows the plot for these linear stat-subjects of the conventional forward curve for responses between the first correct response and last error. Relatively few observations are involved after Trial 9, so the later estimates are somewhat unstable. The over-all mean percentage correct during the intermediate trials was .71. The estimates in Fig. 21 are obviously nonstationary

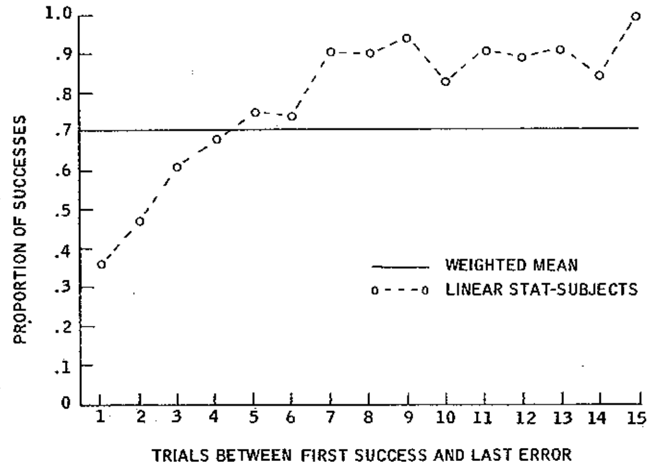

FIG. 21. Proportion of success over intermediate trials for 50 stat-subjects generated by the linear model.

Grouping some of the later points for the χ^2 test (to keep expected frequencies above 5), on the test for stationarity we obtain $\chi^2 = 60.60$; with 12 df, $P < .001$. Similarly, the rank-order correlation between trials 1 to 15 and probability of success was $+ .81$. This gives a Student's $t = 4.43$, df = 13, $P < .001$. Finally, a comparison was made within each stat-subject of the first and second halves of its intermediate trials. This necessarily restricts analysis to those stat-subjects having at least two intermediate responses. Thirty-seven of the pool of 50 met this restriction. When we compared the number of successes in the second and first halves for each stat-subject, we found that 30 went up, 2 went down, and that for 5 there was no difference in number of successes in the two halves. Leaving out the ties, the sign test assigns a one-tailed probability of .0005 to a discrepancy this large or larger on the null hypothesis that signs of + and − are equally likely.

Our purpose in discussing these linear stat-subjects is to demonstrate two points: first, that the argument of selection-artifacts does not really rescue the continuous models from the stationarity data; and second, that the statistical tests we routinely use to assess stationarity of intermediate responses have considerable power to reject the null hypothesis when it is false, even when only 50 learning sequences are involved. All the experiments reported here have involved at least 100 learning sequences; therefore the power of the tests for stationarity has been very high.

In principle, one could argue that there might exist some peculiar distribution of learning rates for some unspecified continuous model that would produce a horizontal series of p-estimates over intermediate trials. One cannot gainsay such an existence claim, but it devolves upon the claimant to demonstrate its existence and to show the plausibility of such unique distributions occurring throughout our varied and numerous experiments. In the absence of such a demonstration, we propose the simple alternative that performance in these experimental situations can be adequately represented by three discrete performance levels.

REFERENCES

Bower, G. H. Application of a model to paired-associate learning. *Psychometrika*, 1961, **26**, 255–280.

Bower, G. H. An association model for response and training variables in paired-associate learning. *Psychol. Rev.*, 1962, **69**, 34–53. (a)

Bower, G. H. Some experiments related to a learning model. Paper read at Western Psychol. Ass., April 1, 1962. (b)

Bush, R. R. Sequential properties of linear models. In R. R. Bush and W. K. Estes (Eds.), *Studies in mathematical learning theory*. Stanford, Calif.: Stanford Univer. Press, 1959. Pp. 215–227.

Bush, R. R., and Mosteller, F. *Stochastic models for learning*. New York: Wiley, 1955.

Coppock, H. W., and Freund, J. E. All-or-none versus incremental learning of errorless shock escapes by the rat. *Science*, 1962, **135**, 318–319.

Estes, W. K. Component and pattern models with Markovian interpretations. In R. R. Bush and W. K. Estes (Eds.), *Studies in mathematical learning theory*. Stanford, Calif.: Stanford Univer. Press, 1959. Pp. 9–52.

Maatsch, J. L. Learning and fixation after a single shock trial. *J. comp. physiol. Psychol.*, 1959, **52**, 408–410.

McNemar, Q. *Psychological statistics*. (2nd ed.) New York: Wiley, 1960.

Moore, J. W., and Gormezano, I. Yoked comparisons of instrumental and classical eyelid conditioning. *J. exp. Psychol.*, 1961, **62**, 552–559.

Siegel, S. *Nonparametric statistics for the behavioral sciences*. New York: McGraw-Hill, 1956.

Suppes, P., and Atkinson, R. C. *Markov learning models for multiperson interactions*. Stanford, Calif.: Stanford Univer. Press, 1960.

Suppes, P., and Ginsberg, Rose. A fundamental property of all-or-none models; binomial distribution of responses prior to conditioning, with applications to concept formation in children. Psychology Series, Technical Report No. 39, Institute for Mathematical Studies in the Social Sciences, Stanford Univer., 1961. (Published in *Psychol. Rev.*, 1963, **70**, 139–161.)

Theios, J. A three-state Markov model for learning. Psychology Series, Technical Report No. 40, Institute for Mathematical Studies in the Social Sciences, Stanford Univer., 1961.

Theios, J. Simple conditioning as two-stage all-or-none learning. *Psychol. Rev.*, 1963, in press.