A Model for Temporal Generalization and Discrimination[a]

ERIC G. HEINEMANN

Department of Psychology
Brooklyn College of the City University of New York
Brooklyn, New York 11210

This paper describes a model for memory and decision processes that has been applied to the acquisition and performance of discriminations by pigeons in a variety of situations. The phenomena to which the model has been applied include probability learning, the acquisition of easy and difficult discriminations, discrimination reversal, probabilistic discrimination learning, stimulus categorization,[1] and absolute identification of stimuli.[2]

The situations considered in the past have involved stimuli that differ in intensity. In the present paper I consider the application of the model to stimuli that differ in duration, and shall focus on variables governing performance at asymptotic levels rather than on acquisition processes.

The prototypical situation considered in the original development of the model involves the presentation, on each of a series of discrete trials, of one of two stimuli, $S1$ or $S2$, that differ in a single attribute. The subject is rewarded for making one response, $R1$, when presented with $S1$, and an alternative response, $R2$, when presented with $S2$. It is assumed that the acquisition of the discrimination between $S1$ and $S2$ involves two stages. The first stage is the presolution period, represented by a series of trials, at the beginning of training, during which there is no evidence of a developing discrimination. The animal is regarded as functioning, during this stage, as a detector of the statistical association the experimenter has arranged among the environmental stimuli, $S1$ and $S2$, selected behaviors (responses), $R1$ and $R2$, and the consequences of these behaviors (reward and nonreward, for example). During the presolution period the subject learns to attend to the stimulus dimension that is correlated with the differential consequences of its behavior. A formal treatment of this matter, based on the theory of sequential analysis developed by A. Wald[3] and his associates, has been presented by Heinemann.[4]

The presolution period ends when the subject discovers that variation in the dimension along which $S1$ and $S2$ differ is related to the outcome of its choice behavior. The theory assumes that the processes occurring after the end of the presolution period involve a memory that has a limited storage capacity (to be referred to as the limited capacity memory or LCM). The principal assumptions of the second-stage model are as follows:

1. Storage and the representation of events in memory. Information gathered on each trial of an experiment is assumed to be placed in a memory that has a limited number of storage locations. Each record placed in this memory during the course of discrimination training contains information concerning the *discriminative stimuli* that analysis done during the presolution period has shown to be predictive of the

[a]This work was supported by Grant MH18246 from the National Institute of Mental Health and by Grant 14002 from the Professional Staff Congress—City University of New York Research Award Program.

outcomes of behavior, the *response* made, and the *reward* received. One such record is entered into the LCM on every trial and is said to occupy a "storage location." The location to which each record is sent is selected randomly, and any record occupying a storage location will be destroyed ("overwritten") when a new record is entered at that location.

The experiments considered in this paper involve the presentation of sets of stimuli that differ in duration. It is assumed that the duration of each stimulus is measured on an internal clock. The durations read from this clock constitute stimuli to a sensory system that has an output proportional to the logarithm of the stimulus. This is a form of Fechner's law,

$$S = c \log t,$$

where S is subjective time, t is the duration read from the clock, the constant $c = 1/\log(1 + W)$, and W is the Weber fraction.

The clock is assumed to be perfectly accurate, but the response of the sensory system to a constant input from the clock varies from trial to trial. It is assumed that, over repeated presentations, each stimulus induces sensory effects, S, that are normally distributed with different means, $\mu_1, \mu_2, \ldots, \mu_n$, and a common standard deviation σ. The sensory effect experienced on each trial is stored in the LCM.

The responses made just before receipt of reward are assumed to be represented in memory in the form of sensory information associated with these responses, such as the visual characteristics of the device that was pecked or pressed, such as its color, position, and so forth. In some situations, such as those considered by Chase,[2] the manner in which response information is represented plays an important role in the analysis. In the present paper, however, the responses will be treated as though they were represented in memory simply by the labels $R1$ and $R2$. For now, outcomes are treated as though they were represented in memory simply as reward (positive) and nonreward (negative).

2. *Retrieval.* It is assumed that during each trial the subject draws a small sample of records from the LCM. The choice of response is based solely on the information contained in this sample of records. The samples are assumed to be independent random samples with the following restriction: The subject is assumed to draw records from the LCM one at a time until a fixed number of positive records (records showing that a reward was received) has been obtained. To say that the subject draws a sample of records is not intended to imply that the records in question are removed from the LCM. The idea is that these records are "looked at" or are "copied" for use in a working memory.

3. *Response selection.* The choice of response on each trial involves not only a sample of records from memory, but also the sensory effect present on the continuum that the analysis done during the presolution period has shown to be predictive of outcomes, for example, the currently experienced duration of a visual signal. This sensory effect will be called the "current input."

Before describing how the current input and the records drawn from the LCM are used in response selection, it is necessary to state one further assumption. The sensory effect represented on each record that has been drawn from the LCM is assumed to fluctuate rapidly over time, momentary values falling into a normal distribution whose mean represents the sensory effect experienced on the trial on which the record was formed.

The response the subject selects is the one most likely to be rewarded on the basis of the evidence contained in the sample of records drawn from the LCM. To find it, the subject gets the sum of the probability densities for each response, at the current input, and selects the response for which the density is greatest. The process is illustrated in FIGURE 1, which shows sensory effect distributions representing four records, three of which show that *R*2 was rewarded and one that *R*1 was rewarded. The process amounts to summing the heights of the *R*2 curves above the point representing the current input, doing the same for the *R*1 curves, and then determining which sum is larger.

If the subject were in the presolution period or in an experiment on probability learning, the records drawn from the LCM would contain no information concerning discriminative stimuli. Given a sample of this sort, the subject simply chooses the response with the greater probability of reward. The rule that accomplishes this is a simple one: if the sample contains more records labeled *R*2 than *R*1, make response *R*2; otherwise make response *R*1.

The model described here has been written as a computer program which was used to simulate results to be expected in three experimental situations dealing with discrimination and generalization of durations. The model has three parameters. One of these, the size of the LCM, is the primary determinant of the rate at which learning proceeds, but is virtually irrelevant to the analysis of asymptotic performance. The other parameters are θ, the size of the sample retrieved on each trial, and W, the

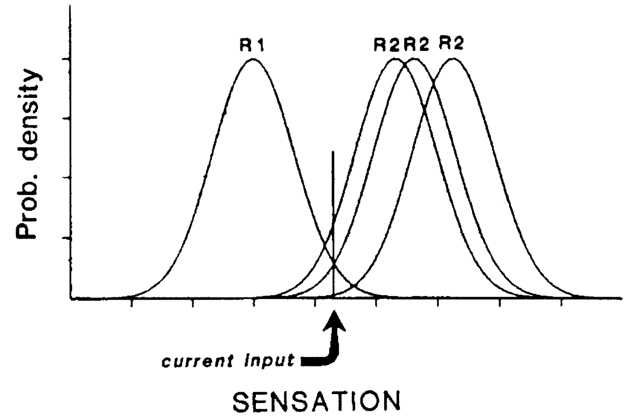

FIGURE 1. A sample of four positive records from the LCM. The choice of response is based on the probability densities at the point labelled "current input."

quantity that represents the Weber fraction in Fechner's law. In the context of the present model W functions as a sensitivity parameter. W was assigned the value of 0.25 for all simulations presented in this paper.

PSYCHOMETRIC FUNCTIONS FOR DURATION

In a well-known set of experiments by Stubbs[5] pigeons were required to classify lights differing in duration. In one experiment the birds were presented on each trial with one of 10 signal durations ranging from 1 to 30 sec. They were rewarded for pecking on one of two keys when presented with any of the durations that exceeded 5 sec, and for pecking on the other key when presented with those durations equal to or less than 5 sec.

The results Stubbs obtained for three pigeons are shown by the dashed lines and data points in FIGURE 2. The solid line represents the results of a computer simulation based on the LCM model. The psychometric function predicted by the model is a close approximation to the empirical functions obtained by Stubbs. Psychometric functions of this sort may be used, of course, to obtain indices of differential sensitivity to duration. Stubbs obtained such functions at several absolute duration ranges and found Weber's law to be true when all time values were doubled, tripled, or quadrupled. The simulation program considers only relative durations so, not surprisingly, it also yields Weber's law.

One interesting feature of the functions shown in FIGURE 2 is that the lower and upper asymptotes lie at roughly 0.05 and 0.95, respectively, instead of at 0 and 1.0, as would be expected by most theories of the psychometric function. This feature is fairly typical of psychometric functions obtained in animal experiments. It suggests that on some trials the animal's choice is not under stimulus control, that is, that the animal's choice of response is independent of the stimuli manipulated by the experimenter. One way to describe this situation is to say that the animal is not "attending" to the stimulus dimension on all trials. Heinemann, *et al.*[6] introduced the following equation which may be used to correct the data for the effects of "inattention":

$$p(R) = p(A)p(R \mid A) + (1 - p(A))p(R \mid \overline{A})$$

where $p(R)$ is the probability of one of the two choice responses, $p(A)$ the probability of attention, $p(R \mid A)$ the probability of the response given attention (given by the model for the psychometric function), and $p(R \mid \overline{A})$ the probability of the response given inattention.

The LCM model provides a deeper explanation of the phenomenon under discussion. That the asymptotes of the psychometric function lie at values other than 0 and 1.0 has to do with the fact that the sample of records the subject retrieves from memory on each trial may fail to include information about all stimulus values. Consider a trial on which a stimulus that lies at the asymptote of the psychometric function is presented, say, a very short duration for which $R1$ is the response that would earn a reward. Now assume that the sample of records retrieved from memory contains only records of trials on which $R2$ was made. In this case the subject must make response $R2$ (which is an error, of course). The smaller the size of the sample, the more frequently such instances of "missing records" will occur, and the more the asymptotes

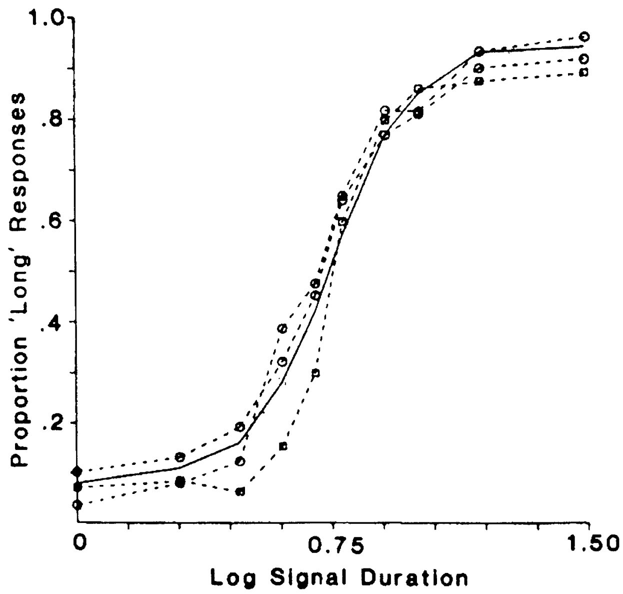

FIGURE 2. Proportion of trials on which the subject made the response designated as correct for the five longest signal durations. The *dashed lines* connecting the circles and squares represent results obtained by Stubbs[5] for three pigeons. The *solid line* represents a computer simulation with $W = 0.25$ and sample size equal to 4.

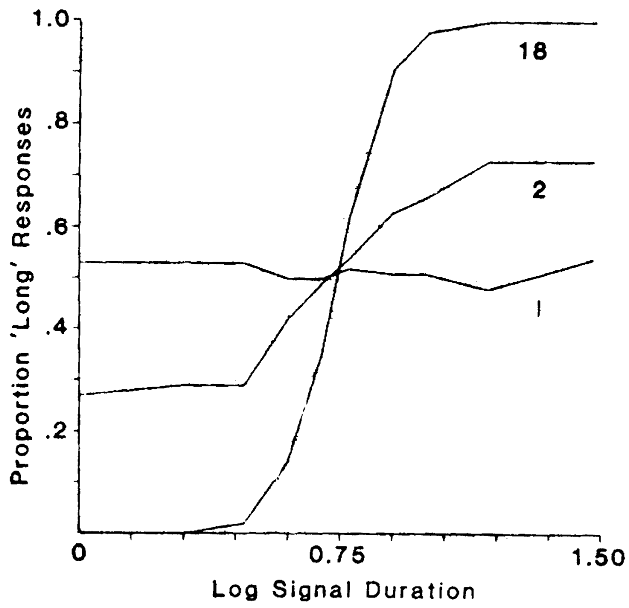

FIGURE 3. Simulated results showing the effects of varying sample size in the situation investigated by Stubbs.[5] $W = 0.25$; the sample size is given by the number written under each curve.

of the psychometric function will depart from values of 0 and 1.0.[b] These effects of varying sample size are illustrated in FIGURE 3, which shows simulated results for sample sizes equal to 1, 2, and 18.

CATEGORIZATION FOLLOWING TWO-STIMULUS TRAINING

Working with rats in a two-choice situation that was formally similar to the one used by Stubbs,[5] Church and Deluty[11] presented only two stimulus durations during training, a short one (for example, 2 sec), after which one response was rewarded, and a long one (for example, 8 sec), after which the alternative response was rewarded. Following training they presented stimuli of intermediate durations in a generalization test. The functions relating proportion of one of the choices to signal durations were very similar to those obtained by Stubbs and shown in FIGURE 2. The stimulus duration corresponding to a response proportion of 0.5, referred to as the "bisection point" corresponded quite closely to the geometric mean of the two stimulus values used in training. This was true for several different sets of training durations, as would be expected if subjective time were logarithmically related to real time. Simulations based on the LCM model also show that the bisection points fall very close to the geometric mean of the training stimuli. For example, in simulations in which the sample size was set equal to 10 and the signal durations presented in training were either 2 and 8 sec or 4 and 16 sec, the bisection points fell at 3.82 sec (geometric mean = 4 sec) and 8.13 sec (geometric mean = 8 sec), respectively.

[b]There are situations in which the asymptotes of the psychometric function fail to lie at 0 and 1.0 for quite different reasons. Among the most important of these, from a theoretical point of view, are situations in which control of behavior is shared by two or more stimulus dimensions. The principles involved have been discussed by Chase and Heinemann,[7] Heinemann and Chase,[8] and Heinemann *et al.*[9,10] Of course, if the appropriate receptors are not stimulated on every trial, then the psychometric function will also have asymptotes other than 0 and 1.0.

CATEGORIZATION OF DURATIONS IN A DISCRETE-TRIAL GO/NO-GO SITUATION

In a recent series of experiments Church and Gibbon[12] studied rats' behavior in a situation first used by Blough[13] and referred to by him as "maintained generalization." In most of the experiments the animals were presented on each trial with a visual signal having one of nine durations. A press on a single lever earned a reward if it occurred when the animal was presented with one of these durations (the $S+$), but had no consequences if it occurred following any of the other eight durations (different values of $S-$). Typically, the $S+$ was the median duration.

This situation differs from all the situations to which the LCM model has heretofore been applied in that it deals with the occurrence or nonoccurrence of a single response rather than a choice between alternative responses that are being recorded by the experimenter. The simulations of this situation assume that the animals are rewarded by the experimenter for not responding to negative stimuli. This is a

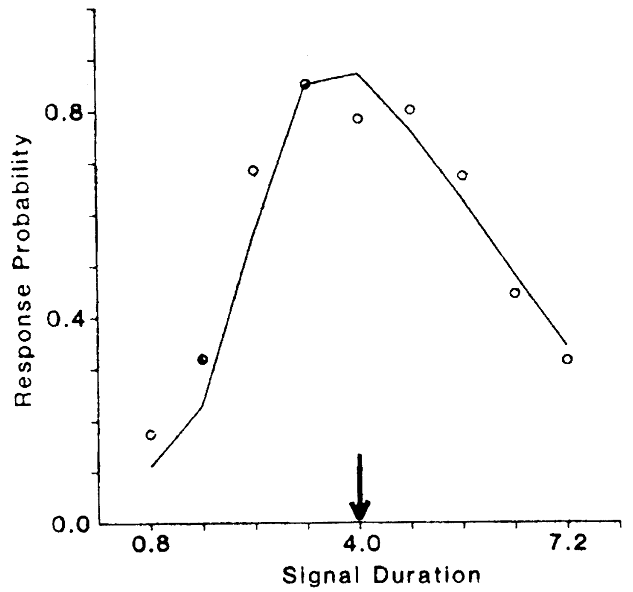

FIGURE 4. Probability of one or more bar-presses as a function of signal duration. The points represent results obtained by Church and Gibbon.[12] The *line* represents a computer simulation with $W = 0.25$ and sample size equal to 10. The *arrow* points to the positive signal duration.

simplifying assumption, probably to be replaced in the future by the assumption that, when the animal does not make the response the experimenter is recording, it is engaged in some other behavior that has some positive value for it. The simulated results that will be shown do not represent "best fits." To obtain them a single parameter, θ, was varied by trial and error until a reasonable approximation to the empirical results was obtained.

FIGURE 4 shows the outcome of one of the basic experiments done by Church and Gibbon[12] (points) together with the results of a simulation (solid line) based on an assumed sample size of 10. (As mentioned, $W = 0.25$ in all simulations shown in this paper.) FIGURE 5 shows the same results plotted on a logarithmically scaled stimulus axis. As Church and Gibbon noted, the generalization curve appears more nearly symmetric on the linear axis than on the logarithmic axis. This may seem surprising in view of the fact that the LCM model assumes Fechner's law. The greater symmetry of the curve on the linear scale reflects the linear stimulus spacing used in this

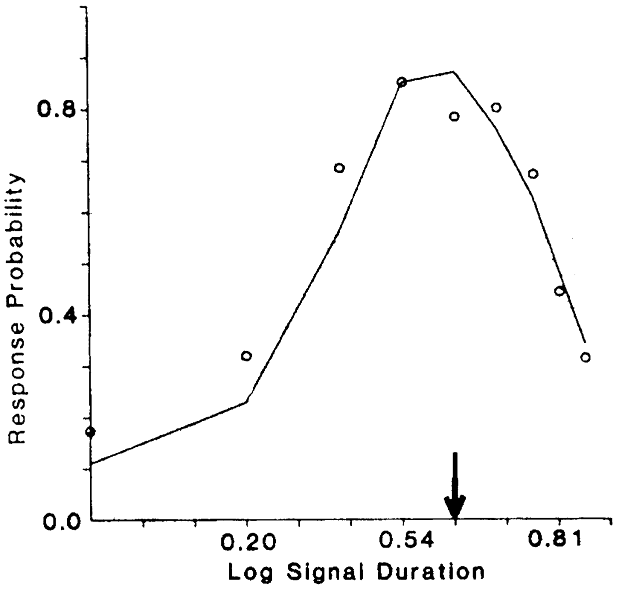

FIGURE 5. Probability of one or more bar-presses as a function of log signal duration. These data are the same as those shown in FIGURE 4.

experiment. If the nine stimuli used in the experiment are spaced at equal logarithmic distances, then the simulated generalization curve is symmetric on a logarithmic stimulus axis.

The last-mentioned theoretical result is in conflict with the empirical results of Church and Gibbon,[12] who found no substantial effect of stimulus spacing on the form of the generalization curve. This may reflect a defect in the LCM model, but it is also possible that the experimental conditions assumed in the simulation do not in fact match the conditions of the actual experiment. For example, the asymptotic levels of performance assumed in the simulation might not have been achieved in the actual experiments. This is of some importance because the same animals were trained first with one spacing and then the other.

These results, particularly as represented in FIGURE 5, suggest that the generalization curves may be leveling off at a response probability well above zero. Noting this, Church and Gibbon decided to check on the matter by extending the range of durations presented during the experiment. Their results with the extended stimulus range are shown in FIGURE 6. The squares represent the results obtained with the "normal" range of stimuli represented in FIGURE 5, the circles the results obtained with the extended range. The solid line represents the results of a simulation based on the LCM model for sample size equal to 5. Both empirical and theoretical curves clearly level off well above a response probability of zero.

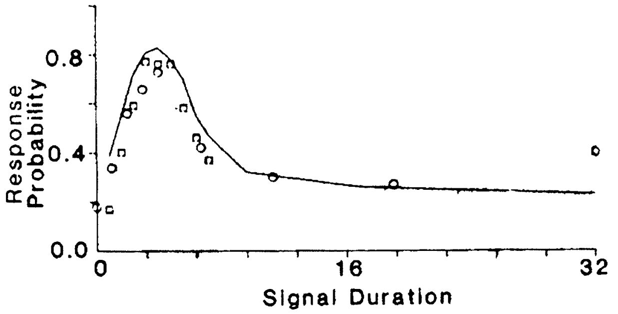

FIGURE 6. Extended stimulus range: The *circles* represent performance of rats on the extended range; the *squares* show performance on the standard range. Data are from Church and Gibbon.[12] The *solid line* represents the result of a simulation. $W = 0.25$ and the size of the sample was 5.

Heinemann and Chase[14] showed how the correction for "inattention" may be applied to go/no-go situations of the sort under discussion here, and this correction is part of the theoretical analysis proposed by Church and Gibbon.[12] According to this interpretation, when the subjects are attending to the relevant stimulus continuum (duration), they base their choice of response upon the experienced value of subjective time in accordance with a specified decision rule; when they are not attending, they base the choice of responding on factors unrelated to the value of the stimuli manipulated by the experimenter. The LCM model elucidates the notion of "inattention" and eliminates the need to estimate from the data two free parameters, the probability of attention and the probability of the response given inattention.

According to the LCM model the cause of the raised asymptote shown in FIGURE 6 is again missing records. Consider a trial on which a very long duration (say 32 sec) is presented, but the sample of records contains information concerning the response made in the presence of $S+$, but no records of responses to any durations longer than $S+$. In this case the subject will make the response appropriate to $S+$.

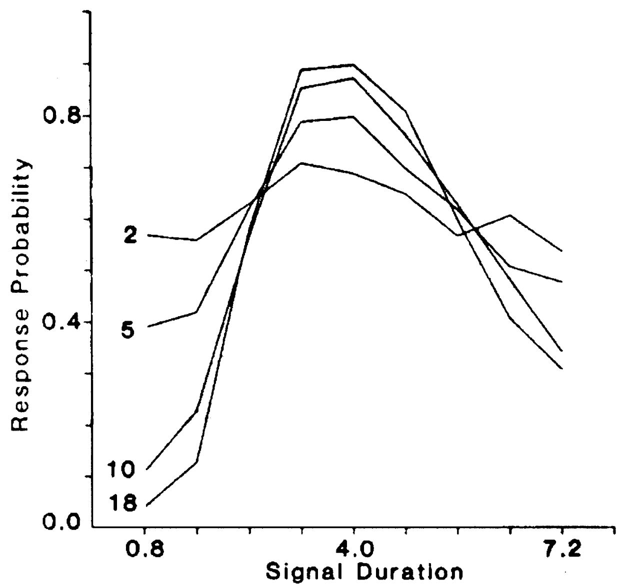

FIGURE 7. Theoretical effects of sample size on the generalization curves. Results of simulations assuming the sample sizes that are written next to each curve.

The probability that the sample will not contain information concerning responding to stimuli longer than S+ increases as the size of the sample decreases. The theoretical effect of sample size on the generalization curves may be seen in the simulated results shown in FIGURE 7. The theoretical curves become flatter and level off at progressively higher values of response probability as the size of the sample decreases. For a sample size of 1, the generalization function becomes a horizontal line at a level that matches the presentation probabilities of the positive and negative stimuli (assumed to be 0.5 in the present simulation).

Variation in the proportion of trials on which the S+ is presented has a pronounced effect on the generalization curve as shown by Church and Gibbon.[12] The results they obtained with two different presentation probabilities are shown by the points and dashed lines in FIGURE 8. The solid lines also shown in FIGURE 8 represent the results of a series of simulations assuming the presentation probabilities written next to each curve. With respect to the form of the curves and the direction of their vertical displacement as a function of presentation probability, there is good agreement

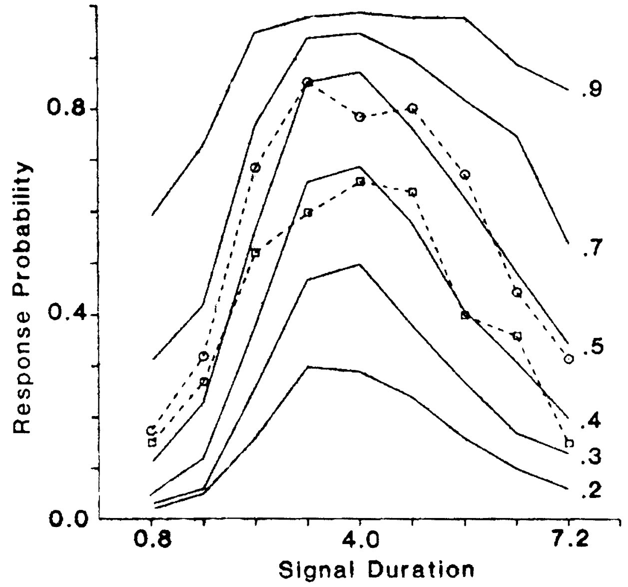

FIGURE 8. The effects of presentation probability on the generalization curves. The *solid lines* are the results of simulations assuming the sample size written next to each curve. The *dashed lines* connect data from Church and Gibbon.[12] *Squares:* presentation probability = .125; *circles:* presentation probability = .5.

between the theoretical and empirical curves. That the numerical values of the presentation proportions represented by the empirical curves and the nearest theoretical curves do not agree should not be surprising in view of the approximations made in the simulations and the fact that the experimental procedures assumed in the simulations do not quite match the actual experimental procedures.

FIGURE 9 shows the effects of varying the probability of reward for correct responses to $S+$. Again, the points represent results obtained by Church and Gibbon[12] for two values of reward probability. The solid lines represent the results of simulations for the reward probabilities written under each curve.

According to the LCM model, the effects of variations in stimulus presentation probabilities and reward probabilities have essentially the same cause. Changing the

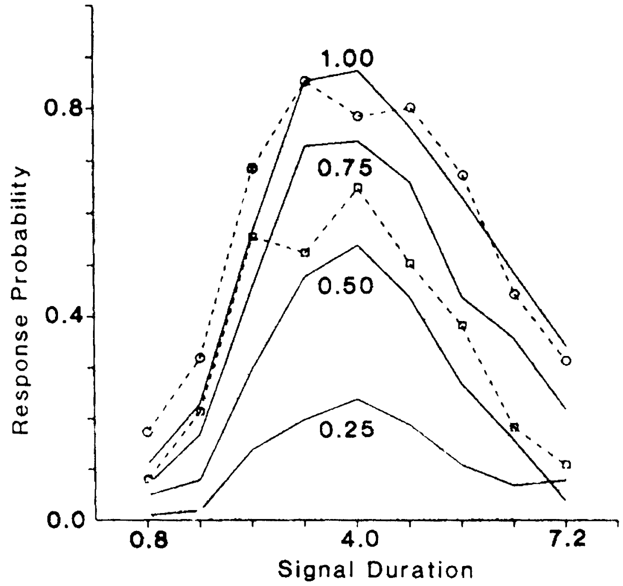

FIGURE 9. The effect of reward probability on the generalization curves. The *solid lines* represent the results of simulations assuming the reward probabilities written above each line. The *dashed lines* connect results obtained for rats by Church and Gibbon.[12] *Squares:* reward probability = .25; *circles:* reward probability = 1.0.

proportion of trials on which the $S+$ is presented is reflected directly in the number of records of successful responding to $S+$ that are contained in the memory. Since the records upon which the choice of response on any trial is based represent a random sample of the records in the LCM, any change in the proportion of one type of record in the LCM is reflected in a corresponding change in the average proportion of these records in the samples. Thus, an increase in the presentation probability of $S+$ results in an increase in the average proportion of $S+$ records in each sample and therefore an increasing probability of responding. Correspondingly, a decrease in the presentation probability of $S+$ will result in a decrease of responding.

Turning to the effect of reward probability: each unrewarded response to $S+$ results in placing a negative (uniformative) record in the memory. As a result the average sample retrieved from memory will contain fewer $S+$ records than it would if all responses to $S+$ were rewarded; and the probability of obtaining samples that contain no records of responding to $S+$ at all is correspondingly larger.

CONCLUDING REMARKS

The LCM model, which describes a large number of phenomena of intensity discrimination and generalization, also gives a reasonable account of some basic phenomena of temporal discrimination and generalization. It appears that information concerning the intensity and duration of stimuli is processed in much the same manner.

One of the virtues of the model that has been discussed is parsimony. As applied to asymptotic performance, the LCM model uses just two parameters, W and θ. It gives an account of the effects of "inattention" and it provides a rational account of the effects of "motivational" variables, such as the presentation probability of the positive and negative signals and the probability of obtaining a reward for responding to $S+$.

ACKNOWLEDGMENT

I wish to thank Sheila Chase for important advice on theoretical matters.

REFERENCES

1. HEINEMANN, E. G. 1984. A memory model for decision processes in pigeons. *In* Quantitative Analyses of Behavior, Vol. 4: Discrimination Processes. M. L. Commons, R. J. Herrnstein & A. R. Wagner, Eds. Ballinger. Cambridge, MA.
2. CHASE, S. 1984. Pigeons and the magical number seven. *In* Quantitative Analyses of Behavior, Vol. 4: Discrimination Processes. M. L. Commons, R. J. Herrnstein & A. R. Wagner, Eds. Ballinger. Cambridge, MA.
3. WALD, A. 1947. Sequential Analysis. Dover. New York, NY.
4. HEINEMANN, E. G. 1984. The presolution period and the detection of statistical associations. *In* Quantitative Analyses of Behavior, Vol. 4: Discrimination Processes. M. L. Commons, R. J. Herrnstein & A. R. Wagner, Eds. Ballinger. Cambridge, MA.
5. STUBBS, A. 1968. The discrimination of stimulus duration by pigeons. J. Exp. Anal. Behav. **11:** 223–238.
6. HEINEMANN, E. G., E. AVIN, M. A. SULLIVAN & S. CHASE. 1969. Analysis of stimulus generalization with a psychophysical method. J. Exp. Psychol. **80:** 215–224.
7. CHASE, S. & E. G. HEINEMANN. 1972. Decisions based on redundant information: An analysis of two-dimensional stimulus control. J. Exp. Psychol. **92:** 161–175.

8. HEINEMANN, E. G. & S. CHASE. 1970. Conditional stimulus control. J. Exp. Psychol. **84:** 187–197.
9. HEINEMANN, E. G., S. CHASE & C. MANDELL. 1968. Discriminative control of "attention." Science **160:** 553–554.
10. HEINEMANN, E. G., S. CHASE & C. MANDELL. 1969. Discrimination control of "attention:" Technical note. Science **164:** 198.
11. CHURCH, R. M. & M. Z. DELUTY. 1977. Bisection of temporal intervals. J. Exp. Psychol. Anim. Behav. Processes **3:** 216–228.
12. CHURCH, R. M. & J. GIBBON. 1983. Temporal generalization. J. Exp. Psychol. Anim. Behav. Processes **8:** 165–186.
13. BLOUGH, D. S. 1975. Steady state data and a quantitative model for operant generalization and discrimination. J. Exp. Psychol. Anim. Behav. Processes **1:** 3–21.
14. HEINEMANN, E. G. & S. CHASE. 1975. Stimulus generalization. *In* Handbook of Learning and Cognitive Processes, Vol. 2; Conditioning and Behavior Theory. W. K. Estes, Ed. Erlbaum. Hillsdale, NJ.