From M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.). 1983. *Quantitative Analyses of Behavior: Discrimination Processes.* Cambridge, MA: Ballinger (Vol. IV, Chapter 2, pages 21-36).

The Presolution Period and the Detection of Statistical Associations

Eric G. Heinemann
Brooklyn College of the City University of New York

When studying the acquisition of discriminations, one often finds an initial period during which there is no measurable change in the dependent variable under study (e.g., percent correct choices). This period is usually called the "presolution period" (PSP) (Krechevsky, 1932). Numerous experiments on the presolution period have been done to decide between discontinuity theory (Lashley, 1929; Krechevsky, 1932, 1938) and continuity theory (Spence, 1936, 1945), most often by studying the effect that the reversal of the discrimination during the presolution period has on the course of acquisition. Reviewers of this work (Goodrich, Ross, and Wagner, 1961; Sutherland and Mackintosh, 1963) agree that discrimination reversal during the presolution period retards acquisition of the reversed discrimination, provided that the discriminative stimuli stimulate the appropriate receptor organs. In spite of the large research effort organized around the discontinuity-continuity controversy, relatively little is known of the variables that govern the length of the presolution period. Theories from which such effects may be predicted in a quantitative way have not been developed. It is the purpose of this chapter to present a quantitative model for the processes occurring during the presolution period, together with some relevant data. The model is intended to apply to a specific type of experimental situation, namely, a discrete-trial experiment in which one of two stimulus values from a continuum, such as sound intensity, is presented on each trial. The subject is rewarded for making one response, R_1, in the presence of one of the stimuli, S_1, and an alternative response, R_2, in the presence of the other stimulus, S_2.[1] In the experiments to be described in this chapter, the subjects were White Carneaux pigeons, the stimuli, S_1 and S_2, were levels of intensity of white noise, and R_1 and R_2 were pecks on two response keys in a standard pigeon training chamber.

The theoretical treatment regards the animal as a detector of the statistical association that has been arranged among the stimulus intensity levels, the stimulation associated with the responses that are made, and the consequences of the responses. As is commonly done in signal-recognition theory, it is assumed that over indefinitely repeated presentations, the two stimulus intensities, S_1 and S_2, produce sensory effects that are normally distributed with different means, μ_1 and μ_2, and a common standard deviation σ. It is further assumed that over a series of trials, the subject remembers the sensory effects that were present when R_1 was followed by reward and those present when R_2 was followed by reward. The subject's task is to decide whether the mean sensory effects, M_1 and M_2, which preceded rewarded R_1 and R_2 responses differ reliably. If they do differ reliably, then the stimuli under consideration predict the outcome of the responses, and it would be profitable to use them as a basis for response choice.

The heart of the present theory is a formal model of the procedure that the subject uses to decide whether or not to consider the difference between M_1 and M_2 significant. More specifically, the subject tests the hypothesis, H_0, that

$\left|\frac{\mu_1 - \mu_2}{\sigma_{\mu D}}\right| = 0$ against the hypothesis, H_1, that $\left|\frac{\mu_1 - \mu_2}{\sigma_{\mu D}}\right| \geq \delta$, where δ is a positive number that is only negligibly different from zero. This model is the sequential probability ratio test developed by Wald (1947). It is of very special interest here because Wald has been able to show that, for testing certain simple statistical hypotheses, the sequential probability ratio test is an optimal test in the sense that no other test of the same strength can lead to rejection or acceptance of the hypothesis on the basis of fewer observations. (The strength of the test is defined by α and β the probabilities of errors of the first and second kind, respectively.)

The basic evidence the subject has to work with consists of the sensory effects experienced when R_1 and R_2 are followed by reward. If these sensory effects were presented to the subjects in pairs, one of each kind, the difference between the members of each pair could be treated as an observation. Indeed, this is how situations of this type are generally treated by Wald. However, in the experimental situations to be considered here, S_1 and S_2 are presented in random order. Furthermore, provision needs to be made for dealing with the case in which S_1 and S_2 are not presented equally often. In the treatment that follows, it is assumed that the subject computes the *average* sensory effects experienced when R_1 and R_2 were followed by reward. The average is taken over n trials where, for now, n is considered fixed.[2] An observation x is defined as the difference between the mean sensory effects associated with successful R_1 and R_2 responses. The distribution of the observations is the sampling distribution of the differences between the two means, and the hypothesis to be tested is that the mean of this sampling distribution is equal to zero.

In a sequential test, the observations used to test the hypothesis are taken one at a time. As each new observation is taken, the observer makes one of three decisions: (1) to end the test by accepting H_0, (2) to end the test by rejecting H_0, or (3) to continue the test by taking another observation. The process is continued until either the first or second decision is made. With respect to the present application, the presolution period is assumed to end when the subject rejects H_0. Further, the criterion for accepting H_0 is assumed to be so strict that for all practical purposes it cannot be attained within the lifetime of the subject.

In somewhat more detail, the decision made at each stage of the process (after the mth observation) is based on P_{1m}, the joint probability of the m observations when H_1 is true divided by P_{0m}, the joint probability of the m observations when H_0 is true. The test is conducted by choosing two positive constants, A and B $(B < A)$, and following the rule:

If	$B < P_{1m} / P_{0m} < A$	Continue the test taking another observation.	
If	$P_{1m} / P_{0m} \geq A$	Reject H_0	(accept H_1).
If	$P_{1m} / P_{0m} \leq \boldsymbol{A}$	Accept H_0	(reject H_1).

Wald shows how to select the constants A and B so that the test will have any desired strength (α, β).

If successive observations are independent, then the joint probabilities P_{1m} and P_{0m} will each be the product of m factors. It is convenient to work with the logarithm of P_{1m}/P_{0m} because this logarithm can be written as the sum of m terms. For this reason, the actual test procedures described by Wald are based on cumulative sums.

The specific form that the test takes when the observations are normally distributed and the statistical hypotheses referred to above are being tested is given by Wald (1947, pp. 134-137) and more sketchily in Appendix A. A graphic representation of the test procedure is given in

Figure 2-1. Z_m the quantity on the vertical axis, is a transform of the cumulative sum of m observations (see Appendix). The test ends (the presolution period ends) when Z_m crosses the decision line U.

Figure 2-1. Schematic representation of the sequential probability ratio test. The value of Z_m as a function of the number of observation, m, when H_0 is not true.

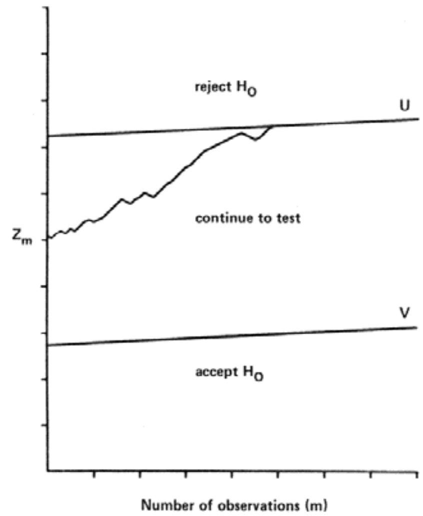

Since the observation x is a random variable, the number of trials needed before the test is ended will also be a random variable. The relation between the average number of observations required to end the test and the true mean of the observations $(\mu D = \mu_1 - \mu_2)$ is called the average sample number function (ASN function) of the test. This ASN function is of primary interest in the present context because, on the assumptions that have been made, it should describe empirically obtained relationships between stimulus discriminability and the number of trials in the PSP. It can be shown that, for fixed stimulus-presentation frequencies and fixed n (number of trials over which the subject averages sensory effects), the mean of the observations, expressed in units of standard deviations, $\mu_D / \sigma_{\mu D}$, is proportional to $d' (S_1, S_2)$. Here $d' (S_1, S_2) = \mu_1 - \mu_2 / \sigma$ where σ is the common standard deviation of the normal distributions of sensory effects induced by repeated presentations of S_1 and S_2.

If the slope of lines U and V in Figure 2-1 were equal to zero, then the average number of observations that must be made before the test is ended would be inversely proportional to $d' (S_1, S_2)$ (Appendix A). Another way to express this relationship is as the reciprocity rule $m \times d' (S_1, S_2) = k$. The product of the number of trials in the presolution period and $d' (S_1, S_2)$ is constant. In fact the decision lines U and V have a positive, non-zero slope (equal to $\delta^2 / 2$; see Appendix A, a), but numerical estimates indicate that, for the application considered in this chapter, the slope of these decision lines is very close to zero (Appendix A, b). As a consequence, the reciprocity rule can be expected to be a close approximation to empirical data.

STIMULUS DISCRIMINABILITY AND THE LENGTH OF THE PSP

The first experiment was designed to examine the relationship between stimulus discriminability and the length of the PSP. Three groups of pigeons were trained to discriminate between two levels of intensity of white noise delivered through a speaker located in the response panel of an LVE pigeon chamber. The intensities were 71 and 72 dB sound pressure level for one group, 70 and 73 dB for the second group, and 69 and 74 dB for the third. The chamber contained three response keys located in a horizontal row at the approximate height of the pigeons' heads. Training sessions consisted of eighty trials separated from each other by intervals of 10 seconds. At the start of a trial the center key was illuminated by white light, and a subsequent peck on the center key caused the white noise to come on at one of the two levels. At the same time, the side keys became illuminated. A single peck on one of the illuminated side keys was considered a correct response if it occurred in the presence of the lesser of the two intensities, and a peck on the other side key was considered correct for the greater intensity. Correct responses were rewarded by a 2.5-second period of access to a tray of mixed grain. Incorrect responses were followed after 10 seconds by a repetition of the trial.

The acquisition curves for one bird from each group are shown in Figure 2-2. At the beginning of training, the two curves shown in each figure are inter. twined, indicating the absence of any discrimination. It is clear that the length of the presolution period depends strongly on the difference between the intensity levels.

To obtain estimates of the length of the presolution periods, the learning curves for each bird were first replotted with proportion of errors as the dependent variable. During the presolution period, these error proportions scatter about 0.5. After the end of the presolution period, the proportion of errors decreases in a fashion quite well described by a negatively accelerated learning curve taken from Estes and Straughan (1954, Eq.6). Accordingly, a two segment curve, a horizontal line at $p = 0.5$ followed by the negatively accelerated learning curve was fitted to the results of each bird. No interpretation of the learning process in terms of Estes and Straughan's theory is intended. Rather, their equation is used as a convenient empirical one. The fitting processes consisted of a computer-aided search for the length of the horizontal line and the parameters of the learning curve that together yielded the least sum of squared deviations of the points from the fitted function. The length of the horizontal line segment found in this fashion is accepted as an estimate of the length of the presolution period.

The data for the last twenty days of training, representing asymptotic performance, were used to estimate a value of d' (S_1, S_2) for each bird. The relation between these values of d' (S_1, S_2) and the lengths of the presolution periods is shown in Figure 2-3. Each point represents the results of a single pigeon, and the fitted curve represents the ASN function of the sequential test. The function derived from Wald's sequential-decision theory appears to be a respectable approximation to the data.

Figure 2-2. Each panel shows the proportion of R_1 responses made in the presence of S_1 (solid line) and the proportion of R_1 responses made in the presence of S_2 (dashed line) for a single pigeon. Proportions are based on 40 trials per day for each stimulus and reflect only trials following correct responses.

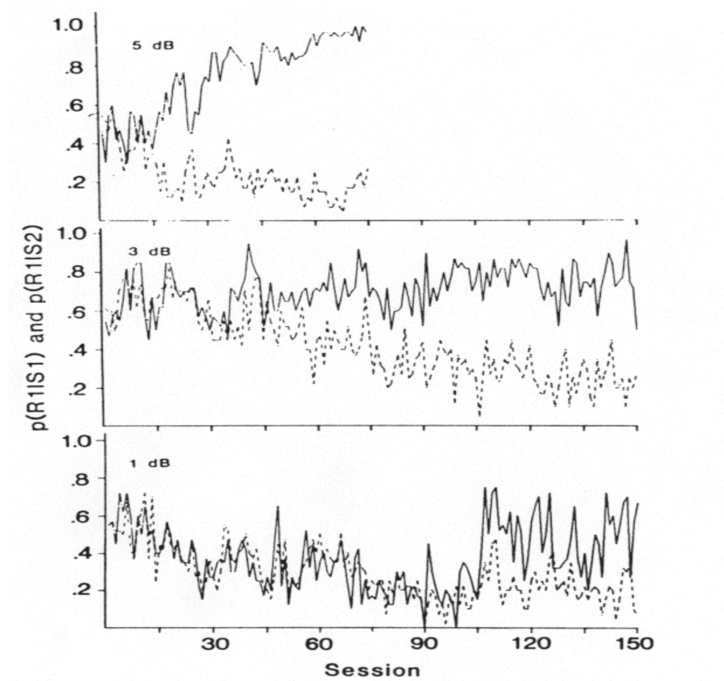

Figure 2-3. Number of trials in the presolution period as a function of d'. Each point represents results for a single pigeon. The smooth curve is the ASN function.

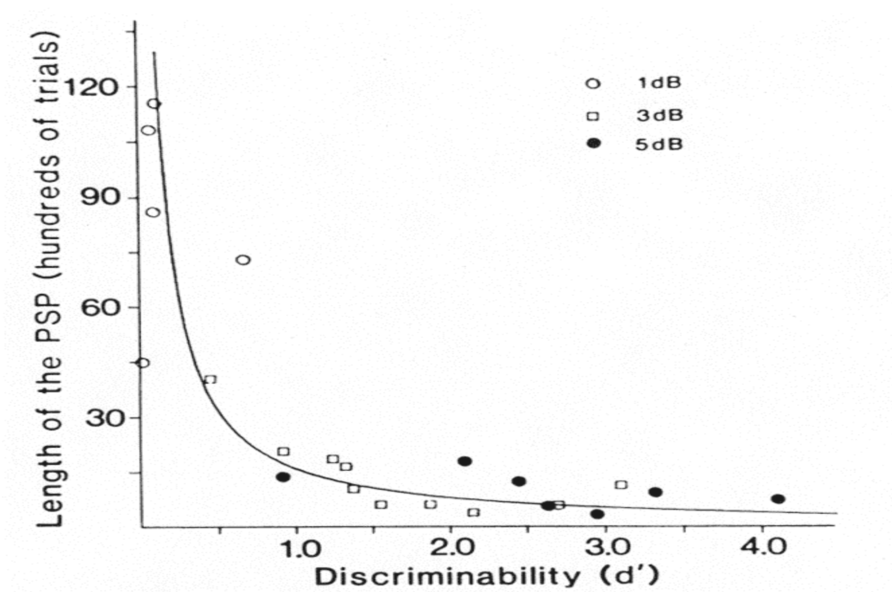

THE EFFECTS OF EASY-TO-HARD AND REVERSAL SHIFTS ON THE PSP

It is possible to derive from sequential-decision theory a large number of predictions concerning the effect that manipulations such as discrimination reversal, probabilistic discrimination training, easy-to-hard or hard-to-easy shifts, etc., will have upon the length of the presolution period. The effects of several such manipulations were examined in a recent Ph.D. dissertation by Mary Knopp (1980).

There were seven groups of eight pigeons in Knopp's experiments. Two were control groups trained to discriminate between levels of white noise differing from each other by 30 and *5* dB, respectively. Of the remaining groups, three began the experiment with 40, 80, or 160 trials of training on the 30-dB discrimination and were then shifted to the 5-dB one. During the 30-dB and 5-dB phases of this procedure, the relation between correct responses and the relative stimulus intensities remained unchanged. For example, if a pigeon was rewarded for pecking on the left key when presented with the more intense stimulus during the 30-dB phase, it was also rewarded for pecking the left key when presented with the more intense of the two stimuli presented during the 5-dB phase. According to the results provided by the 30-dB control group (and those of many other pigeons trained under the 30-dB condition in our laboratory), the presolution period for the 30-dB discrimination exceeds 160 trials, so the shift from easy to hard occurred before the end of the presolution period.

Figure 2-4. Predicted effects of an easy-to-hard shift during the presolution period (see text).

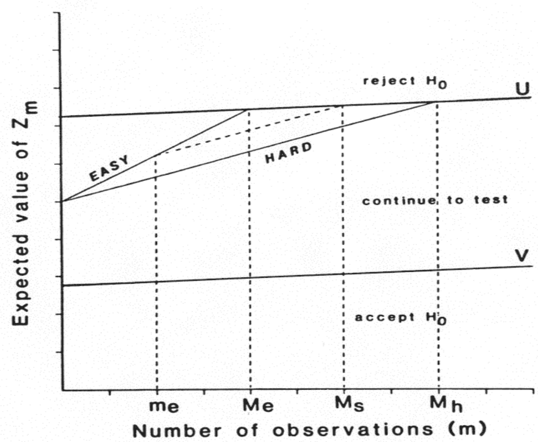

The remaining two groups were also trained on easy and hard discriminations during the PSP, but the easy discrimination was reversed with respect to the hard one. Specifically, one of these groups received 400 trials on the hard (5 -dB) discrimination, followed by 160 trials on the reversed easy (30-dB) discrimination (concurrent reversal and hard-to-easy shifts) followed by a shift back to the original ("re-reversed") *5* -dB discrimination. The second group began its training with 160 trials on the easy discrimination (30 dB) and was then shifted to the reversed *5*-dB discrimination (a concurrent reversal and easy-to-hard shift). Thus, one of these groups received 160 trials of training on the reversed easy discrimination in the middle of the PSP, the other received the same number of trials at the beginning of the PSP. The purpose of this particular manipulation was to test the prediction that the effect of the reversal training is independent of when it occurs, provided it occurs within the PSP.

Figure 2-5. Predicted effects of an easy-to-hard reversal shift during the presolution period (see text).

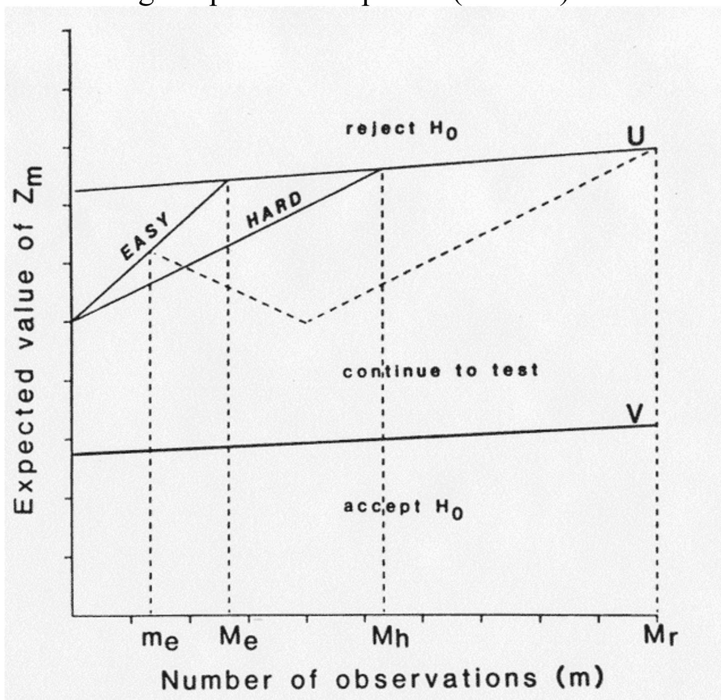

The way the sequential-decision model handles the effect of easy-to-hard shifts and easy-to-hard reversal shifts is illustrated in Figures 2-4 and 2-5. In these figures, the lines U and V represent the decision criteria. The lines labeled "easy" show how the expected value of the decision variable Z_m increases with *m* for an easy discrimination, with the PSP ending at $m = M_e$. The lines labeled "hard" show the same thing for the hard discrimination, with the PSP ending at $m = M_h$. The broken lines show how the expected value of Z_m varies with *m* after the shift from easy to hard or the reversal shift from easy to hard. After the easy-to-hard shift, the PSP ends at $m = M_s$ (Figure *2- 4).* After the concurrent easy-to -hard and reversal shifts it ends at $m = M_r$ (Figure 2-5).

Figure 2-6. Summary of Knopp's results on the effect of easy-to-hard and easy-to-hard reversal shifts on the length of the PSP (see text).

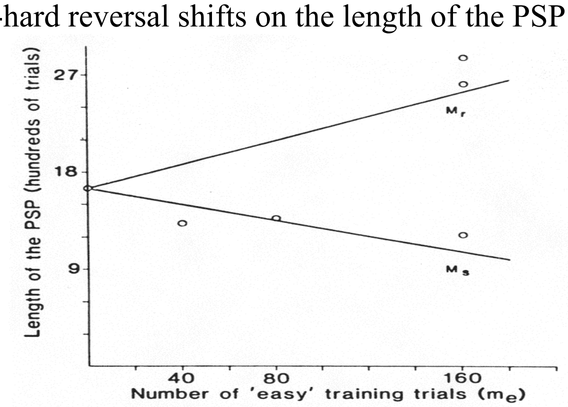

It can be shown (Appendix B) that for the easy-to-hard shift, the value of M_s the total number of trials in the PSP is given by

$$M_2 = (1 - M_h / M_e)\, m_e + M_h \qquad (2.1)$$

For the easy-to-hard reversal shift, the value of M_r is given by

$$M_r = (1 - M_h / M_e)\, m_e + M_h \qquad (2.2)$$

In both equations, m_e represents the number of training trials on the easy discrimination.

Equations (2.1) and (2.2) are the equations for the straight lines labeled M_s and M_r respectively, in Figure 2-6. The points scattered about line M_s represent the average PSPs of Knopps' three easy-to-hard groups, and the two points that lie above line M_r (at $m_e = 160$) are the results for her two reversal groups. In evaluating the fit of the model to the data, it should be noted that the predictive equations (2.1) and (2.2) involve only the parameters M_s and M_r that are estimated from the results of the control groups. The model predicts the results of the experimental groups without any free parameters.

Past research has shown that easy-to-hard shifts facilitate learning of the hard discrimination, whereas reversal shifts retard discrimination learning. Knopp's findings are in reasonable qualitative agreement with the findings of such research, however, detailed comparisons with prior experiments are difficult to make because of differences in procedure and in the dependent variables studied.

Two differences are particularly important. The first is that in Knopp's experiments the shifts from one condition to the other were always made while the subjects were still in the PSP, whereas in most other experiments (for example, Lawrence, 1952; Logan, 1966; Sutherland and Mackintosh, 1963), the subjects were not shifted to the hard task until they showed a reliable discrimination in the easy one (after the end of the PSP). The second difference is that Knopp's primary dependent variable was the trial on which the PSP ended, whereas other investigators have typically used some measure of discrimination accuracy. The essential point is that the portion of Knopp's work considered here deals only with the PSP. Comparison with other research clearly requires consideration of processes that follow the PSP as well.

CONCLUDING REMARKS

Implicit in the treatment of the PSP that has been presented here is the assumption that the PSP is the first stage of a more extended acquisition process. During the PSP stage, the subject comes to "discover" the predictive-stimulus continuum or to "attend to the relevant-stimulus continuum." A theory of the processes that follow the end of the PSP is presented in Chapter 1. It is important to understand that the theory of the PSP does not specify the overt behavior of the subject during the PSP. For example, it does not specify the choice of responses. Response selection is assumed to be governed by processes of the type discussed in Chapter 1.

In order to specify more precisely what happens at the end of the PSP, that is, to specify what is meant by saying that the subject "attends" to the relevant-stimulus continuum, it is necessary to make some minimal reference to the theory presented in Chapter 1. This theory assumes a limited-capacity memory, the contents of which play an important role in the guidance of overt behavior. It is assumed that when the PSP ends (with the rejection of H_0 by the subject), the subject begins to store information concerning the relevant stimuli, behaviors engaged in, and the outcome of these behaviors in the limited-capacity memory just mentioned. Actual behavior is governed entirely by the contents of this memory. In particular, it is not assumed that as a result of the processes occurring during the PSP the subject has learned which response to make in the presence of each stimulus.

There is one more matter to be considered. Subjects can detect statistical associations involving any of a large number of stimulus continua, such as a luminance and wavelength of lights, frequency of pure tones, curvature of lines, etc. It must therefore be assumed that the subject carries out a large number of sequential tests of the sort described above, one for each physical continuum that is providing differential stimulation.

APPENDIX A

(a) The Sequential Probability Ratio Test

The *assumed* sequential-test procedure is that proposed by Wald for "testing that the mean of a normal distribution with known variance is equal to a specified value." (Wald, 1947, pp. 134-137). The specific hypothesis to be tested in the present application is that the mean is equal to zero ($\mu D = \mu_1 - \mu_2 = 0$). Wald's test is based on the assumption that there will be a small positive value 6 such that the acceptance of the hypothesis that *uD* = 0 is regarded as an error of practical importance whenever | μD | *oD* | > δ. For the test under consideration, the probability that the hypothesis $\mu D = 0$ will be rejected does not exceed a small preassigned value a when μD =0, and the probability of accepting the hypothesis $\mu D = 0$ does not exceed a small preassigned value β whenever | μD | *oD* | > δ. The test procedure is as follows:

After each new observation is made, compute

$$Z_m = \ln \cosh\left[\frac{\delta}{\sigma_D} \sum_{a=1}^{m} x_a\right] \qquad (2.3)$$

Continue making new observations as long as

$$\ln B + \frac{\delta^2}{2} m < Z_m < \ln A + \frac{\delta^2}{2} m \, . \qquad (2.4)$$

Reject H_0 if

$$Z_m \geq \ln A + \frac{\delta^2}{2} m \, . \qquad (2.5)$$

Accept H_0 if

$$Z_m \leq \ln A + \frac{\delta^2}{2} m \, . \qquad (2.6)$$

In the above $A = \frac{(1-\beta)}{\alpha}$ and $B = \frac{\beta}{(1-\alpha)}$.

(b) The Slope of the Decision Line

It can be seen from Eq. (2.4) above that the value Z_m must attain to end the test; varies linearly with *m,* the slope of the decision lines being 82/2. In the application made in this chapter, it is possible to estimate the largest value 6 might have from knowledge of the smallest stimulus difference that subjects are known to discriminate. Numerical calculation shows that δ does not exceed 0.4, yielding a slope of 0.08, a value sufficiently close to zero so that in practice predictions differ only slightly from those based on a slope of zero.

(c) The Reciprocity Rule

Given: $d' = \frac{\mu_1 - \mu_2}{\sigma}$ and the observations x are normally distributed with mean $\mu 1$ - $\mu 2$ and variance σ_D^2.

The expected value of Z_m is

$$E(Z_m) = 1n \cosh\left[\delta m \frac{(\mu_2 - \mu_1)}{\sigma_D}\right]. \quad (2.7)$$

Since $\frac{\mu_2 - \mu_1}{\sigma_D} = d' \frac{\sqrt{n}}{2}$.

where n is the number of trials on which each observation is based, Eq. (2.7) may be rewritten as

$$E(Z_m) = 1n \cosh\left[\frac{\delta \sqrt{n}}{2} md'\right]. \quad (2.8)$$

This shows that the expected value of Z_m depends on the product of m and d'. Since, for practical purposes, the test ends when Z_m attains a constant value (Z_m = 1 *n A*), the reciprocity rule will describe the relation between d' and the number of trials in the PSP.

APPENDIX B

Effects of easy-to-hard and reversal shifts.

Wald shows that if $\left[\frac{\delta}{\sigma}\sum_1^m x_a\right]$ is > 3, Z_m is very nearly equal to $\left[\frac{\delta}{\sigma}\sum_1^m x_a\right] - 1n\ 2$. The expected value of this approximation to Z_m is a linear function of *m*. This approximation is used as the decision variable in Figures 2-4 and 2-5. Equations (2.1) and (2.2) are readily derived from the geometrical relations shown in these figures. The derivation of these equations does not require the assumption that the slope of the decision lines be equal to zero.

REFERENCES

Estes, W.K., & J.H. Straughan. (1954). Analysis of a verbal conditioning situation in terms of statistical learning theory. *Journal* of *Experimental Psychology* 47: 225-34.

Goodrich, K.P., Ross, L. E. & Wagner, A.R.(1961). An examination of selected aspects of continuity and noncontinuity positions in discrimination learning. *Psychological Records II: 105 -1 17.*

Knopp, M. (1980). The effects of intradimensional shifts on the length of the presolution period for discrimination learning. Ph.D. dissertation, City University of New York.

Krechevsky, I. (1932). Hypotheses in rats. *Psychological Review 39: 516-32.*

Krechevsky, I. (1938). A study of the continuity of the problem solving process. *Psychological Review 45: 107-33.*

Lashley, K.S. (1929). *Brain Mechanisms and Intelligence: A quantitative study of injuries to the brain.* Chicago: University of Chicago Press. Reprinted by Dover, New York, 1963.

Lawrence, D.H. (1952). The transfer of a discrimination along a continuum. *Journal* of *Comparative and Physiological Psychology 45: 511- 16.*

Logan, F.A. (1966). Transfer of discrimination. *Journal* of *Experimental Psychology 71: 616-18.*

Spence, W.K. (1936). The nature of discrimination learning in animals. *Psychological Review 43: 427-49.*

Spence, W.K. (1945). An experimental test of continuity and noncontinuity theories of discrimination learning. *Journal* of *Experimental Psychology 35: 253-66.*

Sutherland, N.S., & Mackintosh, N. J. (1971). *Mechanisms* of *Animal Discrimination Learning.* London: Academic Press.

Wald, A. (1947). *Sequential Analysis.* New York: Dover Publications, Inc.

ACKNOWLEDGEMENTS

This article is reproduced from the book with some minor changes, e.g. in formatting and in references to other chapters in the book. The research reported here was supported by NIMH Grant No. MH 18246 and PSC-CUNY Grant No. 14002. I thank Mary Knopp who gathered the data for the first experiment and whose Ph.D. dissertation is discussed.

NOTES TO CHAPTER 2

[1] The situation mentioned is the only one for which a detailed treatment has been developed. However, it seems likely that presolution periods occur in virtually all learning situations and that the basic theory presented here will be applicable to many of them. Many investigators have noted that acquisition curves for “simple” learning tasks, including classical conditioning, frequently have an initial flat segment. A notable example is autoshaping in which number of *trials to first response* is perhaps the most common dependent variable.

[2] When the model is applied, it is assumed that the subject makes one test per 80-trial session. Varying the assumed value of n below 80 has only a very weak effect on the predictions of the model. Given the random sequence in which stimulus information is presented to the subject, with possible unequal presentation frequencies, basing the decision process on average sensory effects intuitively seems sensible and is perhaps close to optimal in some reasonable sense. In any case, it should be understood that because it treats the nature of the observations as given, Wald’s procedure is an optimal way of testing hypotheses concerning them.