

PSYCHOPHYSICS OF REMEMBERING

K. GEOFFREY WHITE AND JOHN T. WIXTED

UNIVERSITY OF OTAGO, NEW ZEALAND AND
UNIVERSITY OF CALIFORNIA, SAN DIEGO

We present a new model of remembering in the context of conditional discrimination. For procedures such as delayed matching to sample, the effect of the sample stimuli at the time of remembering is represented by a pair of Thurstonian (normal) distributions of effective stimulus values. The critical assumption of the model is that, based on prior experience, each effective stimulus value is associated with a ratio of reinforcers obtained for previous correct choices of the comparison stimuli. That ratio determines the choice that is made on the basis of the matching law. The standard deviations of the distributions are assumed to increase with increasing retention-interval duration, and the distance between their means is assumed to be a function of other factors that influence overall difficulty of the discrimination. It is a behavioral model in that choice is determined by its reinforcement history. The model predicts that the biasing effects of the reinforcer differential increase with decreasing discriminability and with increasing retention-interval duration. Data from several conditions using a delayed matching-to-sample procedure with pigeons support the predictions.

Key words: remembering, conditional discrimination, forgetting functions, discriminability, reinforcer probability, delayed matching to sample, pigeon

The foundation for a psychophysical analysis of remembering was laid over a century ago, with Fechner's quantitative analysis of sensation and Ebbinghaus' experimental analysis of memory. Fechner (1860) had proposed that sensory experience is a logarithmic function of stimulus intensity. Although the function relating experience and environment has since been interpreted as following a power law (Stevens, 1961) Fechner's fundamental contribution was to render sensory experience amenable to scientific analysis. Fechner had also described a theory for discrimination that predated the more recent development of signal-detection theory (Link, 1994). The field of psychophysics now

offers well-established methods for quantifying the changes in behavior that result from measurable changes in the physical environment.

Ebbinghaus (1885/1964) showed that changes in remembering were amenable to quantification. His measure of retention, the percentage savings in time to relearn a list of nonsense words, decreased logarithmically as retention interval increased over a period of several weeks. It has since been shown that Ebbinghaus' data are better described by a power function (Anderson & Schooler, 1991; Wixted & Ebbesen, 1991). Ebbinghaus' early attempt to quantify memory nevertheless established the possibility that remembering may be amenable to analysis in the same terms as sensing and perceiving.

The aim of the present paper is to offer an analysis of remembering that follows the general approach of signal-detection theory and to apply it to remembering in nonhuman animals. The analysis assumes that the effect of the stimulus can be represented in terms of discriminative processes of the kind suggested by Thurstone (1927). Unlike decision-theoretic approaches, however, a decision criterion is not assumed; instead, the animal's choice is assumed to be influenced directly by the payoff ratio. This approach makes the interesting prediction that the effect of payoffs in biasing an animal's choice depends on

Earlier versions of the model described in this paper were presented to the meetings of the Society for Quantitative Analyses of Behavior, Washington, D.C., June, 1995, and the Behavior Symposium, Christchurch, New Zealand, August, 1996. Preparation of the manuscript was facilitated by the generous hospitality provided by Peter Killeen during the first author's leave at Arizona State University. Killeen's constructive comments considerably improved the manuscript. We are grateful to Angela Ruske for her outstanding assistance during the conduct of the experiments reported here, to members of our laboratory groups, especially Emily Cooney and Deirdra Dougherty, for their helpful discussions, and to Barry Dingwall, for his essential technical expertise.

Requests for reprints should be addressed to the first author at the Department of Psychology, University of Otago, Dunedin, New Zealand (E-mail: kgwhite@otago.ac.nz or jwixted@ucsd.edu).

the discriminability of the stimuli to be discriminated.

Detection Models of Recognition

The treatment of remembering in the same terms as sensory discrimination was not developed until 80 years after Ebbinghaus, when Murdock (1965) suggested that remembering was a matter of discriminating familiar from novel events (see also Banks, 1970; Lockhart & Murdock, 1970; Parks, 1966). An important assumption was that familiar and novel items were assumed to vary in "memory strength" or "familiarity," with the mean of the distribution of novel events at zero and the mean of the distribution of familiar events at a value of memory strength greater than zero. The familiarity of an event increases with training or practice. The increase in familiarity (or item strength) tends to follow a power function of the number of training trials (Anderson, 1995). According to the theory of signal detection (Green & Swets, 1966), the task of deciding whether a particular event is familiar is accomplished by using a decision rule: Events of memory strength greater than a criterion value are categorized as familiar and hence are responded to as remembered, whereas those of weaker strength than the criterion are categorized as novel.

Discriminability and Bias

The signal-detection approach to recognition generated an extensive empirical literature that has benefited from perhaps the most influential theoretical assumption of psychophysics: the independence of discriminability and response bias. The approach yields a measure of discriminability or memorability, measured by the distance d' between the means of the distributions of novel and familiar events. Memorability is independent of the location of the criterion value for deciding between familiar and novel. The location of the criterion value can be used as a measure of response bias, that is, the tendency to report familiar versus novel independently of the discriminability of familiar from novel. A major determinant of the location of the criterion is relative payoff, although specific rules relating criterion value to relative payoff have yet to be developed (Macmillan & Creelman, 1991). If the payoff favors reporting an event as familiar, the criterion is

adjusted so that more events are reported as familiar. In principle, and in the absence of criterion variance, the analysis prescribed by the signal-detection approach leaves the measure of discriminability untainted by errant changes in response tendencies, and reveals pure memorability.

Criterion Location

Egan (1975) and Macmillan and Creelman (1991) have summarized possibilities for choice of a decision rule for models that rely on the assumption of a decision criterion. (a) A probability matching rule was proposed by Parks (1966) and elaborated by Thomas and Legge (1970) and Thomas (1975). The probability matching rule is assumed to apply in cases in which observers have incomplete information about underlying distributions in contrast to the "ideal observer" who has complete information and is thus able to utilize a likelihood ratio criterion. For symmetrical payoff matrices, the probability matching rule assumes that the observer reports occurrence of the target items with a probability that matches their a priori occurrence. Thus, for example, if Stimulus A occurs on 80% of the trials and Stimulus B occurs on 20% of the trials, the criterion will be placed on the decision axis in such a way that the probability of reporting Stimulus A is 80%. Creelman and Donaldson (1968) showed that for judgments of line length, changes in the prior probability of the stimuli did not affect discriminability (i.e., the distance between the distributions) but did affect the placement of the decision criterion such that the highly trained subjects matched response proportions to relative stimulus probability. Because correct responses produced monetary rewards, it is also possible that response proportions were sensitive to the relative monetary payoff. In other studies, response proportions undermatched prior stimulus probability (Dusoir, 1975). (b) An alternative decision rule is one that maximizes expected value. If expected value is to be maximized, the rule takes account of *both* the prior probabilities of occurrence of the stimuli and the payoff matrix. It assumes that the decision rule is based on the likelihood ratio, that is, the ratio of probabilities of occurrence of one sample versus the other (Egan, 1975). Healy and Kubovy (1981) orthogonally varied pay-

off and prior probability in a numerical categorization task. They reported that prior probabilities had larger effects on the likelihood ratio criteria than did payoff, but that a probability matching rule that included a constant based on the payoff matrix could account for their data.

Conditional Discrimination

The separation of discriminability from bias in psychophysics has a parallel in the study of conditional discrimination learning, where the effects of the discriminative stimuli may be separated from the effects of the differential reinforcer probabilities that maintain the discrimination (Nevin, 1981; White, 1986; Williams, 1988). In the context of choice procedures, differential reinforcement biases the choice towards one alternative versus another. Variation in the reinforcer probabilities for two choices results in a power function relation between the ratios of the choices and the reinforcer ratios (Baum, 1974). Variation in the disparity of the discriminative stimuli of a conditional discrimination allows an assessment of the extent to which stimulus disparity limits the discrimination (Nevin, 1969). In conditional discriminations, however, the effects of the reinforcer differential are modulated by stimulus disparity. When choice relies on both stimulus and reinforcer differences, large stimulus differences attenuate the effect of the reinforcer differential and small stimulus differences amplify the reinforcer effect (White, 1986). Signal-detection procedures, which are also classed as conditional discriminations (McCarthy & White, 1987), are associated with a similar problem. The extent of discriminability may modulate the effects of the reinforcers that otherwise bias choice towards one or the other alternative. In situations in which discriminability is expected to vary, the effect of the reinforcer differential may interact with the effects of the discriminative stimuli.

Discriminability Decrement in Remembering

The interaction of reinforcer effects with stimulus effects is especially evident in remembering by virtue of its primary characteristic, the systematic decrement in discriminability with increasing retention-interval duration. Remembering (or forgetting) in-

volves continual changes in discriminability or stimulus control. At very short retention intervals, remembering is easy and the biasing effects of differential payoffs are attenuated, whereas at very long retention intervals, remembering is difficult and the effects of differential payoffs are magnified (Wixted, 1989). The model we describe here addresses this issue and at the same time deals with another more fundamental difficulty associated with the earlier signal-detection approach to remembering, namely the problem of specifying how the subject's decision rule is determined by the location of the criterion on the memory strength continuum.

The model we describe does not rely on the assumption of a criterion. It combines the time-honored assumption that stimulus effects are represented by random values drawn from normal density functions, as in Thurstone's (1927) *discriminal processes*, with the more recent generalization that the preference or choice between alternatives is a function of the payoffs they produce (Baum, 1974; Herrnstein, 1970).

Interaction of Reinforcer and Stimulus Control

Quantification of the effects of payoffs on remembering is easily achieved in terms of the matching law. Jones and White (1992) examined performance in a standard delayed matching-to-sample procedure in which five responses to a red or green sample stimulus initiated a delay interval of variable duration, followed by a choice between red and green comparison stimuli. The probability of reinforcers for correct (matching) choices following red and green sample stimuli was varied over several conditions. When the probability of a reinforcer for a correct choice of red was higher than for a correct choice of green, the tendency to choose red was greater than the tendency to choose green. More generally, the ratio of red to green choices following red and green samples was a power function of the ratio of reinforcers obtained by correct red versus green choices.

The result reported by Jones and White (1992) is summarized in Figure 1, for 1 of their birds, X1. Panel A shows the standard decrement in matching accuracy with increasing delay-interval duration. Matching accuracy was assessed in terms of a measure of

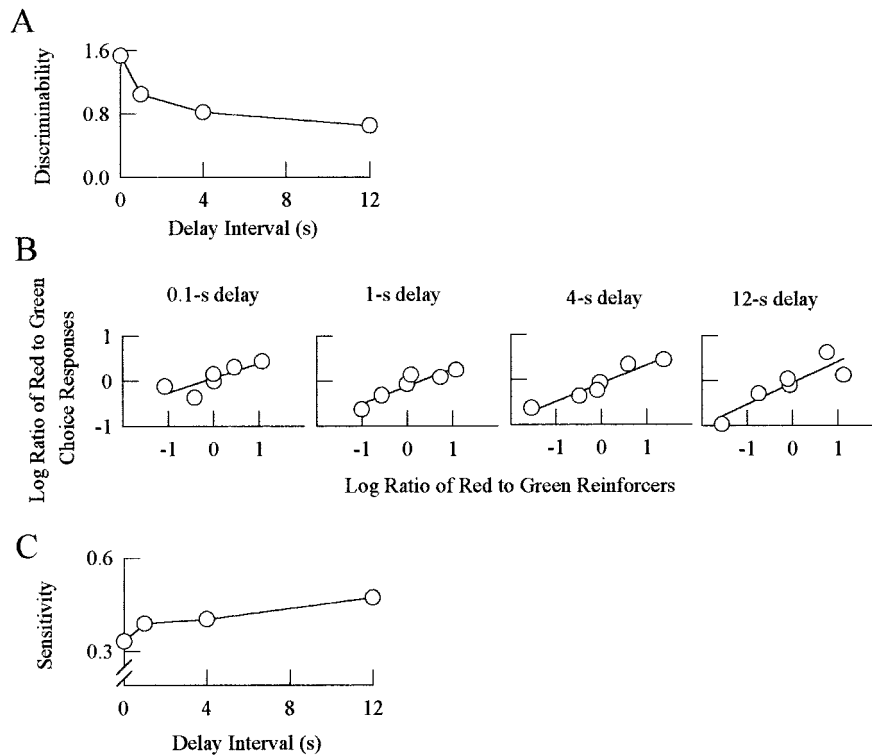


Fig. 1. Data for Bird XI replotted from results reported in Appendix A of Jones and White (1992). In (A), the forgetting function is defined by the decrease in discriminability as a function of delay-interval duration. In (B), the slope of the function relating the log ratio of red to green choices to the log ratio of red to green reinforcers steepens with increasing delay, as summarized in (C).

discriminability (see below). Panel B shows that the tendency to choose red versus green is a function of the ratio of reinforcers obtained by red versus green choices, as expected from the power function version of the matching law (Baum, 1974). The slopes of the functions in Panel B are a measure of the exponent of the power function relating the ratio of choices to the ratio of reinforcers. The exponent is usually interpreted as an index of sensitivity to reinforcement (Davison & McCarthy, 1988). Panel C summarizes these slopes and, more importantly, indicates that sensitivity to reinforcement is a function of delay-interval duration. The slopes of the functions in Panel B gradually increase with increasing retention-interval duration.

The result that the biasing effect of the reinforcer ratio is weak at short delays when discriminability is high and is strong at long delays when discriminability is weak has a parallel in an earlier result for conditional discriminations (White, 1986; White, Pipe, &

McLean, 1985). Here, reinforcer control of the choice was weak for an easy line-tilt discrimination and strong for a difficult discrimination. Nevin, Cate, and Alsop (1993) have also reported a stronger effect of varying reinforcer probability for a smaller disparity between discriminative stimuli in a discrete-trials conditional discrimination. That is, there is a parallel between reinforcer sensitivity effects for perceptual and memorial procedures. If a discrimination is easy because of large stimulus disparity or a short retention interval, there is little effect of varying the reinforcer ratio. If a discrimination is difficult, as with small stimulus disparity or long retention intervals, behavior is very sensitive to variation in the reinforcer ratio. This result is consistent with the model of delayed matching-to-sample performance described by Wixted (1989). According to this model, the influence of sample stimuli and the ratio of reinforcers obtained by correct choices have separate effects that depend on the de-

lay interval. At short delays the conditional discrimination is influential and the reinforcer effect is weak. At long delays when the conditional discrimination is weaker, the reinforcer effect is strong.

Response Measures

The measures of discriminability and bias that we adopt are, respectively, the geometric mean of the ratio of correct to error responses following each sample and the geometric mean of the ratio of red to green choices following each sample. These measures can be derived from different theoretical assumptions, but can otherwise be treated as theory-free measures of performance (as is our preference here). The measures were originally proposed by Luce (1963) in the context of choice theory, and as anticipated by Nevin (1969), Davison and Tustin (1978), and Nevin (1981) in the context of behavioral detection theory. From a theory-free perspective, the measures reflect the likelihood that the subject makes correct responses versus errors, or reports one alternative versus the other. By taking ratios, the measurement scale is not constrained in the way that, for example, a proportion scale is bounded by 1.0. The discriminability measure ($\log d$, or $\log \alpha$) is linearly related to the discriminability measure d' derived from signal detection theory, and satisfies the requirement that both hits and false alarms contribute to the measure of discriminability (Macmillan & Creelman, 1991, pp. 11–13). The discriminability measure is calculated by taking the logarithm (base 10) of the ratio of correct (c) to error (e) responses following each sample (subscripts r and g for red and green samples) and averaging them according to

discriminability,

$$\log d = .5 \cdot \log[(c_r/e_r) \cdot (c_g/e_g)]. \quad (1)$$

The tendency or bias to choose red versus green is calculated by averaging the logarithms of the ratios of choices of red to choices of green following each sample. In Panel B of Figure 1, it is plotted as a function of the red to green reinforcer ratio, according to

$$\begin{aligned} \log \text{red/green choices} \\ &= 0.5 \cdot \log[(c_r/c_g) \cdot (e_r/e_g)] \\ &= a \log(R_r/R_g) + c. \end{aligned} \quad (2)$$

Consistent with the generalized matching law, the log ratio of red to green choices is a linear function of the log ratio of reinforcers obtained by correct red versus green choices (R_r, R_g). The slope of the function, a , measures sensitivity of choice to variation in the reinforcer ratio, and c is a constant describing overall (unexplained) preference for one or the other choice alternative.

The result that the biasing effects of the reinforcer ratio depend on the delay suggests that discriminability and bias may not be independent, as otherwise assumed by the standard signal-detection approach or early versions of the behavioral detection approach (e.g., Davison & Tustin, 1978; but see Alsop & Davison, 1991). Both approaches assume separate and independent influences of the discriminative stimuli and factors that generate bias. In the signal-detection model, the location of a decision criterion along the evidence variable (the determinant of bias) is independent of the distance between signal and noise distributions (the determinant of discriminability). In the behavioral detection model, the biasing effects of stimulus disparity on choice are independent of the biasing effects of the reinforcers obtained by the choices. In a more recent version of the behavioral detection model (Alsop & Davison, 1991; see also Nevin et al., 1993, especially Equation 8), although the two free parameters describing stimulus and reinforcer effects are said to be independent, sensitivity to reinforcement, as measured by Equation 2, reflects the joint effects of stimulus difference and differential reinforcement under some conditions. An interesting issue for the behavioral detection model is that it assumes that the choice between two alternatives following one sample is determined by the ratio of reinforcers obtained by correct choices following both samples. That is, it treats the multiple concurrent schedule as if it were a concurrent schedule. McLean and White (1983) have claimed that in multiple concurrent schedules (including detection procedures), the reinforcers obtained by correct choices following one sample do not influence choice following the other sample (with the exception of reallocation of extraneous reinforcers from one time to another). An alternative assumption was made by Nevin (1981) and is consistent with Luce's (1963) choice theory:

The choice following one sample is influenced by the reinforcers obtained by correct choices following that sample relative to the *generalized* effects of reinforcers obtained by correct choices following the other sample.

A Criterion-Free Model of Remembering

We model the interaction between discriminability and the biasing effects of reinforcers by assuming that the individual chooses between available response options on the basis of which is more likely to be reinforced in a given instance. The model may be characterized as a behavioral theory of remembering, in that choice is directly determined by the reinforcer ratio and no decision criterion is assumed.

Distributions of stimulus effect. The effect of the sample stimulus presumably varies from trial to trial. The stimulus effect at the time the choice is required is assumed to be a random value drawn from a normal density function distributed along a dimension of stimulus effect. This dimension may be interpreted as Thurstone's (1927) psychological continuum (Luce, 1994), and the distributions correspond to Thurstone's discriminational processes. Consecutive occurrences of a stimulus always involve some variation in their effect on behavior (Green & Swets, 1966), and we assume that this is particularly the case when the stimulus is temporally separated from the behavior. At present we are not committed to the view that the variation in stimulus effect is the result of a psychological process; more simply, the stimulus effect is a reflection of the environment (Fetterman, 1996; White, 1991). The distribution of stimulus effect is specified at the time of remembering. The passage of time weakens the discrimination and contributes to the trial-by-trial variability in stimulus effect. Thus the discriminational processes represent the potential stimulus effect *at a given delay interval*. This issue is discussed further in the General Discussion section below.

This assumption is illustrated in Figure 2 (Panel A). The distance D between the means of the discriminational processes is directly related to stimulus disparity. Overall higher levels of discriminability are reflected in larger values of D . The variances of the distributions are usually assumed to be equal, but unequal variances may result from specific manipulations of the attributes of the to-be-remembered

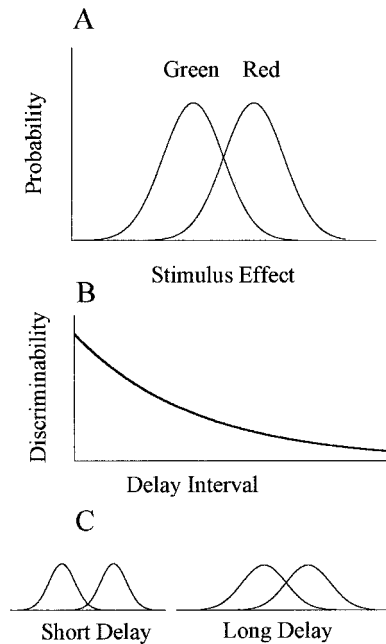


Fig. 2. (A) assumed distributions of stimulus effect for red and green sample stimuli (discriminational processes). (B) decrement in discriminability with increasing delay interval that characterizes the forgetting function. (C) greater variance (and hence greater overlap) of the discriminational processes is illustrated for a long delay interval compared to a short delay.

events. In studies of human recognition memory, for example, receiver-operating characteristics that relate hit rates to false alarm rates are typically asymmetrical. That is, there are different variances for the normal distributions for new and old items on the familiarity or memory strength dimension (Ratcliff, McKoon, & Tindall, 1994; Yonelinas, 1994).

The assumption depicted in Figure 2 (Panel A) follows the standard detection approach (cf. White & Cooney, 1996; Wixted, 1993), in that the discriminational processes have equal variance and are separated by a distance that is related to stimulus disparity. A specific assumption we make, however, is that the variance of the distributions increases monotonically with increasing delay-interval duration. This assumption is illustrated in Figure 2 (Panel C). Two pairs of discriminational processes are shown, both with the same distance D between their means. The pair with less overlap illustrates the case at a short retention interval. The pair with the greater overlap, owing

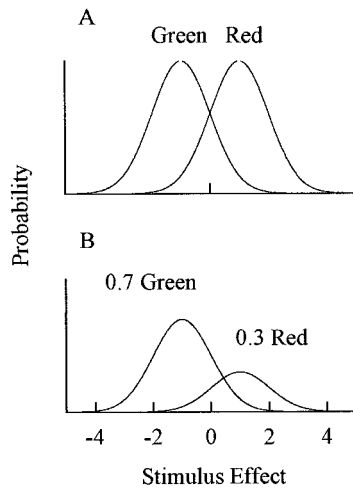


Fig. 3. Probability distributions of stimulus effect (Panel A) multiplied by reinforcer probabilities of .7 (green) or .3 (red) (Panel B).

to larger standard deviations, illustrates the case at a longer retention interval. Thus the reduction in discriminability with increasing retention-interval duration (Figure 2, Panel B) is the result of a continuous increase in the standard deviations of the discriminational processes as the retention interval lengthens. At this stage it is not necessary to specify the function relating variance to delay interval, because the main aim of the model is to predict an inverse relation between discriminability and sensitivity to reinforcement. That is, the model predicts that reinforcer sensitivity increases as discriminability decreases. The reduction in discriminability can be achieved by increasing the variances of the distributions or by decreasing the distance between their means (or both).

Distributions of reinforcer probability. In experimental procedures, reinforcing consequences follow accurate remembering with a defined probability. Each distribution on the continuum of stimulus effect is associated with a given reinforcer probability. For example, assume that the left distribution in Figure 3 (Panel A) describes the variation in stimulus effect at a given delay interval for the green sample, and assume that the right distribution describes the variation in stimulus effect at the same delay for the red sample. (Note that the abscissa is scaled in z-score units, although the variance of the distributions may differ from 1.0.) Also assume that

the reinforcer probability for correctly reporting green is .7 and is .3 for correctly reporting red. The distributions of stimulus effect in Panel A may be translated into distributions of reinforcer probability shown in Figure 3 (Panel B) simply by multiplying the values for the green distribution by .7 and the values of the red distribution by .3. The resulting distributions are distributions of reinforcer probabilities along the dimension of stimulus effect. (Because the probabilities are not required to sum to 1.0, the distributions are not probability density functions.) Each of the possible values of stimulus effect x is therefore associated with a pair of reinforcer probabilities, one for the green distribution, $p(G_x)$, and one for the red distribution, $p(R_x)$.

Choice determinants. We assume that at a given instant or on a particular trial, the effect of a prior event is represented by a stimulus value x randomly determined by either of the discriminational processes in Panel A of Figure 3. The individual's task is to choose green or red on each trial. Assuming that the individual has prior experience in the procedure, for a particular value of x choices of red have been reinforced with probability $p(R_x)$, and choices of green have been reinforced with probability $p(G_x)$. We suppose that on each trial (i.e., for each x) the relative tendency to choose red versus green is directly determined by the ratio of probabilities of reinforcers, $p(R_x)/p(G_x)$. This assumption follows the matching law, according to which choice proportions match reinforcer proportions (Herrnstein, 1970). Our application of the matching assumption differs from the usual use of the matching law, however, in that we predict the tendency to choose one of two alternatives on a single trial on the basis of relative reinforcers accumulated over a large number of prior trials.

It should be noted that unlike various versions of detection theory, our model does not incorporate the notion of a *decision criterion*, and, indeed, could be said to be free of decision rules in that the response tendency is determined directly by relative reinforcer frequency. In signal-detection models, the x axis can be interpreted as a ratio of likelihoods that either signal may have been presented (Macmillan & Creelman, 1991). A decision criterion is established at a specific likelihood ratio, which remains constant over all trials

(within the bounds of criterion variance). For values of x greater than the decision criterion, one response alternative is chosen, and for values of x smaller than the criterion, the other alternative is chosen. In the present model, choice is not related to the likelihood ratio because it depends on the current ratio of reinforcer probabilities that change with x . The probability of choosing one versus the other alternative varies over the x axis and is not all or none, as in the detection models. Whereas the adoption of a constant decision criterion in the signal-detection models results in the derivation of independent measures of discriminability and bias, the assumption that choice is determined by the reinforcer ratio for a given value of x in the present model allows the prediction that bias and discriminability interact.

Separate *measures* of discriminability and bias can be derived from the model, which simulates events in a delayed matching-to-sample procedure on a trial-by-trial basis. The model generates frequencies of correct and error responses separately for trials with red and green prior events as a function of the relative reinforcer frequencies that vary over trials as a result of varying the values of stimulus effect.

Predictions from the Model

Predictions from the model were generated by running many simulations, each for 5,000 trials of a delayed matching-to-sample procedure. The discriminational processes are defined by normal (or, in practice, logistic approximations to normal) distributions with variances S_r^2 and S_g^2 and means of $-0.5D$ (green) and $+0.5D$ (red). The parameter D is disparity on the stimulus effect dimension. To model an experimental procedure, initially a red or green sample stimulus is chosen with a particular probability (usually .5). For the chosen sample, say red, a value x is randomly selected from the associated distribution. In the simulation, x is specified in z -score units, so that $x = S_r * z_i + D/2$, or if the green distribution is selected, $x = S_g * z_i - D/2$, where z_i is a randomly selected z score, that is, a value randomly selected from a normal distribution with a mean of 0 and a standard deviation of 1.

To summarize thus far, on each trial a red or green sample is chosen with a signal pre-

sentation probability of .5, and a value x is randomly sampled from the stimulus effect dimension according to the probability density function defined for each sample stimulus. The reinforcer distributions are defined by multiplying the discriminational processes by the reinforcer probabilities associated with the distributions. For a given value of x , the proportion of the heights of the resulting distributions (at x) directly predicts the probability of choosing red versus green on that trial. The heights (or probabilities) are given by $p(R_x) = p_r * N(D/2, S_r)$ and $p(G_x) = p_g * N(D/2, S_g)$. These distributions represent the probabilities that a reinforcer is arranged on the red and green choice alternatives for a given value of x . For example, assume that p_r (the probability of reinforcement for a correct red response) and p_g (the probability of reinforcement for a correct green response) are both 1.0 (cf. Figure 3, Panel A). Further assume that x on a given red trial happened to equal 0 (i.e., x on this trial happened to fall at the intersection of the two distributions). In the past for that value of x , a reinforcer was as likely to be set up on red as on green, so an unbiased bird would be equally likely to choose either alternative. On the next red trial, imagine that x equaled $D/2$ (i.e., x fell at the mean x value for red trials). Note that, at that point, the height of the red distribution exceeds that of the green distribution by about 6 to 1. That is, in the past, this value of x has occurred on red trials about six times as often as green trials. Because p_r and p_g are both 1.0, this means that, for this particular value of x , a reinforcer is six times as likely to be arranged on the red choice alternative. Thus, according to the matching law, an unbiased bird would be about six times as likely to choose red over green given this value of x .

The process works the same way when p_r and p_g are unequal. For example, assume that $p_r = .3$ and $p_g = .7$. Multiplying both distributions by these values yields the distributions shown in Figure 3 (Panel B). Assume once again that on a given red trial the value of x happened to equal 0. Although that value is as likely to occur on a red trial as on a green trial (as before), it is no longer the case that a reinforcer is as likely to be arranged on the red choice alternative as on the green choice alternative. Instead, as illustrated in

Figure 3, a reinforcer is now about twice as likely to be arranged on the green alternative as on the red alternative. This can be most easily appreciated by noting that the height of the green distribution at 0 is about twice the height of the red distribution (Figure 3, Panel B). According to the matching law, an unbiased bird will, under these conditions, be twice as likely to choose green as red. Note that the indifference point (i.e., the value of x that yields indifference between red and green) has shifted to the right. A bird that was previously indifferent between red and green at an x of 0 (when p_r and p_g both equaled 1.0) and that is now indifferent between red and green when x is about +0.5 (the indifference point when p_r and p_g equal .3 and .7, respectively) has exhibited some sensitivity to reinforcement. As described in more detail later, the degree to which that indifference point shifts with changes in the scheduled probabilities of reinforcement is directly related to D and the standard deviations of the distributions. In general, the greater the overlap between the distributions by virtue of smaller D or larger standard deviations, the more a given change in p_r and p_g changes the indifference point. This is another way of saying that the model predicts that sensitivity to reinforcement will be inversely related to discriminability. The experiments described later provide a test of that prediction.

For the first 100 trials of the simulation, the reinforcer frequencies associated with that value of x are given by the programmed probabilities. However, for later trials, the obtained reinforcer proportions (in the simulation) were used to compute reinforcer probabilities. That is, for a particular value of x , the choice of red or green on that trial was predicted by the ratio of obtained reinforcers for that value of x on previous trials. If, for example, x equals 1, and the number of obtained reinforcers for choosing red and green in the past for an x of 1 equaled 20 and 40, respectively, then on the current trial the probability of choosing red would be $20 / (20 + 40)$, or .33. That is, the simulated birds were assumed to match.

Once the choice response, red or green, is selected, reinforcer occurrence is determined by whether the sample stimulus on that trial was red or green and whether a re-




Fig. 4. Values of discriminability predicted by the model as a function of increasing standard deviation with disparity D set at 1 (top panel) and as a function of increasing disparity with standard deviations set at 1 (bottom panel).

inforcer had been set up on a probabilistic basis. Accordingly, the result of a particular trial would include the sample selected, the choice response made, and whether a reinforcer was obtained, just as in a standard experimental procedure. The values of discriminability, $\log d$, obtained from the simulations for different values of the standard deviations with $D = 1$ (top panel) and for different values of the disparity parameter, D , with $S_r = S_g = 1$ (bottom panel) are shown in Figure 4. Whereas the change in predicted discriminability with increasing standard deviation tends to mimic a forgetting function, the increase in discriminability with increasing D is virtually linear.

Figure 5 shows the results of two sets of simulations, conducted as described above, and just as if an actual delayed matching-to-sample procedure was being conducted. For one set, the disparity parameter was set at $D = 4$, as might be the case for an overall easier discrimination. For the other set, D was set at 1, as might be the case for a harder discrimination. The standard deviation values were $S_r = S_g = 1$. For both sets, 12 pairs of reinforcer probabilities for correct red and green choices were varied over values ranging from .9, .1 to .1, .9. From the 5,000 trials of each simulation, frequencies of reinforcers for correct red and green choices were used to cal-

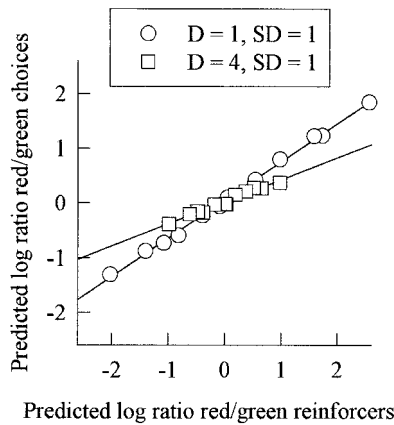


Fig. 5. Predictions of the model for low ($D = 1$) and high ($D = 4$) levels of disparity and constant standard deviations, for log ratios of red to green choices as a function of log ratios of red to green reinforcers obtained by the correct choices.

culate log ratios of red to green reinforcers. Corresponding frequencies of red and green choices on red-sample and green-sample trials were used to calculate the mean of the log ratio of red to green choices following the red sample and the log ratio of red to green choices following the green sample. This measure is the same as the bias measure proposed by Luce (1963) and Davison and Tustin (1978), and is described by Equation 2 above. The log ratio of choices from the simulations was satisfactorily described by a linear function of the log ratio of reinforcers, consistent with the generalized matching law (Baum, 1974). That is, the present model predicts a power function relation between choice and reinforcer ratios.

Three aspects of the predictions in Figure 5 are noteworthy. First, the present model predicts undermatching in the conditional discrimination, on the basis of the assumption that the simulated bird matches choice proportions to reinforcer proportions at each value of x on the stimulus effect continuum. Second, with decreasing distance between the discriminative processes (decreasing D), the reinforcer ratios are predicted to cover a wider range (despite entering the same values of programmed reinforcer probabilities for each D). Third, the function for the easy discrimination is flatter than the function for the hard discrimination. That is, sensitivity of the choice ratio to changes in the reinforcer




Fig. 6. The relation between sensitivity to reinforcement predicted by the model and discriminability predicted by the model, for instances when changes in both were produced by varying the standard deviation with $D = 1$ (top panel) and varying the disparity parameter D , with $SD = 1$ (bottom panel).

ratio, as given by the slope of the straight line, is smaller for the easy discrimination.

Predicting the Interaction Between Discriminability and Bias

Figure 6 shows predictions from the model for sensitivity to reinforcement and discriminability, one plotted as a function of the other. These results were obtained in the same way as for Figure 5 but with two pairs of reinforcer probabilities (p_r and p_g) for each instance of standard deviation and D ($p_r/p_g = .2/.8$ or $.8/.2$). The top panel shows the inverse relation between reinforcer sensitivity and discriminability when standard deviations were varied over 14 values with $D = 1$. The bottom panel shows the inverse relation obtained by varying D over 17 values with the standard deviations of the discriminative processes set at 1. Each 5,000-trial simulation generated a matrix of obtained response and reinforcer frequencies, which were used to calculate log ratios of red to green reinforcers and log ratios of red to green choice responses. Reinforcer sensitivity was estimated from the slope of the function relating ratios of

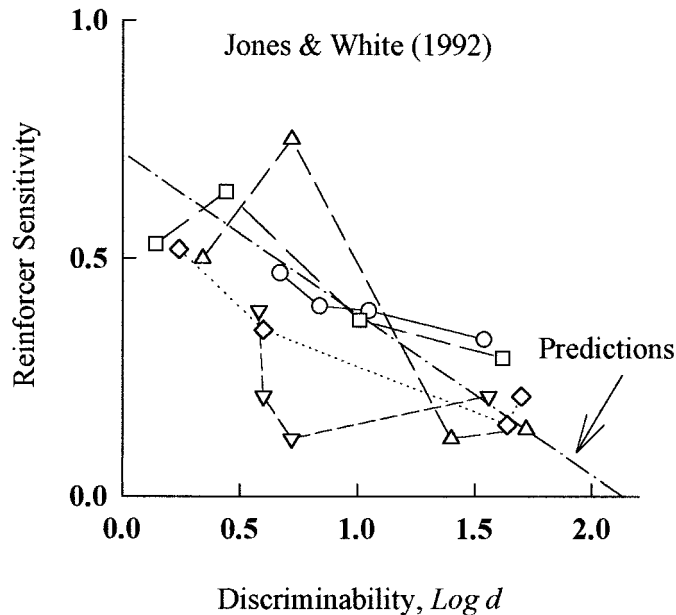


Fig. 7. The relation between sensitivity to reinforcement and discriminability at different delay intervals for individual birds (shown by different symbols) in the study by Jones and White (1992) and the regression line for the predictions from the present model (from Figure 6, top panel).

choices to ratios of reinforcers (cf. Figure 5). Discriminability was calculated from the log ratios of correct to error responses (Davison & Tustin, 1978; Luce, 1963; McMillan & Creelman, 1991). Figure 6 shows the inverse and virtually linear relation between reinforcer sensitivity and discriminability predicted by the model. The results in the top panel obtained when the standard deviations were increased, illustrate the pattern expected when retention-interval duration is increased. This is the interaction evident in the data reported by Jones and White (1992). The results in the bottom panel are what might be expected when other factors that influence overall task difficulty are manipulated.

Figure 7 shows the data for individual birds from Jones and White (1992) along with the regression line that best fits the predictions from our simulation for standard deviation varied with the disparity parameter arbitrarily set at 1 (shown in the top panel of Figure 6). The predictions are generally consistent with the data.

Sources of Discriminability

An important assumption of the model is that the distance D between the means of the

discriminal processes is influenced by several sources of discriminability. These include the wavelength disparity (for example) of the sample stimuli. Retention-interval duration also influences predicted discriminability, but by increasing the variances of the distributions. Other factors that influence discriminability are the fixed-ratio requirement for sample-key responding and sample-presentation duration (White, 1985). We now report a set of experimental conditions for pigeons working in a delayed matching-to-sample procedure in order to examine further the inverse relation between discriminability and the biasing effects of the reinforcer differential. In all cases, retention-interval duration was varied in order to generate further evidence for the inverse relation predicted when standard deviations of the discriminational processes are increased. In other conditions, factors influencing discrimination difficulty were included in order to assess the relation predicted when the distance between the means of the discriminational processes changes. Reinforcement rate was determined by the independent probability that a correct choice was followed by a reinforcer. Correct choices that were not followed by reinforcers had no other conse-

quence, consistent with the typical method for arranging delayed matching-to-sample procedures and successive discriminations. In the latter, for example, when responses in one component are reinforced at variable intervals, nonreinforced responses have no other consequence that distinguishes them from the nonreinforced responses in the other component in which extinction may be arranged. In some very few studies in which all correct choices in delayed matching to sample are followed by hopper illumination and only some also produce grain (McCarthy & Davison, 1991), hopper illumination could serve as a conditional reinforcer, and its probability could be confounded with the probability of the food reinforcer.

EXPERIMENT 1: ABSOLUTE RATE OF REINFORCEMENT

In the signaled magnitude effect described by Nevin and Grosch (1990), accuracy is overall higher on trials in which larger reinforcers follow correct choices than on trials in which small reinforcers follow correct choices. The two types of trials are differentially signaled. The result was confirmed by McCarthy and Voss (1995) and Jones, White, and Alsop (1995). Our generalization of the result is that accuracy may be overall higher when higher absolute rates of reinforcement are arranged in delayed matching to sample.

We compared the effects of overall rich versus overall lean reinforcement rate in a delayed matching-to-sample procedure. We asked whether we could replicate the effect reported by Jones and White (1992) and predicted by our model, in which sensitivity to reinforcement increased with increasing delay-interval duration, and whether the inverse relation between reinforcer sensitivity and discriminability was generalizable to the different performance levels generated by the manipulation of absolute reinforcer rate.

METHOD

Subjects

Three adult homing pigeons with prior experience in delayed matching-to-sample procedures were maintained within 12 g of 80% of their free-feeding body weights by supplementary feeding with mixed grain following

experimental sessions. Water and grit were freely available in the home cages. The colony room was naturally illuminated with a photoperiod of approximately 14:10 hr.

Apparatus

A light- and sound-attenuating experimental chamber, 33 cm wide, 33 cm deep, and 34 cm high, was painted matte black. A ventilation fan provided masking noise. Three translucent response keys, each 2.9 cm in diameter, were mounted on one wall 9.2 cm center to center and 25.6 cm above the grid floor. Each key could be lit red or green and could be operated by a minimum force of 0.15 N. A central hopper opening below the center key and 5 cm from the grid floor allowed 3-s access to wheat. The only illumination in the chamber was provided by red or green on center or side keys or by a white hopper light during 3-s wheat presentations.

Procedure

Sessions were conducted 7 days per week. Each session consisted of 80 trials, separated by 15-s intertrial intervals (ITIs) during which the chamber was dark and responses were ineffective. On each trial, a fixed ratio (FR) of five responses to a red or green sample stimulus presented on the center key terminated the sample and initiated a dark delay period. The delay lasted for 0.2 s, 1 s, 4 s, or 12 s, with delay durations mixed within sessions. Presentation of red and green comparison stimuli on side keys followed the delay. A single choice response to one of the side keys darkened both keys. Correct (matching) choices produced 3-s access to grain with a given probability. Unreinforced matching responses and incorrect choices produced 3-s blackout periods, followed by the dark ITI. The order of red and green sample stimuli, whether the red and green comparison stimuli appeared on the left or the right, and the order of delay intervals were random within each session, with the constraint that each combination of sample, comparison-stimulus location, and delay occurred equally often and with no more than three consecutive trials with the same sample stimulus.

Reinforcer probabilities were fixed at certain values for each of 20 sessions per condition. In five conditions, the probabilities were overall rich, and in five conditions they

Table 1
Reinforcer probabilities for correct choices of red and green comparison stimuli. 20 sessions were conducted for each condition.

Condition	Order	Red	Green
Rich FR 5 15-s ITI			
1	1	.3	.9
2	2	.9	.3
3	3	.6	.6
4	9	.96	.24
5	10	.24	.96
Lean FR 5 15-s ITI			
6	4	.1	.3
7	5	.3	.1
8	6	.2	.2
9	7	.075	.3
10	8	.3	.075
FR 1 15-s ITI			
11	11	.3	.075
12	12	.075	.3
13	13	.3	.1
14	15	.1	.3
15	14	.2	.2
FR 1 1-s ITI			
16	16	.3	.075
17	17	.075	.3
18	18	.3	.1
19	19	.1	.3

were overall lean. The probabilities are shown in Table 1, which indicates that in some conditions, correct red choices were reinforced with higher probabilities than were correct green choices, and vice versa for other conditions. Table 1 shows the reinforcer probabilities for all experiments reported in the present paper; Conditions 1 to 10 contributed to Experiment 1.

RESULTS

Data analyses were based on frequencies of responses and reinforcers obtained at each delay summed over the last eight sessions per condition. Eight sessions allowed a maximum response frequency of 80 correct responses in the cells of the 2 × 2 matrix for red and green responses following red and green sample stimuli at each delay, for each reinforcer-ratio condition. Figure 8 shows that when responses were pooled over the five different reinforcer conditions, discriminability ($\log d$, calculated according to Equation 1) decreased with increasing delay, and was overall

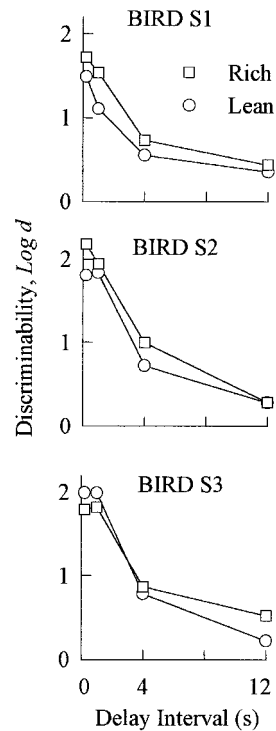


Fig. 8. Discriminability as a function of delay-interval duration for rich and lean reinforcement-rate conditions in Experiment 1.

higher when the overall reinforcer rate was rich.

The functions in Figure 8 are standard forgetting functions, with a clear reduction in discriminability with increasing delay duration (White, 1985). For the pooled data for each bird at each delay, the frequencies of correct red and correct green responses were each higher in the rich reinforcement conditions than in lean, and the frequencies of red and green error responses were each lower in rich than in lean conditions (except for Bird S3 at the two shortest delays). Accordingly, for each bird, the forgetting function for the lean reinforcement conditions was consistently lower than that for the rich conditions, although the difference was small. A similarly small improvement in overall accuracy as a result of increasing overall reinforcement probability has been reported by Blough (1998) for a matching-to-sample procedure with no delays and categories of hues as samples.

The effect of variation in the reinforcer

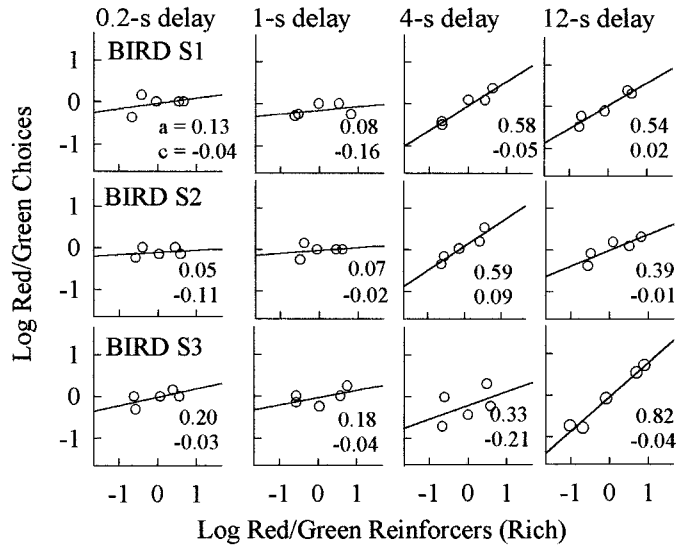


Fig. 9. Log ratio of red to green choices as a function of log ratio of red to green reinforcers for the rich reinforcer-rate conditions of Experiment 1, as a function of delay interval. Values of the parameter estimates for the slope (reinforcer sensitivity) and intercept are shown.

probability ratio is shown in Figure 9 (for the rich reinforcement conditions) and Figure 10 (for the lean reinforcement conditions). Here, the log ratio of red to green choices provides a measure of the tendency to choose red versus green, and was calculated according to Equation 2. For both rich and lean re-

inforcer probabilities, the same effect as was reported by Jones and White (1992) is clear in Figures 9 and 10. At the short delays, the slopes of the matching lines are relatively flat, and at the long delays they are steep.

The slopes, which provide an index of sensitivity to reinforcement, are compared in Fig-




Fig. 10. Log ratio of red to green choices as a function of log ratio of red to green reinforcers for the lean reinforcer-rate conditions of Experiment 1, as a function of delay interval. Values of the parameter estimates for the slope (reinforcer sensitivity) and intercept are shown.

ure 11 for rich and lean reinforcement conditions. With the exception of the lean condition for Bird S1, as the delay lengthens, reinforcer sensitivity increases. This result is consistent with an inverse relation between reinforcer sensitivity and discriminability. There was no consistent difference in reinforcer sensitivity between rich and lean conditions, as indicated by the overlapping standard error bars in Figure 11. In the context of the inverse relation between reinforcer sensitivity and discriminability that is being explored here, this result is perhaps not surprising in view of the small difference in discriminability between the rich and lean conditions. In addition, although the general outcome of the model is an inverse relation between reinforcer discriminability and sensitivity, the model predicts that variation in the absolute reinforcer rate should have no effect on sensitivity because the same proportions of obtained reinforcers are maintained for rich and lean conditions.

EXPERIMENT 2:
SAMPLE RATIO
REQUIREMENT

In Experiment 2 we examined the interaction between reinforcer sensitivity and the change in discriminability caused by decreasing the sample-key ratio requirement. A standard result is that discriminability is overall higher at all delay intervals when more responses are required to the sample stimulus (Roberts, 1972; White, 1985). The same 3 birds and the same procedure as for the lean conditions in Experiment 1 were used. An additional five reinforcer probability conditions were conducted, but with a ratio requirement of FR 1 for sample-key responding. Compared to the function for the FR 5 sample requirement in Experiment 1, an overall lower level of discriminability was expected for the FR 1 conditions, together with an overall higher level of reinforcer sensitivity.

METHOD

The 3 birds of Experiment 1 continued under the same procedural conditions as in the lean reinforcer probability conditions of Experiment 1, but with just a single response to the red or green sample stimulus being required to initiate the delay interval. Table 1




Fig. 11. Sensitivity to reinforcement (values of the slopes in Figures 9 and 10) as a function of delay-interval duration for rich and lean reinforcer-rate conditions of Experiment 1.

shows the order in which Conditions 11 to 15 were conducted. As for the earlier conditions, each was conducted for 20 sessions. Analyses were based on response and reinforcer frequencies summed for the last eight sessions of each condition.

RESULTS

Figure 12 shows the forgetting functions for the FR 1 and FR 5 conditions, for data pooled over the five reinforcer conditions in each set for each bird. The result shows the typical overall reduction in discriminability that results from decreasing the sample-key ratio requirement (Roberts, 1972; White, 1985). Figure 13 shows that the matching law functions for the FR 1 conditions behaved in very much the same way as for the FR 5 lean conditions in Experiment 1 (Figure 10). The slopes of the functions tended to increase with increasing delay-interval duration.

Figure 14 summarizes the change in reinforcer sensitivity over delay-interval duration for the FR 1 and FR 5 conditions. With the exception of the 0.1-s and 4-s delays for Bird S1 and the 12-s delay for Bird S2, reinforcer




Fig. 12. Discriminability as a function of delay-interval duration for conditions with sample ratio requirements of FR 5 (from Experiment 1) and FR 1 (Experiment 2).

sensitivity at the different delays is lower for the FR 5 conditions than for the lower discriminability FR 1 conditions. The tendency for reinforcer sensitivity to increase with increasing delay duration, which was reported by Jones and White (1992) and was observed in the two sets of conditions in Experiment 1, is just as apparent for the FR 1 conditions of Experiment 2 for each of the 3 birds

EXPERIMENT 3: INTERTRIAL INTERVAL DURATION

An inconsistency between our result and model predictions and some data reported by McCarthy and Davison (1991) and McCarthy and Voss (1995) has left us with an interesting puzzle. They reported that reinforcer sensitivity *decreased* with increasing delay-interval duration. McCarthy and Davison varied reinforcer probability over three values in a delayed matching procedure in which the samples were different brightnesses. For delays of 0 s, 1 s, 3 s, and 25 s, respectively, reinforcer sensitivities averaged .75, .52, .56, and .37 (with standard errors of about .09). This result is not consistent with those reported by Jones and White (1992) and in the present Experiments 1 and 2, and is inconsistent with the expectation that when control

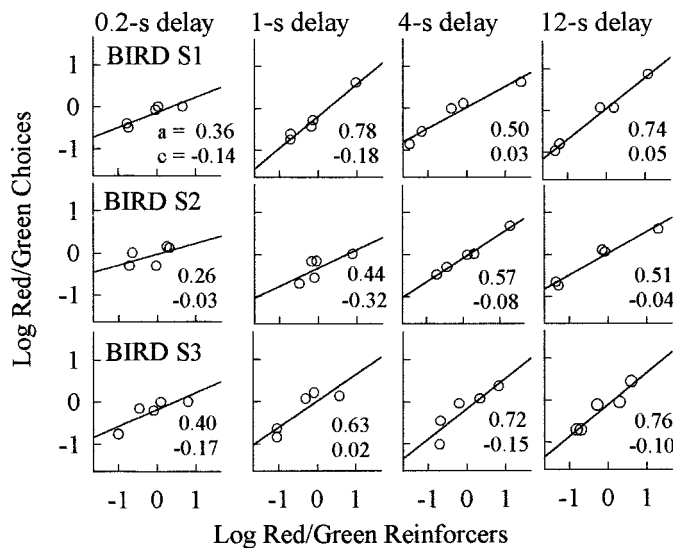


Fig. 13. Log ratio of red to green choices as a function of log ratio of red to green reinforcers for the FR 1 ratio-requirement condition of Experiment 1, as a function of delay interval. Values of the parameter estimates for the slope (reinforcer sensitivity) and intercept are shown.

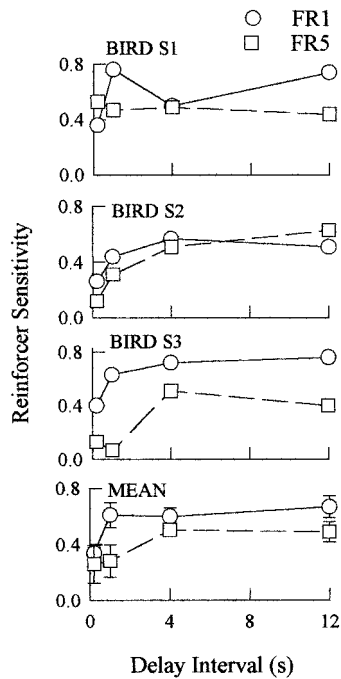


Fig. 14. Sensitivity to reinforcement (values of the slopes in Figures 10 and 13) as a function of delay-interval duration for FR 1 and FR 5 sample ratio-requirement conditions in Experiments 1 and 2.

by the conditional discrimination is minimal, choice should be governed by the reinforcer ratio (cf. Wixted, 1989).

In the study by McCarthy and Davison (1991), each condition was conducted with just one delay interval in each session (whereas delays were mixed within sessions in the present study). In addition, they employed a reinforcement procedure in which a reinforcer for a correct choice on one alternative could not be obtained until the reinforcer that had been set up for a correct choice on the other alternative had been obtained in a previous trial. This procedure resulted in a reduction in total obtained reinforcers with increasing delay, where total reinforcers obtained at the 25-s delay were about half those at the three shorter delays. Furthermore, Jones and White (1992) have shown that the reinforcement scheduling procedure, perhaps along with the use of just one delay per session, results in a large response bias for choice between left and right keys.

The procedure used by McCarthy and Voss (1995) was more similar to that used by Jones

and White (1992), except for (a) use of the procedure to control the ratio of obtained reinforcers, (b) presentation of the hopper light following every correct response, and (c) the use of a short ITI, among other differences. Their general result was consistent with that reported by McCarthy and Davison (1991), namely a decrease in reinforcer sensitivity with increasing delay. In some instances for long delays when there was zero discriminability, there was also zero reinforcer sensitivity. That is, variation in the reinforcer ratio for correct choices had no effect on the choice, a puzzling result. In the experiments by McCarthy and her colleagues, there were generally low levels of discriminability, compared to those reported by Jones and White (1992). In the experiments by McCarthy and her colleagues, short intertrial intervals were used, which are known to lower discriminability (Edhouse & White, 1988; Roberts, 1972; White, 1985). Therefore, in Experiment 3 we compared the effect of a 1-s ITI to the effect of a 15-s ITI using the same procedure as for Experiment 2.

METHOD

The subjects, apparatus, and procedure were the same as those in Experiment 2 with the FR 1 sample ratio requirement except that the intertrial interval was reduced to 1 s. Four ratios of reinforcer probabilities were used. The reinforcer probabilities and the order of conditions are given in Table 1. Twenty sessions were conducted for each condition. Analyses were based on response and reinforcer frequencies summed over the last eight sessions for each condition.

RESULTS

Figure 15 shows that the usual ITI effect on the forgetting functions was obtained. The functions with the 1-s ITI exhibit overall low discriminability, compared to the functions for the 15-s ITI from Conditions 11 through 15 of Experiment 2. The matching law functions in Figure 16 are systematic, but compared to those in Experiments 1 and 2, exhibit steep slopes at all delays.

Figure 17 compares the change in reinforcer sensitivity with increasing delay-interval duration for the 1-s ITI conditions to the function for the corresponding 15-s ITI conditions. Whereas the function for the 15-s ITI




Fig. 15. Discriminability as a function of delay-interval duration for ITIs of 15 s (Experiment 2 data) and 1 s (Experiment 3 data).

increases, the function for the 1-s ITI decreases. We therefore suppose that short ITIs may be responsible for the high reinforcer sensitivity at short delays seen here and in the data of McCarthy and Voss (1995) and McCarthy and Davison (1991). Note, however, that at longer delays reinforcer sensitivity remains high, whereas in the data reported by McCarthy and Voss and McCarthy and Davison it drops to near-zero levels. That is, even with a short ITI we were unable to produce the result seen in the studies by McCarthy and Voss and McCarthy and Davison.

The overall high levels of reinforcer sensitivity in Experiment 3 are predicted by the present model in that discriminability was overall very low. But in order to predict the high sensitivity at the shortest delays, an additional assumption is needed. One possibility is that with short delays, short ITIs, and very low discriminability levels, there may be a stronger proactive effect of the reinforced response from the preceding trial. This situation was modeled by running simulations as above in order to generate a set of reinforcer sensitivities that corresponded to the similarly small discriminability values obtained for the 1-s ITI condition. The additional assumption was incorporated by amplifying the tendency to make a choice on trials preceded by a reinforcer, in the direction of the previously re-

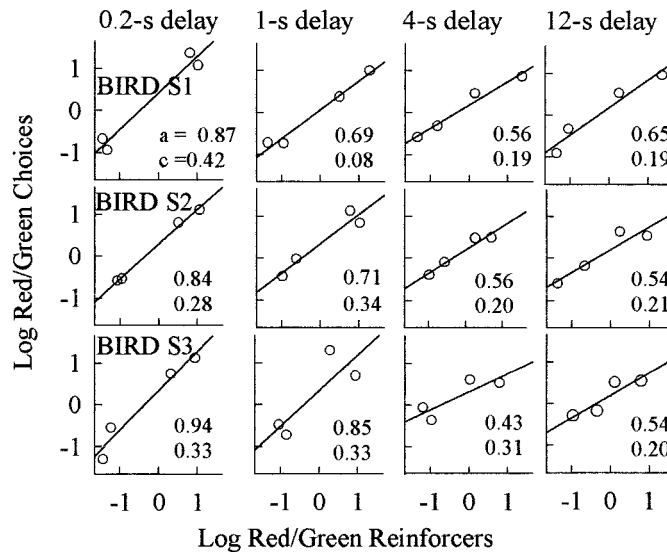


Fig. 16. Log ratio of red to green choices as a function of the log ratio of red to green reinforcers obtained for correct choices as a function of delay, for the 1-s ITI condition of Experiment 3.

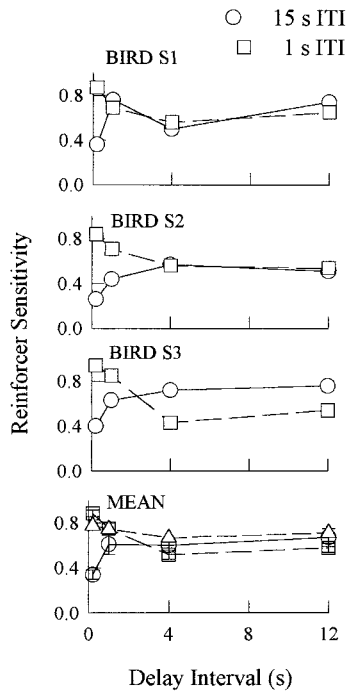


Fig. 17. Sensitivity to reinforcement as a function of delay interval for conditions with 15-s ITIs (Experiment 2) and 1-s ITIs (Experiment 3). Triangles in the panel for the mean data are the predictions from the model.

inforced response, by an arbitrarily chosen factor of five. (By comparison, the level of left-right bias at long delays in the study by McCarthy & Davison, 1991, ranged from 1.26 to 8.91, and in the study by McCarthy & Voss, 1995, they ranged from 1.15 to 5.99.) The resulting predictions are shown as triangles along with the mean data in Figure 17. The predicted values are satisfactorily close to the obtained values for reinforcer sensitivity for the 1-s ITI condition. Our model therefore copes well with the present data.

SUMMARY OF RESULTS

The main approach that we took to evaluating the present model was to vary the level of discriminability in delayed matching to sample by varying delay-interval duration and other factors, and to ask whether there was an inverse relation between reinforcer sensitivity and discriminability as predicted by the model. In order to summarize the relation between reinforcer sensitivity and discriminability for Experiments 1 and 2, in which there was clear variation in discriminability,




Fig. 18. The relation between reinforcer sensitivity and discriminability in Experiments 1 and 2. For each bird, individual data points are shown for the four delay intervals for each of four sets of conditions combining FR 1 or FR 5 with rich or lean reinforcer rates and 1-s or 15-s ITIs. The straight line is the prediction of the model taken from Figure 6 (top panel) and also drawn in Figure 7.

Figure 18 shows values of reinforcer sensitivity for each delay and each of the three sets of conditions, plotted against the corresponding values of discriminability (12 points per bird). The predicted straight-line function included in each panel is exactly the same function as in Figure 7, in which the data from the study by Jones and White (1992) were plotted. That is, the same predicted function is shown with the data for the different birds and the three different procedures in the present study. Using the straight line in Figure 18 to predict values for reinforcer sensitivity, the slopes of the regression (through the origin) relating obtained to predicted val-

ues were 1.03, 0.90, and 1.07 for Birds S1, S2, and S3, respectively. The results of Experiment 3 were not included in Figure 18 because the lack of variation in discriminability did not contribute to assessment of the inverse relation predicted by the model. Nevertheless, the overall high levels of reinforcer sensitivity in Experiment 3 (Figure 17) are consistent with the present model in that they are expected at low discriminability levels. When the data from Experiment 3 are plotted in Figure 18, they cluster at the lowest levels of discriminability, and around the line predicted by the model. Because the same predicted function as in Figure 7, generated by varying standard deviations with D arbitrarily fixed at 1, is drawn for the different birds in Figure 18, the correspondence of the data to the predicted function provides convincing evidence that our model successfully predicts the interaction between reinforcer sensitivity and discriminability.

GENERAL DISCUSSION

To summarize the main point of the model we propose here, the individual's trial-by-trial tendency to choose one or the other comparison stimulus in a delayed matching procedure is determined by the relative probability of previously obtained reinforcers associated with the value x of the stimulus effect on the current trial. The value x is the stimulus effect at the time of remembering. The variation in stimulus effect is defined in terms of a pair of hypothetical discriminative processes or distributions. With increasing retention-interval duration, the variances of the distributions (discriminative dispersions) increase. Discriminability is predicted to decrease as a result of the increased overlap between the distributions. Discriminability may also decrease as a result of other factors that lead to a reduction in the separation of the means of the distributions.

The model can be characterized in terms of a conditional discrimination in which there is a continuum of values of the discriminative stimuli. That is, although only two stimuli are actually used (red and green), their effects vary from trial to trial, thereby creating a distribution of effective stimulus values for each stimulus. Each effective stimulus value has associated with it a unique re-

inforcement history. Thus, instead of two physical stimuli, the model takes account of a whole range of effective stimuli that happen to be normally distributed. But the model is otherwise just the familiar matching law. The model predicts that the effect of the reinforcer ratio in biasing choices is amplified by decreasing the level of discriminability that occurs with longer retention intervals. For example, if the distributions overlap extensively, variations in the arranged reinforcer ratio result in large changes in obtained reinforcer ratios and hence choice ratios across the range of stimulus effect values. But with virtually nonoverlapping distributions, obtained reinforcer ratios and hence choice ratios uniformly favor red for x values under the red distribution and uniformly favor green under the green distribution, relatively independently of variation in the arranged reinforcer ratios. In the experiments reported here, there was clear evidence for this prediction, and also for the more general prediction that factors that decrease overall levels of discriminability increase sensitivity to reinforcement.

Our behavioral theory of remembering bears obvious similarities to signal-detection theory, but it differs in one critical respect. Detection theory assumes that the individual arrives at a decision by setting a criterion somewhere along the stimulus effect continuum. If x on a given trial exceeds that criterion, the pigeon chooses red; otherwise, green is selected. According to this account, every trial involves a decision on the part of the bird (namely, does x exceed the criterion or not?). The decision criterion plays no role in the present nonmediational theory of remembering. Instead, on each trial, the bird's behavior is governed by the history of reinforcement for choosing red or green under the prevailing conditions. If, given a value of x , a response to green has been reinforced more often than a response to red, the bird will be more likely to choose green than red according to the matching law.

The version of the model described above has two parameters, one for the distance between the discriminative processes and another for their variance. Further possibilities include different variances for the two distributions, as may occur in delayed matching for asymmetrical sample stimuli (Wixted &

Dougherty, 1996). Bias and undermatching in the effect of the reinforcer ratio on the trial-by-trial choice probability may also be included. These factors have not been incorporated here because they involve adding free parameters to an otherwise simple model. Quantitative fits of the model to data would be enhanced, however, by inclusion of the additional parameters.

The present model bears some apparent similarity to detection models in that both assume distributions of stimulus effect along a psychological continuum, following Thurstone's notion of discriminial processes (Luce, 1994). The question of how to interpret this continuum is of importance to the issue of how the present model differs from other models of recognition memory.

Signal-detection theory was applied to recognition memory by Parks (1966) as suggested by Murdock (1965) and later by Banks (1970) and Lockhart and Murdock (1970). Snodgrass and Corwin (1988) give a clear summary of these models and discuss the problem of measuring response bias in recognition memory. In the model described by Parks, old and new items vary in their degree of "familiarity" or, in other models, "memory strength." Familiarity or memory strength varies from item to item according to a normal distribution. The mean of the probability distribution for old items is greater in familiarity or memory strength than that for new items, if old is discriminable from new. The dimension along which item strength varies is a psychological continuum (Luce, 1994), in that it reflects variation in how stimuli may be represented in the nervous system or stored in long-term memory. As applied to the matching-to-sample paradigm or similar two-alternative choice procedures, however, there is little sense in defining a zero point for memory strength or familiarity, unless the stimulus effect continuum is interpreted as extending from "greenness" to "redness."

In detection models, the dimension is portrayed as a decision axis (Green & Swets, 1966). A decision based on a criterion value on the dimension provides the main basis for the recognition response. But in the present model, there is some difficulty in identifying the continuum as a decision axis, because no criteria are assumed. One possibility is to define the stimulus effect dimension in terms

of a composite of the factors that influence the overall level of performance. That is, the disparity parameter D may be related to wavelength disparity (for a hue discrimination) in addition to factors such as the FR sample requirement or ITI duration. A related issue is whether the stimulus effect dimension should be construed in sensory or memorial terms, an issue related to our notion that stimulus effect is defined at the time of remembering. Although a memorial interpretation is an obvious possibility, the notion that the effect of the temporally distant sample stimulus may be direct is consistent with a sensory interpretation (White, 1991). That is, if all of the variation in stimulus effect can be attributed to factors in the environment, a "direct remembering" approach can be sustained. It may be more plausible, however, to adopt the view that organismic processes contribute to stimulus variation. In either case, the stimulus effect dimension seems to be related to hue (e.g., for red and green sample stimuli), in which case the means of the stimulus effect distributions should not drift towards each other with increasing delay, as was assumed by White and Cooney (1996). For this reason, it was assumed in the present model that distribution variances increased with increasing delay, with the location of the means remaining constant. A mediational view consistent with signal-detection interpretations of the stimulus effect dimension is that a value x on a given trial provides "evidence" on the basis of which the choice response is made. For example, for red versus triangle samples (for stimuli on different dimensions), the value of x provides evidence for the redness versus triangularity of the previously presented sample. As it happens, the present model is silent on these alternatives, but does emphasize that the stimulus effect value on which the choice is conditional in the model is defined at the time of remembering, consistent with the treatment of remembering as delayed stimulus control.

As a final note, we offer a comment on the scope of the model. The aim was to present a model that incorporated the effects of the reinforcer ratio in delayed matching-to-sample procedures, or other remembering procedures in which the choice alternatives were explicit. The general prediction that the reinforcer effect becomes more influential as

the delay interval lengthens has been noted before (Wixted, 1989) and was given a quantitative basis in the present model. Although the present paper has focused on remembering, the model may be extended in future applications to account for stimulus–reinforcer interactions in signal-detection procedures and may be compared to other possible models in relation to the prediction of such interactions (Nevin et al., 1993). It is of interest to note that the model described by Alsop and Davison (1991, Equation 3) predicts that reinforcer sensitivity decreases with increasing stimulus discriminability when discrimination between the choice alternatives is imperfect. Future research involving manipulation of comparison–stimulus disparity in delayed matching and signal-detection procedures (cf. Nevin et al., 1993; White, 1986) might assist in determining whether the present model and that described by Alsop and Davison are equally successful in accounting for the inverse relation between reinforcer sensitivity and discriminability described in the present paper.

Given programmed reinforcer probabilities, the present model predicts obtained reinforcer frequencies as well as response frequencies, and on the basis of an assumption of strict matching of the ratio of choice probabilities on a given trial to reinforcer probabilities (given a value of x), the model predicts the power function relation between response and reinforcer ratios. Also to be explored in future analyses, the model predicts change in performance from the early trials in which the reinforcer effects are variable owing to few obtained reinforcers, to the “steady-state” performance shown here. The main prediction offered by the model, as emphasized in the present paper, is the inverse relation between discriminability and reinforcer sensitivity. Although it is possible to incorporate in the model factors such as constant response bias and proactive effects of reinforcers on prior trials, the simplest version of the model has two free parameters, one for the distance between the distribution means and one for the standard deviation of the distributions. Although there is no free parameter that describes the reinforcer effect, the model predicts systematic variation in reinforcer effect on the basis of factors affecting discriminability.

REFERENCES

- Alsop, B., & Davison, M. (1991). Effects of varying stimulus disparity and the reinforcer ratio in concurrent-schedule and signal-detection procedures. *Journal of the Experimental Analysis of Behavior*, *56*, 67–80.
- Anderson, J. R. (1995). *Learning and memory: An integrated approach*. New York: Wiley.
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, *2*, 396–408.
- Banks, W. P. (1970). Signal detection theory and human memory. *Psychological Bulletin*, *74*, 81–99.
- Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, *22*, 231–242.
- Blough, D. S. (1998). Context reinforcement degrades discriminative control: A memory approach. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 185–199.
- Creelman, C. D., & Donaldson, W. (1968). ROC curves for discrimination of linear extent. *Journal of Experimental Psychology*, *77*, 514–516.
- Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum.
- Davison, M. C., & Tustin, D. (1978). The relation between the generalized matching law and signal-detection theory. *Journal of the Experimental Analysis of Behavior*, *29*, 331–336.
- Dusoir, A. E. (1975). Treatments of bias in detection and recognition models: A review. *Perception & Psychophysics*, *17*, 167–178.
- Ebbinghaus, H. (1964). *Memory: A contribution to experimental psychology* (H. A. Ruger, C. E. Bussenius, & E. R. Hilgard, Trans.). New York: Dover. (Original work published 1885)
- Edhouse, W. V., & White, K. G. (1988). Sources of proactive interference in animal memory. *Journal of Experimental Psychology: Animal Behavior Processes*, *14*, 56–71.
- Egan, J. P. (1975). *Signal detection theory and ROC analysis*. New York: Academic Press.
- Fechner, G. T. (1860). *Elemente der psychophysik*. Leipzig: Breitkopf & Hartel.
- Fetterman, J. G. (1996). Dimensions of stimulus complexity. *Journal of Experimental Psychology: Animal Behavior Processes*, *22*, 3–18.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Healy, A. F., & Kubovy, M. (1981). Probability matching and the formation of conservative decision rules in a numerical analog of signal detection. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 344–354.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, *13*, 243–266.
- Jones, B. M., & White, K. G. (1992). Stimulus discriminability and sensitivity to reinforcement in delayed matching to sample. *Journal of the Experimental Analysis of Behavior*, *58*, 159–172.
- Jones, B. M., White, K. G., & Alsop, B. (1995). On two effects of signaling the consequences for remembering. *Animal Learning & Behavior*, *23*, 256–272.
- Link, S. W. (1994). Rediscovering the past: Gustav Fechner and signal detection theory. *Psychological Science*, *5*, 335–340.

- Lockhart, R. S., & Murdock, B. B. (1970). Memory and the theory of signal detection. *Psychological Review*, *74*, 100–109.
- Luce, R. D. (1963). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1, pp. 103–189). New York: Wiley.
- Luce, R. D. (1994). Thurstone and sensory scaling: Then and now. *Psychological Review*, *101*, 271–277.
- Macmillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide*. New York: Cambridge University Press.
- McCarthy, D., & Davison, M. (1991). The interaction between stimulus and reinforcer control on remembering. *Journal of the Experimental Analysis of Behavior*, *56*, 51–66.
- McCarthy, D., & Voss, P. (1995). Delayed matching-to-sample performance: Effects of relative reinforcer frequency and of signaled versus unsignaled reinforcer frequencies. *Journal of the Experimental Analysis of Behavior*, *63*, 33–52.
- McCarthy, D., & White, K. G. (1987). Behavioral models of delayed detection and their application to memory. In M. L. Commons, J. Mazur, J. A. Nevin, & H. C. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 5. The effect of delay and intervening events on reinforcement value* (pp. 29–54). New York: Erlbaum.
- McLean, A. P., & White, K. G. (1983). Temporal constraint on choice: Sensitivity and bias in multiple schedules. *Journal of the Experimental Analysis of Behavior*, *39*, 405–426.
- Murdock, B. B. (1965). Signal-detection theory and short-term memory. *Journal of Experimental Psychology*, *70*, 443–447.
- Nevin, J. A. (1969). Signal detection theory and operant behavior: A review of David M. Green and John A. Swets' *Signal Detection Theory and Psychophysics*. *Journal of the Experimental Analysis of Behavior*, *12*, 475–480.
- Nevin, J. A. (1981). Psychophysics and reinforcement schedules. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 1. Discriminative properties of reinforcement schedules* (pp. 3–27). Hillsdale, NJ: Erlbaum.
- Nevin, J. A., Cate, H., & Alsop, B. (1993). Effects of differences between stimuli, responses, and reinforcer rates on conditional discrimination performance. *Journal of the Experimental Analysis of Behavior*, *59*, 147–161.
- Nevin, J. A., & Grosch, J. (1990). Effects of signaled reinforcer magnitude on delayed matching-to-sample performance. *Journal of Experimental Psychology: Animal Behavior Processes*, *16*, 298–305.
- Parks, T. E. (1966). Signal-detectability theory of recognition-memory performance. *Psychological Review*, *73*, 44–58.
- Ratcliff, R., McKoon, G., & Tindall, M. (1994). Empirical generality of data from recognition memory receiver-operating characteristic functions and implications for the global memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 763–785.
- Roberts, W. A. (1972). Short-term memory in the pigeon: Effects of repetition and spacing. *Journal of Experimental Psychology*, *94*, 74–83.
- Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, *117*, 34–50.
- Stevens, S. S. (1961). To honor Fechner and repeal his law. *Science*, *133*, 80–86.
- Thomas, E. A. C. (1975). Criterion adjustment and probability matching. *Perception & Psychophysics*, *18*, 158–162.
- Thomas, E. A. C., & Legge, D. (1970). Probability matching as a basis for detection and recognition decisions. *Psychological Review*, *77*, 65–72.
- Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review*, *34*, 273–286.
- White, K. G. (1985). Characteristics of forgetting functions in delayed matching to sample. *Journal of the Experimental Analysis of Behavior*, *44*, 15–34.
- White, K. G. (1986). Conjoint control of performance in conditional discrimination by successive and simultaneous stimuli. *Journal of the Experimental Analysis of Behavior*, *45*, 161–174.
- White, K. G. (1991). Psychophysics of direct remembering. In J. A. Commons, M. C. Davison, & J. A. Nevin (Eds.), *Models of behavior: Signal detection* (pp. 221–237). New York: Erlbaum.
- White, K. G., & Cooney, E. B. (1996). The consequences of remembering: Independence of performance at different retention intervals. *Journal of Experimental Psychology: Animal Behavior Processes*, *22*, 51–59.
- White, K. G., Pipe, M.-E., & McLean, A. P. (1985). A note on the measurement of stimulus discriminability in conditional discriminations. *Bulletin of the Psychonomic Society*, *23*, 153–155.
- Williams, B. A. (1988). Reinforcement, choice, and response strength. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology: Vol. 2. Learning and cognition* (pp. 167–244). New York: Wiley.
- Wixted, J. T. (1989). Nonhuman short-term memory: A quantitative reanalysis of selected findings. *Journal of the Experimental Analysis of Behavior*, *52*, 409–426.
- Wixted, J. T. (1993). A signal detection analysis of memory for nonoccurrence in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *19*, 400–411.
- Wixted, J. T., & Dougherty, D. H. (1996). Memory for asymmetric events. In D. Medin (Ed.), *The psychology of learning and motivation* (Vol. 35, pp. 89–126). San Diego: Academic Press.
- Wixted, J. T., & Ebbesen, E. B. (1991). On the form of forgetting. *Psychological Science*, *6*, 409–415.
- Yonelinas, A. P. (1994). Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 1341–1354.

Received April 29, 1998

Final acceptance August 13, 1998