

Reaction time signatures of discriminative processes: Differential effects of stimulus similarity and incentive

DONALD S. BLOUGH

Brown University, Providence, Rhode Island

In three experiments with pigeons, the similarity of unreinforced test stimuli to a reinforced stimulus and the frequency of reinforcement associated with a stimulus were varied. The stimulus on each trial was a small spot that appeared in different hues or, in Experiment 3, different forms. Differential response frequency and reaction time (RT) patterns emerged: Changes in similarity affected the percentage of stimuli responded to but left the shape of RT distributions about the same, whereas changes in reinforcement shifted RT distributions but had little effect on the percentage of responses. When the similarity and reinforcement variables were applied to the same stimuli (Experiment 2), their effects were largely independent. A generalization procedure (Experiment 3) replicated the similarity effects of the initial discrimination procedure. The RT distributions were modeled by a diffusion process, and implications for a memory-instance model were suggested.

Those who study discriminative processes often distinguish between two broad classes of independent variables, the first affecting primarily sensory or perceptual processes and the second affecting primarily control or decision processes. The most familiar formalization of this dichotomy is *signal detection theory*, which provides a recipe for separating and measuring the effects of these variables on response emission and choice (e.g., Green & Swets, 1966; Macmillan, 2002). Originally applied to human performance, detection theory soon influenced research with animals (e.g., Blough, 1967; Nevin, 1965). The influence of this approach has continued and is visible in, for example, the recent model of Davison and Nevin (1999), which conjoins measures of stimulus similarity and reinforcement to predict the distribution of responses in operant choice situations. However, although detection theory and its offspring have provided rules by which we may classify and predict the interacting effects of sensory and decision variables on animal behavior, the processes that underlie these rules are poorly understood.

The present experiments were designed to look for clues to discriminative processes by following a strategy that is often encountered in research on human cognition. This strategy depends on the assumption that mental processes take time and that patterns of overt reaction times (RTs) may help to reveal what is going on behind the scenes (see, e.g., Luce, 1986; Van Zandt, 2002). A fruitful tactic in the analysis of RTs has been to derive mathemat-

ical expressions of RT distributions with parameters that can be linked in a reasonable way to hypotheses about psychological processes, and in some of my previous experiments this tactic has been used to relate patterns of RTs to possible mechanisms involved in visual search. A finding of particular interest here is that the similarity between targets and distractors affects the momentary probability of target detection. This conclusion arose from the observation that RT distributions based on searches of varying difficulty could be fit by changing a single parameter representing response probability in a mathematical function, the *ex-Gaussian*, that was fit to the data (see below; see also Blough, 1988, 1989, 2000). However, this rather straightforward interpretation of stimulus-related effects on search RTs did not account for stimulus similarity effects in a nonsearch task (Blough, 1978), nor did it account for search RTs under variations of reinforcement (Blough, 1989, 2000).

The present experiments were intended to clarify the effects of similarity and reinforcement, and their possible interaction, by the use of a task not involving visual search. Instead, pigeons received discrimination training with single stimuli that varied in similarity and in degree of differential reinforcement. A trial-wise go/no-go paradigm yielded response percentages and RTs for several values of these variables, confining each variable to a distinctive target hue in Experiment 1 and measuring their joint effect in Experiment 2. Experiment 3 investigated whether generalization tests with novel stimuli produce the same RT patterns as those from the discrimination procedure used in Experiments 1 and 2.

The three experiments shared several procedural details. Pigeons pecked at a small bright spot that appeared on many brief trials and changed in hue from trial to trial; a single peck terminated the trial and sometimes brought

This research was supported in part by Grant MH61782 from the National Institutes of Health. Correspondence concerning this article should be addressed to D. S. Blough, Department of Psychology, Box 1853, Brown University, Providence, RI 02912 (e-mail: donald_blough@brown.edu).

food. RT was defined as the time from the onset of the target spot until the time that a response, if any, occurred. To minimize visual search time, the target spot appeared alone on the screen, but its location was randomized to discourage rapid, indiscriminate responses triggered simply by stimulus onset; such responses make RT data inhomogeneous and complicate their interpretation (e.g., Blough, 1978, 1996).

EXPERIMENT 1

In Experiment 1, red targets of four slightly different hues appeared during each session; responses to one of the hues were always reinforced and responses to the others were not. A green target also appeared; it was associated with reinforcement on a percentage of trials that changed between blocks of sessions.

Method

Subjects

Six female White Carneaux pigeons served as subjects. The birds had previously served briefly in conditioning demonstrations in an undergraduate laboratory and in two experiments involving search for achromatic targets that differed in form (Blough, 2002). The birds were maintained at approximately 85% of their free-feeding weight, and they received most of their food in the form of 45-mg Noyes pigeon pellets during experimental sessions; if necessary, this ration was supplemented by mixed grain given after the experimental session.

Apparatus and Stimulus Displays

Two experimental boxes contained $34 \times 30 \times 34$ -cm subject chambers. An opening in the front panel of each chamber gave access to an LCD color video monitor; the visible part of the monitor screen measured 18 cm wide \times 11 cm high. Infrared emitters and detectors (Smart-Frame, Carol Touch, Round Rock, TX) detected the positions of pecks to the screen. A pellet dispenser was located immediately below and to the right of the screen. A loudspeaker near the chambers provided white masking noise. Two IBM 486 computers controlled the experiments and recorded response data. RTs were recorded to the nearest 0.03 sec, a limitation imposed by the scan rate of the infrared sensors.

The discriminative stimuli were circular spots 5 mm in diameter. Most of the training and all of the experimental tests used spots of five different hues, only one of which appeared on any given trial. One target spot, the red S+, was a desaturated reddish hue; pecks at this target were always followed by reinforcement. Three other targets were quite similar in hue to the red S+, departing from it by small stepwise reductions in the activation of the red monitor phosphor and increases in the blue; to a human observer, these targets appeared to be reds of a slightly increasing purplish cast. The fourth target, the green S+, was quite different, appearing as a desaturated green to a human observer. The targets were of approximately equal luminance to the human eye, about 40 cd/m². During each experimental trial, one of these five targets appeared on a somewhat less luminous neutral gray background (approximately 13 cd/m²). The luminance of the screen remained at the background level during the intertrial interval (ITI); it rose to 65 cd/m² for 2 sec to signal the delivery of food pellets.

The target spot appeared on the display screen in one of eight locations defined by a matrix 7.6 cm wide \times 2 cm high, made up of four columns and two rows. Each of the four columns corresponded to one of four screen segments within which pecks were localized. In the following, a peck *to* a target means that the peck was directed toward the vertical column on the screen in which the target was located. All

of the targets were highly visible and elicited more rapid responses than targets in previous search tasks involving multiple targets.

Procedure

On each trial a single target spot appeared, followed by a blank screen that defined the ITI. Trials lasted 7 sec, timed from the onset of one target to the onset of the next. The total trial time, rather than ITI, was held constant so that response speed could not affect reinforcements per unit of time. However, as infrequently happened, a peck to any part of the screen during the ITI restarted the trial timer. At the end of the ITI one of the targets appeared, and it remained visible until the bird had pecked it or for a maximum of 3 sec. In either case, the target disappeared, and an RT was recorded; thus, a 3-sec "RT" actually marked a failure to respond. Responses to any part of the screen outside of the target column were ignored, as were responses with RTs less than 100 msec. On most trials the ITI began immediately after the pecks to the target were made; if reinforcement was scheduled the screen turned white, the feeder light came on for 2 sec, two pellets dropped into the feeder, and then the ITI began.

Training. The birds were briefly autoshaped to peck the red S+ and then were trained with five targets that varied in hue from red to reddish violet. This procedure went quickly, since the birds had previous experience in visual search experiments in the same apparatus. Over 20 sessions, the response requirement was reduced from two pecks to one; pecks to all targets were reinforced on 20% of trials. The sessions consisted of 800 trials during which the targets appeared in random sequence. Each target appeared 160 times, it appeared in each possible screen location an equal number of times, and pecks to it were reinforced at each screen location an equal number of times. This scheme of target localization and randomization was followed for the remainder of the three experiments.

A discrimination phase of 40 days followed, during which pecks to the red S+ produced reinforcement 100% of the time, whereas the other red targets were never followed by reinforcement. During these sessions, decreases in the hue range gradually increased the difficulty of the discrimination among the red stimuli. In the last stage of training, the green S+, reinforced on half of its appearances, replaced the most discriminable of the unreinforced red targets, to which responding had virtually ceased.

Testing. Testing consisted of successive blocks of sessions during which pecks at the green S+ yielded reinforcement at a rate that varied from block to block but was constant across the sessions within each block. For all sessions, pecks to the red S+ yielded reinforcement on 100% of the trials, whereas pecks to the other red stimuli were never reinforced. When reinforcement was scheduled, it was rarely missed by failure to respond within the allotted 3 sec except on some green S+ trials at the lowest rate of reinforcement (2.5%). In the case of a missed reinforcement in that condition, reinforcement was rescheduled for the next presentation of the missed target, regardless of the location in which it appeared. In this way, the birds received all or almost all of their scheduled reinforcements even at the lowest rate scheduled. In addition to the above arrangements, each session started with a reinforced presentation of the red S+. The sessions lasted approximately 2 h and were run daily, except for infrequent breaks lasting 1–4 days.

Table 1 displays the sequence of conditions for each bird. Early on, during exploration for appropriate reinforcement conditions, relatively few sessions were run per block of constant green S+ reinforcement. The reduction of reinforcement to the green S+ to 5% produced a substantial change in RT, and beginning at that point each bird was run on a given reinforcement condition until at least eight sessions of stable behavior had been collected. Stability was defined as a change of less than 10% in the mean RT to the green S+ target across the days to be used for RT analysis. Taken together, these eight sessions provided the data used to define RT means and distributions. The birds differed somewhat in the number of sessions they received within a given block, because they differed in the

Table 1
Number of Sessions Run on Each Condition

% Reinf	Green S+	Bird 064	Bird 136	Bird 146	Bird 148	Bird 151	Bird 159
50		5	5	5	5	4	5
20		8	8	8	8	4	4
10		8	8	9	9	9	9
5		13	13	13	13	13	13
2.5		27	27	18	29	16	16
5		14	14	23	13	23	24
100		17	17	17	16	16	16

speed with which they responded to reinforcement changes and also in the stability of their performance.

Results and Discussion

In the following analyses, mean RTs for each bird are represented by the geometric mean RT across sessions for each target. For most analyses these mean RTs were converted to speeds (1/RT). This was done because the speed measure is less sensitive than RT to long outliers and also facilitates comparison with response frequency, since increases in both imply greater responsiveness.

Figure 1 shows mean speeds across birds for each target as a function of daily sessions. For this initial summary of responsiveness, trials on which no response occurred were counted as having a 3-sec RT, the maximum duration of the display. The percentages of green S+ reinforcement for successive blocks of sessions are indicated in the figure. All sessions are shown for the first four conditions (50%, 20%, 10%, and 5% reinforcement on green S+), ex-

cept that these averages omitted a few extra sessions for some birds (see Table 1). For the last three conditions (2.5%, 5%, 100%), in which the block length varied somewhat across birds, the data include the first five sessions of each block from all birds. A break in the curves indicates that a number of days are omitted, differing across birds. After this break, a second segment of the functions represents the means across birds from the eight sessions that were used in the main data analyses, as was specified above.

The data in Figure 1 show that the speed of response to the green S+ target tracked the rate of reinforcement to that target, with the fastest responses occurring at 50% (in the early sessions, left-hand side) and 100% (final sessions, right-hand side). Much of the adjustment to new reinforcement conditions occurred within the first five or six sessions on a given condition. As was expected, the responsiveness to the red targets was greatest to the S+ and dropped with decreasing similarity to the S+. The birds rarely responded to the most dissimilar red stimulus, and

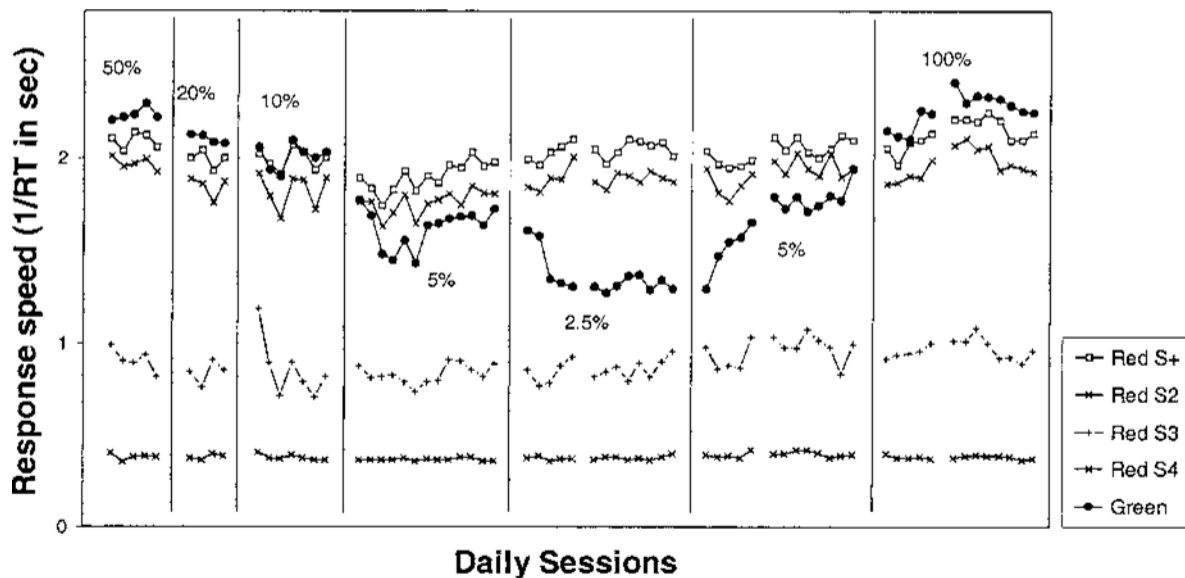


Figure 1. The mean speeds for the five different stimuli for the daily sessions of Experiment 1. Speeds were computed from mean reaction times (RTs) for which a nonresponse was counted as a reaction time of 3 sec, the maximum stimulus-on time. Vertical lines separate blocks of sessions during which reinforcement for the green S+ stimulus was held constant. The red S+ was reinforced on all trials, and the other red stimuli were never reinforced. (Certain sessions have been omitted, which are indicated by breaks in the curves. See the text for details.)

because RTs for nonresponse trials were counted as 3 sec, the mean speed for that target approached a limit of 1/3 responses/sec. Although the overall speed of response drifted somewhat during the experiment, the relative speeds to the red targets stayed fairly constant; for example, the two 5% reinforcement conditions yielded similar relations among the response speeds to different targets, although the overall speed was somewhat higher in the second than in the first of those conditions.

As was noted, the summary in Figure 1 combines RTs for actual pecks with 3-sec default RTs from trials on which no peck occurred. To differentiate these outcomes, trials terminating at 3 sec with no peck were discarded, and the number and geometric mean RT for the remaining trials were determined for each stimulus in each session. These new mean RTs were again converted to speeds, and finally, to compensate for individual differences in overall speed, they were converted to a percentage of each bird's overall mean speed. Figure 2 shows the means across birds of this relative speed of response for the various targets, together with the percentage of trials responded to. (Data from the second of the two 5% reinforcement blocks were used in this and subsequent analyses.) A single-factor analysis of variance (ANOVA) indicated that reinforcement had a significant effect on each measure from the green target trials [N responses, $F(4,25) = 4.55, p < .01$; speed, $F(4,25) = 24, p < .001$]. Mixed-model ANOVAs tested the effects of both variables on responding to the set of red stimuli. For N responses to the red stimuli, similarity had a significant effect [$F(3,12) = 149, p < .001$], the effect of reinforcement of the green stimulus approached significance [$F(4,12) = 2.4, p = .09$], and the interaction was significant [$F(12,60) = 2.1, p = .03$]. For relative speed to the red stimuli, similarity also had a significant effect [$F(3,12) = 4.8, p = .02$], the level of reinforcement to the green stimulus did not [$F(4,12) = 0.18, p = .95$], and the interaction was significant [$F(12,60) = 108, p = .03$].

It is evident that decreases in reinforcement and in similarity were associated with reduced responsiveness, but, as is shown in Figure 2, these effects were expressed in different ways. For the green S+, response speed decreased substantially with reinforcement reduction, yet the birds responded on almost all trials, even at the leanest (2.5%) schedule. This insensitivity of trialwise response probability to low reinforcement has been noted previously (e.g., Killeen & Hall, 2001). For the red targets, the speed of response changed rather little with stimulus similarity, but the percentage of trials responded to dropped to nearly zero for the red stimulus most dissimilar to S+ and reached nearly 100% for the red stimulus that was most similar.

RT distributions enable a more detailed look at these differing effects of reinforcement probability and stimulus similarity. To construct such distributions, RTs from each session were distributed into 0.1-sec bins. Corresponding bins were averaged across the final eight sessions for the last three conditions (those during which the green S+ was reinforced at 2.5%, 5%, and 100%) within and finally

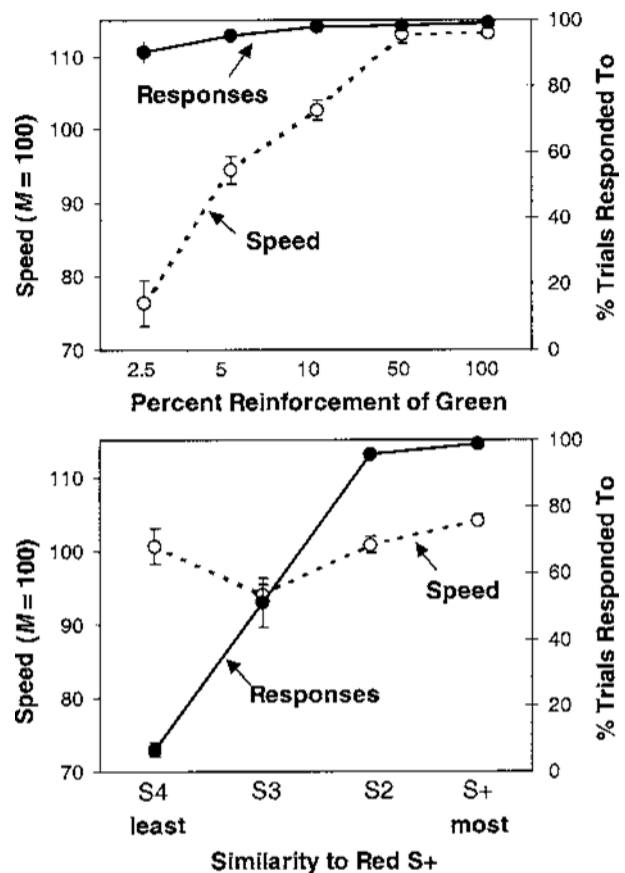


Figure 2. Top panel: Response speeds of emitted responses relative to bird means ($M = 100$; left ordinate) and number of responses (right ordinate) emitted to the green S+ stimulus under the five reinforcement conditions of Experiment 1. Data are means of 8 days for most conditions and 5 days for 50% and 10% conditions. Bottom panel: The same two response measures as in the top panel from trials with the four red stimuli, averaged across the reinforcement conditions described above. Bars show standard errors of the mean. Changing the rate of reinforcement affected primarily response speed and not number of responses emitted (top), whereas changing the similarity of the target to the S+ changed primarily the number of responses emitted and speed to a lesser extent (bottom).

across birds. The distributions for the red stimuli differed very little among these conditions, and they were averaged into a single set. The distributions for the green stimulus differed between the conditions, and they were kept separate.

Figure 3 summarizes the results, with distributions for green and red stimuli displayed in separate panels. Reduced reinforcement caused the distributions from green stimulus trials to shift to the right and to become more skewed to the right. In contrast, the distributions from the four red stimuli have similar shapes; their height varies, as is shown in Figure 2, because fewer responses were made to the dissimilar stimuli. The changes in mean speed that did occur for the red stimuli (Figure 2) were not associated, as for green, with a shift in the mode of the distribu-

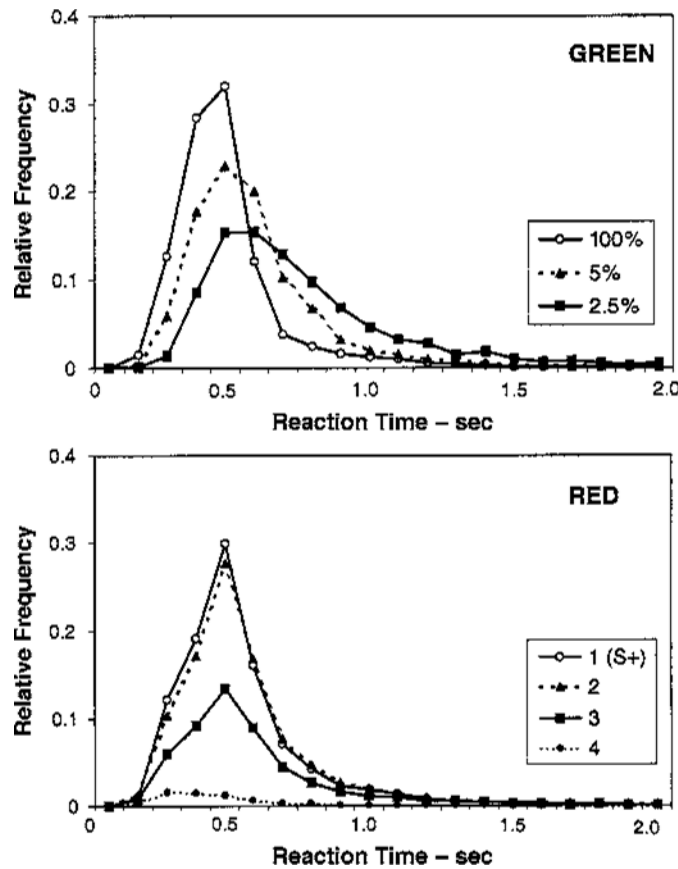


Figure 3. Mean reaction time distributions across the last eight sessions and all birds (Experiment 1). Top panel: Data for the green target at three different reinforcement densities (100%, 5%, and 2.5%). Note the differing modes and curve shapes. Bottom panel: Data for the red S+ and the three other red targets. Note that the distributions are much the same in shape and mode, except for the lowest curve, which represents a stimulus that elicited very few responses (see the text for details).

tion, but resulted primarily from small relative changes in the number of very long RTs or very short RTs. Notably, on average, the birds responded to the most dissimilar stimulus (Stimulus 4) on only 6% of trials, and the resulting low curve was shifted toward shorter RTs. This data pattern has been seen before (Blough, 1978) and probably reflects the presence of some very short RT responses that were triggered by stimulus onset regardless of hue. Although these presumably occurred for all stimuli, they predominate only when there are very few controlled responses. As was mentioned above, the procedure was intended to minimize such responses, but evidently it did not get rid of them entirely. These short RTs were the source of the up-turn at S1 for the speed curve for red (Figure 2, lower panel); few of them were recorded in Experiment 2.

From Figure 3 one can see that the distributions for the green stimulus shifted not only with respect to each other from condition to condition, but also with respect to the

distributions for the red stimuli that, as has already been noted, changed little across blocks of sessions. That is, the reinforcement effect appears both as a between-sessions effect on the green stimulus and as a within-sessions effect displayed by the red–green comparison.

Mean RT distributions sometimes misrepresent individual data and may produce a misleading impression of the data. Therefore, Figure 4 shows key results for each of the six birds during the first 1.2 sec of the trials. Each panel shows two pairs of curves. The solid lines represent RTs to the green S+; the curve to the left shows the results for 100% reinforcement, the one to the right for 2.5%. The dashed lines represent RTs to red stimuli averaged over the same two sets of sessions as those for the green curves. The upper dashed curve in each panel represents RTs to the red S+, the lower one to red S3. Although the shape of the curves differs across birds, each bird produced the essential results noted above: The two green (solid) curves are shifted with respect to each other on the abscissa,

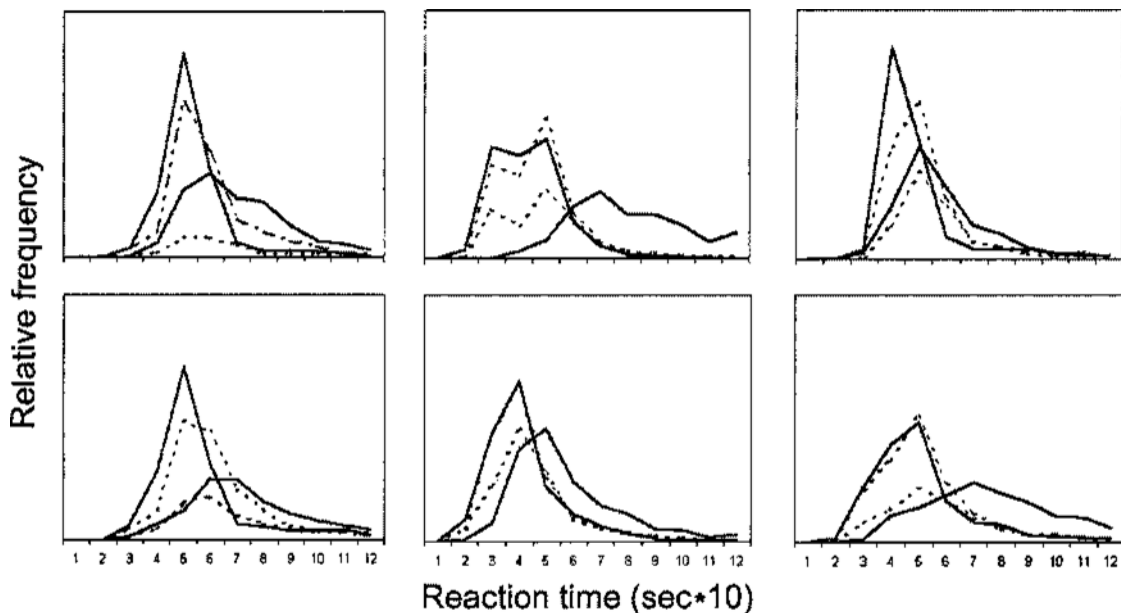


Figure 4. Mean reaction time distributions across the last eight sessions for each bird (Experiment 1). The solid curves come from responses to the green target, with the higher, leftmost function from 100% reinforcement sessions and the lower, rightmost function from 2.5% reinforcement sessions. The dashed curves come from responses to the red targets, with the higher function from responses to the S+ and the lower from responses to a target of intermediate similarity to the S+. The ordinate scale is the same for all the panels. Despite individual differences, each bird shows the overall response pattern summarized in Figure 3. (In the center lower panel, the taller dashed curve is obscured by the very similar solid curve.)

whereas the two red (dashed) curves are similar in shape and differ mainly in the total number of responses represented by each curve.

An important aspect of the data summarized in Figures 3 and 4 is that the RT distributions for the various red targets appear to be approximately proportional to one another. Cumulative distributions provide a test of this observation, for, if it is true, plotting one cumulative function against another should yield a straight line. Figure 5 displays the data in this form. For the green target, mean cumulative curves for stimuli at 5% and 2.5% reinforcement appear as functions of the 100% cumulative distribution. Similarly, curves for each red target are plotted as a function of the curve for the red S+. Straight lines provide a reasonably good fit to the red target data in the lower panel, supporting the idea that the distributions are proportional. In contrast, the curves representing different percentages of reinforcement are clearly nonlinear; the relative displacement of the distributions in Figure 4 is rendered as curvature on these plots in the upper panel of Figure 5. (The reader is invited to mentally flip the functions of Figure 5 from the lower right to the upper left quadrant of the displays. In that configuration, the functions for the green stimulus can represent receiver operating characteristic [ROC] functions of signal detection theory, with RT as a confidence rating. The straight functions from the red stimulus correspond to the ROC curves expected from *high threshold* theory. See Macmillan, 2002, and Blough, 1967, 1978, 2001, for further explanation and discussion.)

The data from this experiment confirm the results of a discrete-trial discrimination study in which RT distributions were little affected by stimulus similarity, despite large changes in response probability (Blough, 1978). The results are also consistent with, and serve to generalize, data showing a reinforcement contrast effect on visual search speed (Blough, 1989, 2000). Most important, the distinctive response patterns suggest that although both of the variables affect responsiveness, they do so via different processes.

EXPERIMENT 2

The results of Experiment 1 support previous evidence that similarity and reinforcement differ in their effects on RT. Experiment 2 explored the possible interaction of these variables by measuring their effects simultaneously on one set of stimuli. A lack of interaction would support the idea that these variables affect independent processes that might be associated with different processing stages.

Method

Subjects and Apparatus

The apparatus and birds that were used in Experiment 1 also served in Experiment 2; the birds were maintained in the same way as that described above. The apparatus was the same as that in Experiment 1.

Procedure

The display and trial details, randomization, and session length were the same as those in Experiment 1. As before, one of five stim-

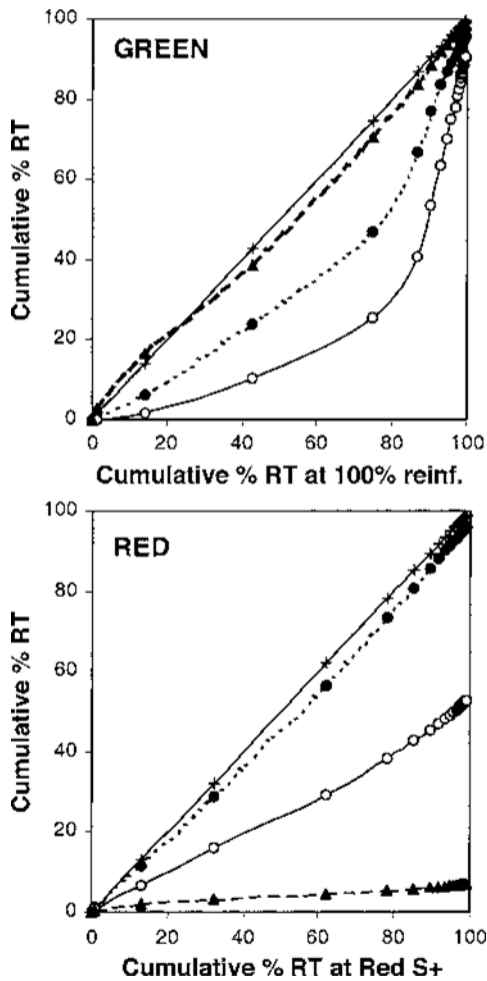


Figure 5. Cumulative distributions corresponding to the mean density functions from Experiment 1, shown in Figure 4. **Top panel:** The distributions for 50%, 5%, and 2.5% reinforcement of the green target are plotted on the ordinate against the 100% distribution on the abscissa. **Bottom panel:** The distributions for the three unreinforced red test stimuli are plotted against the red S+ distribution on the abscissa. These functions approximate the straight lines that would result if the distributions were multiples of one another.

uli appeared on each trial, either a green circular spot or one of four red spots differing slightly in hue. The data from the final condition of Experiment 1 were used for comparison here. During that condition, pecks to both the red S+ and the green S+ brought reinforcement on 100% of the trials. The new conditions of Experiment 2 started immediately after those data had been collected. A peck to the green S+ spot continued to be reinforced on 100% of the trials, but a peck to the red S+ now brought food on a reduced percentage of trials: For two sessions, reinforcement for the red S+ was set at 5%; for the next session it was set at 20%; it was set at 15% for the remainder of the experiment. If a bird failed to peck at the red S+ when reinforcement was scheduled, a pellet was delivered for a peck on the next presentation of the red S+. As previously, the other three red stimuli were never followed by reinforcement. Testing at 15% reinforcement for red S+ continued for 20 sessions (25 sessions for one bird whose behavior was briefly unstable).

Results and Discussion

The data were analyzed as in Experiment 1. Figure 6 shows the mean speeds during the last eight sessions with 100% reinforcement for both red S+ and green S+, followed by mean speeds during the subsequent sessions with reduced reinforcement for the red S+. Like the summary in Experiment 1, this overall summary counted a nonresponse as a 3-sec RT. The speed of response to the red stimuli dropped when reinforcement to the red S+ shifted from 100% to 5% on the first 2 days of the second condition. It recovered somewhat when reinforcement was increased, and at 15% reinforcement it reached a relatively stable intermediate level. Notably, speeds to the four red stimuli changed in tandem; that is, the changes in reinforcement affected all of them, but the relations between them stayed about the same. Also, although pecks to the green target were maintained at 100% reinforcement throughout the experiment, responsiveness to that target dropped somewhat when reinforcement to the red S+ was reduced.

As in Experiment 1, a more detailed analysis was performed to determine the number and RTs of responses actually emitted to each stimulus. The means of these measures for the last 8 days of the 15% condition were compared with those measures for the 100% condition from Experiment 1. As before, speeds were computed from the geometric mean of each bird's RTs, omitting trials with no response, and group means were determined after the individual bird results had been scaled in such a way that each bird's overall mean RT equaled 100. These measures of response frequency and speed under the two reinforcement conditions appear in Figure 7.

Figure 7 shows a pattern of results similar to that in Experiment 1 (Figure 3). As similarity to the red S+ decreased, the percentage of trials that elicited a response fell from near 100% to near 0%. A drop in speed as similarity moved from S+ to S3 is clearer here than in Experiment 1, probably because few uncontrolled short RT responses were recorded. As in Experiment 1, the decrease in reinforcement produced a different pattern: a substantial drop in speed, but almost no change in the response percentage. Confirming these impressions, a mixed-model ANOVA applied to the percentage of response showed an effect of similarity [$F(3,15) = 241, p < .001$] but not of reinforcement [$F(1,5) = 1.2, p = .33$] and no significant interaction [$F(3,15) = 2.4, p = .11$]. The ANOVA for speed showed main effects of reinforcement [$F(1,5) = 41, p < .001$] and similarity [$F(3,15) = 49, p < .001$] but no significant interaction [$F(3,15) = 19, p = .62$]. The reduction in reinforcement of responses to the red stimulus was also associated with a relative reduction in speed of responses to the green stimulus, which a two-tailed *t* test showed to be significant [$t(5) = 2.8, p = .04$]. This change amounted to 5%, about one quarter the size of the mean RT change to the red stimuli.

RT distributions for key stimuli in Experiment 2 appear in Figure 8. Shown are the data based on the red S+, which

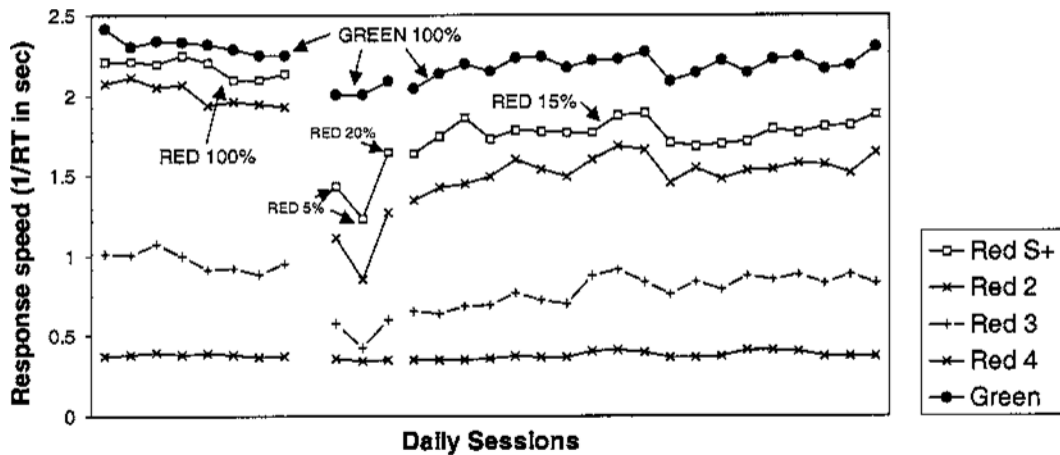


Figure 6. Mean speeds for each of the daily sessions of Experiment 2, computed in the same way as those for Figure 1. Pecks to the green target were reinforced on all trials; reinforcement for the red S+ was initially 100% and finally 15%, as is indicated in the figure; during three transition sessions, set off by spaces, reinforcement was 5%, 5%, and 20%.

elicited responses on almost all trials, and on the red Stimulus 3 (S3), which elicited responses on about half of the trials. The functions repeat the major features of those from Experiment 1 (Figure 3): The drop in reinforcement from 100% to 15% shifted the curves to the right, and the most prominent difference between S+ and S3 was in the height of the curves. More noticeable here than in Experiment 1 is a smaller rightward shift attributable to similarity (the difference between the S+ and S3 curves at

each percentage of reinforcement); this corresponds to the slope in the similarity curves in Figure 7. As has already been noted, short, uncontrolled RTs were largely absent in the data of Experiment 2, perhaps as a result of continued training.

The absence of an interaction between similarity and reinforcement effects is consistent with a good deal of other evidence for the independent operation of stimulus and incentive (e.g., Roberts, 1987). More exactly, if similarity

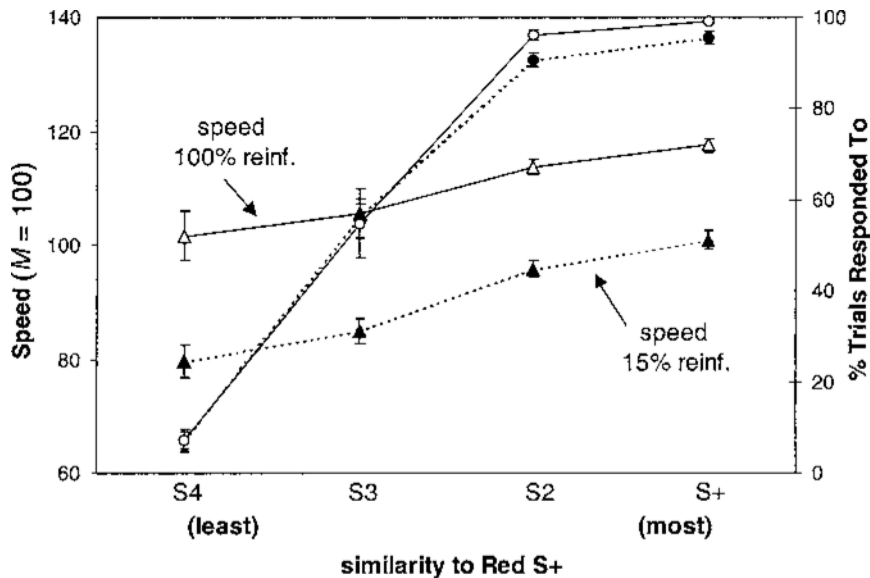


Figure 7. Data from the red stimuli in Experiment 2. The steep curves show the percentage of responding to stimuli of differing similarity, quantified on the right-hand ordinate. The curves labeled "speed" are quantified on the left ordinate, where values have been computed relative to each bird's overall mean speed set at 100. Note that the effect of reinforcement is approximately constant across the different similarities, suggesting that the effect of this variable is independent of similarity. Bars show standard error of the mean; for several points these are too short to be visible.




Figure 8. Selected reaction time distributions from Experiment 2, showing data from the red S+ (open circles, upper curves) and the red S3, to which the birds made about half as many responses. As in Experiment 1, a rightward shift of the curves is caused mainly by the change in reinforcement from 100% (solid lines) to 15% (dotted lines).

affected processes that were independent of those affected by reinforcement, and the processes occurred sequentially, a change in one variable should change RT by the same amount at each level of the other variable (e.g., Sternberg, 1998). The best test of this hypothesis requires the use of untransformed mean RTs rather than speeds. A graph based on that transformation was similar to Figure 7, and a two-factor ANOVA on normalized RT data gave essentially the same outcome as the one reported above for speed. The absence of a significant interaction is consistent with additivity of the effects in time.

In summary, Experiment 2 produced within a single set of red targets the same general pattern of responses that appeared separately for red and green targets in Experiment 1, with reinforcement affecting mainly RT and similarity affecting mainly the number of responses elicited by the target. The principal new finding was that when both variables operated on the same stimuli their effects appeared to be largely independent.

EXPERIMENT 3

In Experiments 1 and 2 small changes in a physical attribute of a stimulus resulted in response patterns unlike those resulting from changes in the reinforcement associated with a stimulus. In those experiments, however, the similarity effect was measured after many sessions of differential reinforcement on the tested discrimination. That training could have been an important determinant of the observed response patterns, for it is widely held that the effects of differential reinforcement differ fundamentally from those produced by generalization tests conducted in extinction with novel stimuli. Behavior theories going back at least to Spence (1937) propose that a generalization decre-

ment reflects reduced excitatory strength, whereas a discrimination decrement often reflects a learned inhibition of responses to negative stimuli. Some cognitive theories make a related distinction, though in a very different conceptual context. A fairly recent example is a distinction between *discriminable* and *confusable* stimuli, which appear to differ in the similarity functions that are found to relate them (Ennis, 1988; Shepard, 1988).

These considerations suggested Experiment 3, which measured RTs during generalization tests to novel stimuli that were similar to a training S+. The test dimension was target shape. The circular red S+ used in the other experiments served as the training stimulus. It appeared in unreinforced generalization tests together with three elliptical forms of the same hue but of differing eccentricity. After the generalization tests, the birds were given discrimination training, in which responding to the circular red target was reinforced and responding to the elliptical targets was extinguished. This was done to confirm that relative responding during the generalization tests did in fact differ substantially from that obtained following discrimination training and that therefore it might be controlled by somewhat different factors.

Method

Subjects

The subjects in Experiment 3 were the same six birds as those in the previous experiments, but one bird died during the study so the data from the remaining five are reported here. Maintenance was performed in the same way as that described previously.

Apparatus and Stimuli

The apparatus was the same as that in the previous experiments. The stimulus displays differed in the following respects. There were four unreinforced test targets, all of which were the same hue as the

former red S+. One of these was circular, identical to the former red S+. The others were ellipses with a horizontal major axis increasing in its length in regular steps. The minor axes were adjusted to yield an approximately constant stimulus area for all four red targets. To help maintain responding during tests, a distinctive reinforced target also appeared. This was a large (12×12 -mm) square with the same hue as the previous green S+ target. Its hue, size, and shape were chosen to minimize any influence it might have on the generalization test with the small red elliptical targets.

Procedure

Pretraining. Prior to this experiment the birds served in a brief pilot study in which the same red S+ was used, but the green hue and the reinforcement associated with it were varied. Several sessions then reestablished baseline responding, with display characteristics, trial timing, randomization, and session length the same as those in Experiments 1 and 2. This pretraining comprised four 800-trial sessions in which the red S+ and the green S+ of Experiments 1 and 2 each appeared on half of the trials; pecks to the red S+ brought reinforcement on 20% of the trials for two sessions, and on 10% for two sessions. Pecks to the green stimulus brought reinforcement on 40%, then 20%, and finally 10% of the trials. At this point the shape of the green S+ was changed to the large block described above. Four sessions were run, with trials divided equally between the new green block S+ and the red S+, each reinforced on 10% of the trials.

Shape generalization. In the generalization tests, the four red shapes described above were used. No pecks to any red target were reinforced except in the first test session, which began with a single reinforced presentation of the circular target, the former red S+. In addition, the green block S+ appeared and yielded reinforcement on 100% of its presentations; this was done to help maintain overall responsiveness. These five targets (one green, four red) appeared equally often in 800-trial sessions structured as in the earlier experiments. Three generalization sessions were scheduled, but two of the birds responded so little on the second session that they were not run on a third; only data from the first two sessions are reported here.

Shape discrimination. In a final series of sessions, a discrimination procedure was used, which was the same as the shape general-

ization test procedure except that pecks to the circular red S+ always brought reinforcement, as did those to the green block S+. Pecks to the elliptical test forms were never reinforced. Each bird was run for 10 discrimination sessions or until it responded less than 20% of the time to the ellipse most similar to the red S+.

Results and Discussion

The data were analyzed as in Experiment 1. Figure 9 displays, for the first two generalization tests, the mean across birds of the cumulative distributions of RTs to each of the elliptical stimuli as a function of the cumulative distribution to the circular stimulus used in training. These functions terminate at 85% on the abscissa on the first test day and at 52% on the second; this number is the average percentage of trials on which the red circle S+ elicited responses. The termination points on the ordinate for the various functions are lower, and each represents the average percentage of trials on which the birds responded to the various test stimuli. It is evident that similarity strongly affected the number of responses emitted to the test stimuli, as in Experiments 1 and 2. The functions also approximate straight lines, except for a small bend at the longest time intervals, which indicates that a small percentage of responses to the generalization stimuli was somewhat delayed. Otherwise, as was described earlier, the functions indicate that the distributions are nearly multiples of one another and thus have approximately the same shape. This result also reproduces the outcome for the different red stimuli in Experiments 1 and 2.

Figure 10 shows the mean number of responses elicited by the test shapes in the first generalization test (upper curve) and in the last shape discrimination training session (bottom curve) as a percentage of responses to the circular red (S+) target. It is apparent that, as was antici-

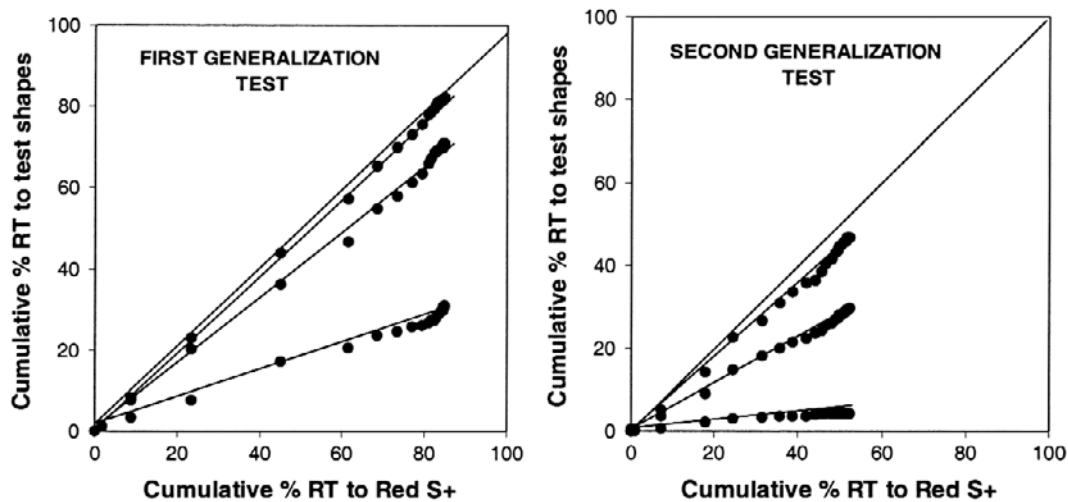


Figure 9. Cumulative distributions corresponding to the mean density functions from the first (left) and second (right) generalization tests of Experiment 3. The distributions for the three elliptical targets are plotted against distribution for the circular target, which was trained as S+. The fact that these functions approximate straight lines indicates that the different stimuli yielded distributions of approximately the same shape during the generalization test (see the text for details).




Figure 10. Responses to each of the different shapes tested in Experiment 3, given as a percentage of responses to the red circular stimulus. The data from the first generalization test appear toward the top of the figure (solid line, open circles), and those from the last discrimination session appear toward the bottom of the figure (dotted line, solid squares). The curves show that relative responding during the two tests differed substantially.

pated, the gradient following discrimination training was much sharper than the generalization gradient, which suggests the possibility, unrealized here, that RT patterns might differ as well. However, this matter remains somewhat open, because one cannot assume that these data represent generalization under all conditions. For example, although the generalization test was along a novel dimension, the red S+ had long been involved in a hue discrimination, and it is well known that discrimination training can sharpen generalization gradients even if such training occurs on dimensions other than the one that is tested. This effect is sometimes attributed to the weakening of control by context stimuli (e.g., Mackintosh, 1983).

In summary, the pattern of responses to different shapes in Experiment 3 was much like the pattern to different red hues in Experiments 1 and 2. This suggests that the generalization test stimuli were processed in the same way as the discriminative stimuli tested in the previous experiments or, at any rate, that during the generalization tests the birds did not shift to a mode of processing in which similarity markedly affected RTs.

GENERAL DISCUSSION

The main results of these experiments may be summarized as follows. Experiment 1 showed these effects: (1) As stimuli became less similar in hue to a red S+, fewer responses occurred; mean response speed also decreased slightly, but RT distributions remained largely unchanged in shape. (2) When the reinforcement associated with a green stimulus diminished, the RT distribution of responses shifted, but the number of responses emitted changed only slightly. Experiment 2 replicated these effects within one set of red stimuli, and its results suggested

that the effects of reinforcement and similarity are largely independent. In Experiment 3 a generalization procedure yielded a pattern of results across stimuli similar to that from the discrimination procedure used in Experiments 1 and 2.

We have seen that the data replicate and expand results from previous studies (e.g., Blough, 1978, 1989, 2000), and I will now consider how the results might shape hypotheses about discriminative processes. RT distributions are featured in this inquiry, having been quite useful in previous analyses of cognitive processes in pigeons (see Blough, 1988, 1989, 1992, 2000) and humans (see Luce, 1986; Ratcliff, 1978, 1988; Van Zandt, 2002).

A useful theoretical position, common to most theories of discrimination and choice, holds that responding is controlled by a comparison of sensory input with preexisting “remembered” information, the latter modulated by incentive and other factors. In the signal detection theory (SDT) implementation of this hypothesis, sensory information is represented by an internal distribution of input values, preexisting information is represented by a criterion, and percentage of responses is determined by the proportion of the sensory distribution that falls beyond the criterion (e.g., Blough, 2001; Macmillan, 2002). This SDT scheme handles the effects of stimulus similarity on the percentage of responses observed in these and related experiments quite well, as I have discussed at length elsewhere (see Blough, 1978, 2001). Also in agreement with SDT is the evidence, noted above, that the similarity and reinforcement effects are independent; the former can be represented by shifts in the sensory distribution and the latter by shifts in the criterion.

However, unlike percentage of responses, the RT data recorded here do not find a natural home in the SDT account. RT was relatively constant across stimuli, and so it

cannot play the role of an SDT “confidence rating,” as can response rate (see Blough, 1978, 2001). Nor did RT decrease with increases in the putative distance between the internal stimulus representation and the criterion, as it does in related human experiments (e.g., Espinoza-Varas & Watson, 1994; Luce, 1986). In short, the RT data seem to reveal little about the sensory or input side of the decision process, except perhaps to suggest that stimulus processing speed does not vary greatly with similarity under these conditions.

On the other hand, the changes in RT with reinforcement may say something about the nonsensory side of the decision process, and the regularity of the RT distributions and their systematic shifts with reinforcement (see, e.g., Figure 3, top) invite the sort of quantitative analysis that has been useful in the previous research cited above. Such analysis requires one to fit a mathematical function to RT distributions; here the choice of that function was guided by three criteria: (1) the function should have no more than three free parameters, (2) the parameters should represent plausible underlying processes, and (3) changing a single parameter should enable the function to fit data from different experimental conditions.

The ex-Gaussian function has helped to account for previous pigeon RT data (see Blough, 1988, 1989, 1992), and it was the first to be tried here. This function, which is the sum of exponential and Gaussian (“normal”) probability density functions, gave good fits to most of the RT distributions in the present experiments, thus meeting the first criterion. However, it was less successful with the other criteria, for the exponential parameter was not clearly meaningful, and attempts to fit RT distributions under different reinforcement conditions required multiple parameter changes for which there was no particular rationale. In retrospect, these failures are not surprising, for the previous

success of the ex-Gaussian function in accounting for pigeon visual search associated the exponential component with the sensory side of discriminative processing, which presumably would be unaffected by reinforcement changes.

The search for a process whose speed might be affected by reinforcement led next to a consideration of a random walk or diffusion model, which can yield RT functions in a rather natural way (Luce, 1986; Van Zandt, 2002). A simple version of this model can be conveniently fit by the Wald density function. The Wald has only two parameters, it yields appropriately shaped RT distributions, and, most important, it can represent a plausible response generation process. Next I will briefly show how a diffusion model implemented with the Wald function might accommodate the present RT data.

A Diffusion Model

Let us imagine a process by which stimulus onset triggers the growth of some quantity in steps of random size until a threshold is reached, at which point a response occurs (see Figure 11). This scheme is called a diffusion process in the continuous case, in which the discrete steps are replaced by continuous time. If we assume the step size to have a Gaussian distribution, the resulting distribution of threshold crossing times (here, RTs) is described by the Wald density function (Luce, 1986). The two parameters of the Wald are algebraically related to the mean and standard deviation of the Gaussian that governs the speed of accumulation. In order to fit the RT distributions from these experiments, I added a third parameter, which controlled the length of a short delay between stimulus onset and the start of accumulation. This delay could arise from those components of RT, such as motor response time, that are unaffected by the experimental variables. The added delay shifted the Wald density function along the time axis,

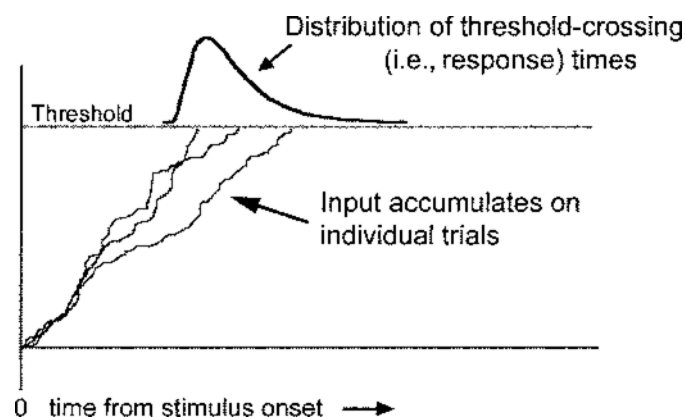


Figure 11. Diagram of a diffusion process that could account for the reaction time distributions collected in these experiments. The distribution of threshold crossing times is described by a Wald density function, as is described in the text and in the Appendix. The data suggest that the rate of accumulation varies as a function of the rate of reinforcement associated with a stimulus.

so I call the result the *shifted Wald function*. The appendix contains details of the function and the fitting process.

The shifted Wald function was fit to the RT distributions from the 2.5%, 5%, and 100% reinforcement conditions of the green target in Experiment 1. First the data from the 2.5% condition were fit by adjustment of all three parameters. Then, to fit the 5% and 100% distributions, only the parameter governing accumulation speed was allowed to change. The results are summarized in Figure 12. The data points are averages of the data distributions across birds; the curves represent the average of the fits of the shifted Wald function to each bird's data. As can be seen, the first fit, to the 2.5% data, was quite close; the mean proportion of variance accounted (ω^2) of the fits for individual birds was 0.97. The fits to the 5% and 100% data, achieved by changing only the accumulation speed, were not quite as good (mean ω^2 of .96 and .94, respectively), but the fitting functions seem to capture the shape and relationships of the distributions reasonably well. The fitting parameters indicate that the average diffusion speed at 5% was 1.35 times as fast as it was at 2.5%, and at 100% it was 1.68 times as fast.

The fits shown in Figure 12 suggest that reinforcement variations may influence response by altering the speed of an accumulation process. It is important to distinguish this conclusion from the idea that reinforcement affects the position of the threshold that is crossed to trigger a response (see Figure 11); in fact, my attempts to fit the distributions by holding accumulation speed constant and changing only the threshold yielded unsatisfactory results. This outcome may be surprising, because, in SDT and similar theories, incentive affects the criterion, which seems analogous to the threshold. Whether or not this analogy is appropriate (which is doubtful), this evidence that reinforcement

affects an accumulation speed and not a threshold is a significant clue to a more complete account. I will next suggest how diffusion might fit into a more complete model.

A Memory Approach

The discussion to this point has left unspecified just what is accumulating in the diffusion model, what the threshold might represent, and why the speed of accumulation might change with variations in reinforcement. These are difficult questions, but tentative answers may be available in the context of existing models. In particular, both the RT data and the diffusion idea appear to be consistent with a memory-instance approach that I and others have suggested to account for a variety of discrimination results (Blough 1998; Chase & Heinemann, 1991; Gibbon & Church, 1984; Heinemann, 1983). These models propose that items carrying information about individual stimulus-response reinforcement events are stored in long-term memory, and response emission is controlled by a process that compares current stimulus input with items retrieved from memory.

One can see that the memory models use a decision process related to that proposed by SDT, the main difference being that information retrieved from memory replaces the criterion against which an incoming sensory representation is compared. My version of the memory model specifies that a certain number of retrieved items are needed for each response decision. Thus, if we assume that retrieval takes time, we may have answers to two of the questions raised above. To the question "What is accumulating?" the model responds, "Items retrieved from memory," and to the question "What determines the threshold?" the model responds, "The number of items required to support a decision."

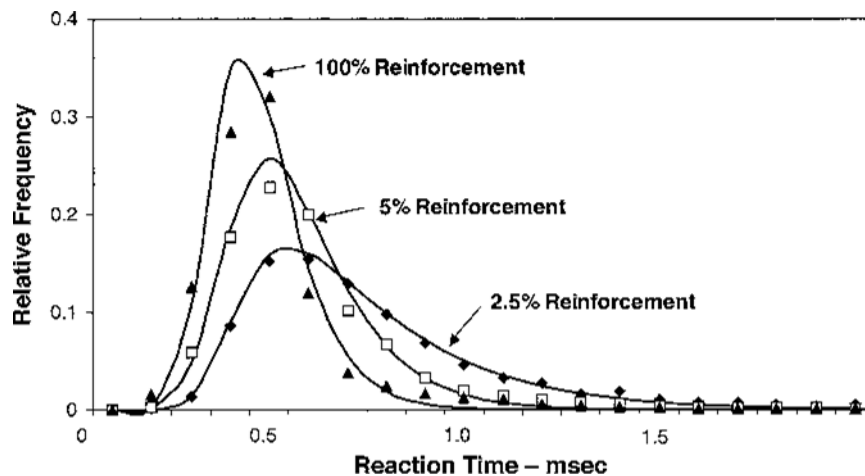


Figure 12. Fits of the shifted Wald density function to reaction time distributions from three reinforcement conditions of Experiment 1. The points represent the average of data across birds; the curves represent the average of least-squares fits to each bird's data. Three parameters were adjusted to fit the 2.5% data; the fits of the other two distributions were derived from that fit by changing a single parameter, which represented the speed of the accumulation shown graphically in Figure 11.

The retrieval process may also be a key to the third question, “Why does accumulation speed change with reinforcement?” What is needed is evidence that the rate of reinforcement affects the speed of retrieval from memory. In fact, Gibbon (1995) proposed just that mechanism to account for a peculiar pattern of response rates recorded in several operant choice experiments. We should note, however, that in hypothesizing an effect of reinforcement on the speed of memory retrieval, Gibbon had in mind an arousal effect that would impact all RTs in a given situation, not just RTs to a stimulus for which reinforcement changes were programmed. To some extent this arousal notion fits the present data, for changes in reinforcement of the green target affected RTs to the red as well as the green targets, and vice versa. However, the largest effect of differential reinforcement was limited to the reinforced stimulus, so if Gibbon’s hypothesis applies it must be modified or supplemented by a stimulus-specific mechanism. This might involve differential storage or retrieval of memory items, but further speculation on this matter would take us far afield.

Various alternatives to the memory approach are available. One example is the stochastic counter model outlined by Killeen and his colleagues (Killeen, Hall, Reilly, & Kettle, 2002). Their “latency machine” generates RT distributions corresponding to the gamma density function, but exploration of this alternative for the present data gave unsatisfactory results: Large changes in both parameters of the gamma were required to fit distributions from different reinforcement conditions, and a conceptual basis for these changes was unclear. Another intriguing possibility is a connection between accumulation models and elemental models of conditioning—in particular, Wagner’s SOP model (Wagner, 1981). Wagner and Brandon (2001) suggested a stochastic process for the activation of a CS node containing many representational units, such that units are probabilistically activated following the onset of a CS and are subsequently deactivated in a similar fashion. If one assumes that a response is initiated when total activation reaches a certain level, this arrangement could generate RT distributions in much the same way as the diffusion model.

Concluding Comment

Theoretical speculation is not to everyone’s taste, and it is surely premature to imagine that a strong case can be made for any of the processes or models suggested above. Close fits by plausible mathematical functions may seem to be compelling evidence, but quite different kinds of models can produce similar sets of RT distributions (e.g., Van Zandt & Ratcliff, 1995). On the other hand, by testing quantitative functions against orderly, detailed data one can eliminate untenable hypotheses and could open directions of inquiry not previously envisioned. For example, it is intriguing that a diffusion process might describe aspects of response initiation in pigeons, given that such processes play a key role in important models of human memory, perception and cognition (e.g., Ratcliff, 1978, 1988; Rat-

cliff & Rouder, 1998; Smith, 1995). Paradoxically, perhaps, the role of diffusion in models of human cognition has more to do with identity than with incentive, involving information as to whether a particular stimulus has been presented or whether the stimulus is represented in memory. As far as I know, no model of human performance involves an accumulation process affected by incentive, but this may be only a symptom of the relative neglect of motivational variables in the human information-processing literature.

At the beginning of this article I suggested that researchers with humans and animals had found common ground in the application of SDT to the analysis of sensory and decision processes. In the discussion above, this common ground was broadened to include the diffusion process and memory models. This is just one of several areas in which quantitative analyses of discriminative processes have facilitated the convergence of models developed with data from different species; the analysis of timing is an even better example (see Church, 2002; Gibbon & Allan, 1984). It will be interesting to see whether such convergence becomes increasingly common.

REFERENCES

- BLOUGH, D. S. (1967). Stimulus generalization as signal detection in pigeons. *Science*, **158**, 940-941.
- BLOUGH, D. S. (1978). Reaction times of pigeons on a wavelength discrimination task. *Journal of the Experimental Analysis of Behavior*, **30**, 133-137.
- BLOUGH, D. S. (1988). Quantitative relations between visual search speed and target-distractor similarity. *Perception & Psychophysics*, **43**, 57-71.
- BLOUGH, D. S. (1989). Contrast as seen in visual search reaction times. *Journal of the Experimental Analysis of Behavior*, **52**, 199-211.
- BLOUGH, D. S. (1992). Effects of stimulus frequency and reinforcement variables on reaction time. *Journal of the Experimental Analysis of Behavior*, **57**, 47-50.
- BLOUGH, D. S. (1996). Error factors in pigeon discrimination and delayed matching. *Journal of Experimental Psychology: Animal Behavior Processes*, **22**, 118-131.
- BLOUGH, D. S. (1998). Context reinforcement degrades discriminative control: A memory approach. *Journal of Experimental Psychology: Animal Behavior Processes*, **24**, 185-199.
- BLOUGH, D. S. (2000). Effects of priming, discriminability, and reinforcement on reaction-time components of pigeon visual search. *Journal of Experimental Psychology: Animal Behavior Processes*, **26**, 50-63.
- BLOUGH, D. S. (2001). Some contributions of signal detection theory to the analysis of stimulus control in animals. *Behavioural Processes*, **54**, 127-136.
- BLOUGH, D. S. (2002). Measuring the search image: Expectation, detection, and recognition in pigeon visual search. *Journal of Experimental Psychology: Animal Behavior Processes*, **28**, 397-405.
- CHASE, S., & HEINEMANN, E. G. (1991). Memory limitations in human and animal signal detection. In M. L. Commons, J. A. Nevin, & M. C. Davison (Eds.), *Signal detection: Mechanisms, models, and applications* (pp. 121-138). Hillsdale, NJ: Erlbaum.
- CHURCH, R. (2002). Temporal learning. In H. Pashler & J. Wixted (Eds.), *Steven’s handbook of experimental psychology: Vol. 3. Learning, motivation and emotion* (3rd ed., pp. 365-393). New York: Wiley.
- DAVISON, M., & NEVIN, J. A. (1999). Stimuli, reinforcers, and behavior: An integration. *Journal of the Experimental Analysis of Behavior*, **71**, 439-482.
- ENNIS, D. M. (1988). Confusable and discriminable stimuli: Comment on Nosofsky (1986) and Shepard (1986). *Journal of Experimental Psychology: General*, **177**, 408-411.

- ESPINOZA-VARAS, B., & WATSON, C. S. (1994). Effects of decision criterion on response latencies of binary decisions. *Perception & Psychophysics*, **55**, 190-203.
- GIBBON, J. (1995). Dynamics of time matching: Arousal makes better seem worse. *Psychonomic Bulletin & Review*, **2**, 208-215.
- GIBBON, J., & ALLAN, J. G. (Eds.) (1984). *Timing and time perception*. New York: New York Academy of Sciences.
- GIBBON, J., & CHURCH, R. M. (1984). Sources of variance in an information processing theory of timing. In H. L. Roitblat, T. G. Bever, & H. S. Terrace (Eds.), *Animal cognition* (pp. 465-488). Hillsdale, NJ: Erlbaum.
- GREEN, D. M., & SWETS, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- HEINEMANN, E. G. (1983). A memory model for decision processes in pigeons. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Discriminative processes* (Vol. 4, pp. 3-19). Cambridge, MA: Ballinger.
- KILLEEN, P. R., & HALL, S. S. (2001). The principal components of response strength. *Journal of the Experimental Analysis of Behavior*, **75**, 111-134.
- KILLEEN, P. R., HALL, S. S., REILLY, M. P., & KETTLE, L. C. (2002). Molecular analyses of the principal components of response strength. *Journal of the Experimental Analysis of Behavior*, **78**, 127-160.
- LUCE, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.
- MACKINTOSH, N. J. (1983). *Conditioning and associative learning*. New York: Oxford University Press.
- MACMILLAN, N. A. (2002). Signal detection theory. In H. Pashler & J. Wixted (Eds.), *Stevens' handbook of experimental psychology: Vol. 4. Methodology in experimental psychology* (3rd ed., pp. 43-90). New York: Wiley.
- NEVIN, J. A. (1965). Decision theory in studies of discrimination in animals. *Science*, **150**, 1057.
- RATCLIFF, R. (1978). A theory of memory retrieval. *Psychological Review*, **85**, 59-108.
- RATCLIFF, R. (1988). Continuous versus discrete information processing: Modeling the accumulation of partial information. *Psychological Review*, **95**, 238-255.
- RATCLIFF, R., & ROUDER, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, **9**, 347-356.
- RATCLIFF, R., & TUERLINCKX, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, **9**, 438-481.
- ROBERTS, S. (1987). Evidence for distinct serial processes in animals: The multiplicative-factors method. *Animal Learning & Behavior*, **15**, 135-173.
- ROBERTS, S., & PASHLER, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, **107**, 358-367.
- SHEPARD, R. N. (1988). Time and distance in generalization and discrimination: Reply to Ennis (1988). *Journal of Experimental Psychology: General*, **117**, 415-416.
- SMITH, P. L. (1995). Psychophysically principled models of visual simple reaction time. *Psychological Review*, **102**, 567-591.
- STERNBERG, S. (1998). Discovering mental processing stages: The method of additive factors. In D. Scarborough & S. Sternberg (Eds.), *Invitation to cognitive science: Vol. 4. Methods, models, and conceptual issues* (pp. 703-863). Cambridge, MA: MIT Press, Bradford Books.
- SPENCE, K. (1937). The differential response in animals to stimuli varying within a single dimension. *Psychological Review*, **44**, 430-444.
- VAN ZANDT, T. (2002). Analysis of response time distributions. In H. Pashler & J. Wixted (Eds.), *Stevens' handbook of experimental psychology: Vol. 4. Methodology in experimental psychology* (3rd ed., pp. 461-516). New York: Wiley.
- VAN ZANDT, T., & RATCLIFF, R. (1995). Statistical mimicking of reaction time data: Single-process models, parameter variability, and mixtures. *Psychonomic Bulletin & Review*, **2**, 20-54.
- WAGNER, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5-47). Hillsdale, NJ: Erlbaum.
- WAGNER, A. R., & BRANDON, S. E. (2001). A componential theory of Pavlovian conditioning. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 23-64). Mahwah, NJ: Erlbaum.

APPENDIX
Curve Fitting: The Wald Distribution

As was noted in the text, the Wald density function can represent an accumulation process that might underlie RT distributions. It is easiest to picture this accumulation as a series of discrete steps that each add an amount to the accumulating total. This discrete form of the Wald density function is as follows (from Luce, 1986, p. 145, with a slight correction in notation):

$$f(n) = (\alpha/2\pi n^3)^{1/2} \exp[-\alpha(n - \beta)^2/2\beta^2 n] \quad (1)$$

In Equation 1, n represents the number of steps in the accumulation process, and $f(n)$ represents the frequency with which the process crosses a response threshold for any given n . The size of each accumulation step varies randomly according to a Gaussian distribution with mean μ and standard deviation σ . These values generate the parameters α and β in Equation 1 according to the relations $\alpha = (C/\sigma)^2$ and $\beta = C/\mu$, where C is a constant representing the value of the threshold. In the continuous Wald density function n is replaced by continuous time t .

It seems certain that some part of the observed RT is contributed by events independent of the accumulation process. This corresponds, in the present analysis, to a delay in the time at which accumulation starts. Doubtless it would be realistic to add this as a another random variable (see Ratcliff & Tuerlinckx, 2002), but for simplicity a constant k is used here. This yields the complete equation used for the fitting function $f(t)$:

$$f(t) = [\alpha/2\pi(t - k)^3]^{1/2} \exp\{-\alpha[(t - k) - \beta]^2/2\beta^2(t - k)\} \quad [f(t) = 0 \text{ for } k \geq t] \quad (2)$$

As was discussed in the text, the RT distributions from one reinforcement condition were fit to the data of each bird by adjusting three parameters, which correspond to a , β , and k in Equation 2. Then distributions from two other conditions were fit by changing only β , with the results shown in Figure 12. Because $\beta = C/\mu$ and C is held constant, changing β is equivalent to changing the mean of five of the Gaussian density function, which specifies the average size of accumulation steps. That is, changing β changed the speed of accumulation.

In making the fits, the value of Equation 2 was computed for $t = 1$ to 20, with each t corresponding to data from a 0.1-sec interval up to 2.0 sec; the small percentage of responses that occurred at longer RTs were discarded, because their values were relatively variable and contributed little to the fits. Fitting was accomplished with the `nlinfit` function of the Matlab computing language; `nlinfit` is a nonlinear least-squares data-fitting algorithm.

To characterize the data in a compact way, average RT distributions are shown in the figures. The distributions from each of the six birds were fit separately; in this way I was able to avoid possible distortions that could arise from fitting an average data function. Because parameters vary somewhat among the subjects, this process can achieve a more realistic fit of the data, but it does not mean that the average functions ultimately displayed necessarily represent a “true” distribution shape.

For information on some of the technical issues involved in curve fitting and the choice of fitting functions see, for example, Luce (1986), Ratcliff and Tuerlinckx (2002), Roberts and Pashler (2000), and Van Zandt (2002).
