TYPES OF RESPONDING IN A SIGNAL-DETECTION TASK

BRENT ALSOP AND RACHAEL ROWLEY

UNIVERSITY OF OTAGO, NEW ZEALAND

Four pigeons were trained to discriminate between two line orientations in a two-alternative forced-choice procedure. The distribution of reinforcers for the two types of correct response was varied across conditions. Performance on each trial was recorded separately, including the time taken to make a choice response. Discriminability and response-bias measures were calculated for overall performance, and, following a median split of the data from each condition, for faster and slower choice responses in each condition. Discriminability between the stimuli did not vary systematically as a function of choice latency. Variations of the reinforcer distributions produced larger response biases for the faster responses than for the slower responses. Responses on trials following reinforcers were faster and showed a greater effect of the reinforcer distribution than did other responses. Behavioral models of signal detection should consider the speed of the choice response as a factor modulating the effects of reinforcer distributions.

Key words: signal detection, stimulus discriminability, response bias, choice latency, pigeon

Nevin (1969) and Davison and Tustin (1978) provided a conceptual and quantitative merging of two hitherto distinct areas of research: signal-detection theory and the free-operant model of choice behavior known as the generalized matching law. The generalized matching law describes a linear relation between the rates of two concurrently available responses, B_1 and B_2 , and their respective rates of reinforcement, R_1 and R_2 , such that

$$\log\left(\frac{B_1}{B_2}\right) = a \log\left(\frac{R_1}{R_2}\right) + \log c. \quad (1)$$

The slope of the line, a , is called the sensitivity parameter because it measures the extent to which changes in the reinforcer distribution (R_1/R_2) produce changes in the allocation of responses (B_1/B_2). Values of a less than 1.0 are the common empirical result, a phenomenon known as undermatching (e.g., Baum, 1974; Taylor & Davison, 1983). The intercept of the line, $\log c$, measures any systematic preference for an alternative that remains unaffected by changes in the reinforcer distribution. This phenomenon is

known as bias, and is often attributed to undetected asymmetries in the organism or apparatus (e.g., Baum, 1974).

Davison and Tustin (1978) extended the generalized matching law to account for performance in two-choice signal-detection procedures. They argued that such tasks can be seen as a complex version of a concurrent schedule. One of two sample stimuli (S_1 or S_2) appears, the subject makes an observing response of some kind, and then the subject chooses between two alternative responses. Correct responses typically produce some payoff. After training on this procedure, the subject will have received some distribution of reinforcers for responses on the two alternatives. Davison and Tustin proposed that this distribution of reinforcers influences behavior in the same manner as the generalized matching law, whereas the effect of the sample stimuli is to bias responding toward correct responses. Following an S_1 presentation, the distribution of choices can be described by the relation

$$\log\left(\frac{B_w}{B_x}\right) = a \log\left(\frac{R_w}{R_x}\right) + \log c + \log d, \quad (2)$$

where B_w and B_x are responses to Alternatives 1 and 2 (correct and incorrect, respectively, in this case), R_w/R_x is the overall reinforcer distribution, and the parameters a and $\log c$ are the same as in Equation 1. The parameter $\log d$ measures discriminability between the sample stimuli, that is, any systematic prefer-

We thank K. G. White for making the equipment available to conduct this study, and the members of his lab for their help running the experiment. Thanks also to Victoria Johnstone, Fraser Wyeth, and Tony Nevin for their thoughtful comments on this article.

Reprints can be obtained from Brent Alsop, Department of Psychology, P.O. Box 56, University of Otago, Dunedin, New Zealand.

ence for the correct response alternative (B_w here) independent of the effects of the reinforcer distribution and inherent biases.

A similar equation describes performance following S_2 presentations:

$$\log\left(\frac{B_y}{B_z}\right) = a \log\left(\frac{R_w}{R_z}\right) + \log c - \log d, \quad (3)$$

where B_y and B_z are responses to Alternatives 1 and 2 (incorrect and correct, respectively, in this case) and all other notation is as above. The discriminability parameter, $\log d$, is subtractive in Equation 3 because the correct response alternative, B_z , is the denominator on the left side of the equation.

Algebraic subtraction of Equations 2 and 3 provides a bias-free measure of discriminability between the two stimuli, and takes the form

$$\frac{1}{2} \log\left(\frac{B_w B_z}{B_x B_y}\right) = \log d, \quad (4)$$

where all notation is as above. In this form, it is clear that $\log d$ is simply the geometric mean of the correct:incorrect response ratios following the two types of stimulus presentation, and is equivalent to the measure of discriminability suggested by Luce's choice theory (1963).

Similarly, algebraic addition of Equations 2 and 3 produces a discriminability-free measure of response bias,

$$\log b = \frac{1}{2} \log\left(\frac{B_w B_y}{B_x B_z}\right) = a \log\left(\frac{R_w}{R_z}\right) + \log c, \quad (5)$$

that incorporates the effects of the payoff or reinforcer distribution, R_w/R_z , and any inherent bias for one alternative over the other, $\log c$. In this form, it is clear that $\log b$ is simply the geometric mean of the ratios of Alternative 1 and Alternative 2 responses, and it is equivalent to the measure of response bias suggested by Luce's choice theory (1963).

Although Davison and Tustin's (1978) model has been successful in describing performance from a variety of signal-detection studies (e.g., McCarthy & Davison, 1979, 1980a, 1980b, 1984), it has encountered a major problem. Davison and Tustin's formulations imply a mutual independence between stimulus discriminability (Equation 4)

and sensitivity to reinforcement (Equation 5). However, several studies have now shown that sensitivity to reinforcement increases as stimulus discriminability decreases (e.g., Alsop & Davison, 1991; Nevin, Cate, & Alsop, 1993; White, 1986). This interaction weakens this model as a theoretical and empirical tool.

Although there have been attempts to reformulate the Davison and Tustin model (e.g., Alsop & Davison, 1991) to maintain the independence of these parameters, such efforts have met limited success. At this point, there are three alternatives for modeling signal-detection performance. The first approach is simply to accept this covariance between these parameters as a natural consequence of signal-detection procedures, and to continue using the current models. The second approach is to use existing signal-detection data to generate alternate models in the hope that they will eventually provide mutually independent bias and discriminability parameters. The third approach is to reexamine some underlying assumptions associated with signal-detection performance. Here we adopt the last approach.

Models such as Davison and Tustin's (1978) describe behavior in terms of the subject's performance in a relatively large number of trials (typically several hundred). However, what does this description imply at a trial-by-trial level? For example, is the behavior of the subject on any given trial under joint stimulus control ($\log d$) and schedule control ($\log b$) in the manner described by the model? Or, is behavior predominantly under stimulus control on some trials and schedule control on others, and the model represents the sum of these trials? There might be a tendency to treat Davison and Tustin's model as if it implied the former, when it is actually silent on this issue. Moreover, there is evidence that discriminability and bias do vary systematically across trials in discrimination tasks, and that this variation seems to depend on the speed with which the subject makes its choice.

Blough (1978) found systematic changes in discriminability as a function of response latency in a go/no-go discrimination task. Pigeons were presented with a series of different wavelengths, but only responses to one wavelength resulted in reinforcement. He found a bimodal distribution of choice latency. The first mode was comprised of re-

sponses with response latencies less than 0.2 s, and for these responses the probability with which a response occurred was independent of wavelength presented. On the other hand, the probability of responses with latencies greater than 0.2 s was proportionally related to the difference between the stimulus presented and the reinforced wavelength. Blough concluded that stimulus onset triggered the quicker responses, and that the particular wavelength of the stimulus influenced only the likelihood of the later responses. Research with human subjects has found similar effects; however, here the results have been more ambiguous. In general, it appears that errors tend to be faster than correct responses when speed is emphasized (e.g., by imposing a deadline for responses) and the discrimination is easy, and that errors tend to be slower than correct responses when accuracy is emphasized and the discrimination is difficult (for review, see Luce, 1986). Furthermore, Lappin (1978) provided evidence that both bias and discriminability vary as a function of choice latency in signal-detection experiments. The human subjects showed near-zero discriminability at shorter choice latencies, but discriminability increased as choice latency increased. However, he found the opposite relation between bias and reaction time; that is, the response bias produced by changes in stimulus presentation probability was greatest at short choice latencies, and it decreased as choice latency increased.

The present study investigated whether there were relations between choice latency, discriminability, and bias in a two-choice signal-detection task when pigeons were used as subjects. If stimulus control and schedule control vary systematically across identifiable types of trial, then this effect might guide future models of signal-detection performance. The subjects were trained to discriminate between two line orientations, and the reinforcer distribution for correct responses (R_w/R_z in Equations 2 and 3) was varied across conditions. The choice latency on each trial was recorded, and discriminability and response bias in each condition were examined as a function of choice latency.

METHOD

Subjects

Four homing pigeons, numbered C1, C3, C4, and C5, were maintained at 85% of their

free-feeding body weights. Water and grit were freely available in their home cages. There was postsession feeding with mixed grain, if necessary, to maintain their body weights. The birds had previous experience in a delayed matching-to-sample experiment, but with different sample stimuli than those used in the present experiment.

Apparatus

The standard pigeon operant conditioning chamber (32 cm wide, 36 cm deep, 29 cm high) contained three circular response keys 3 cm in diameter, 8 cm apart, and 20 cm above the grid floor. A peck of approximately 0.15 N activated a key. The center key could be illuminated with a red light, a line at 90° (S_1 throughout), a line at 60° (S_2 during the experiment), or a line at 30° (S_2 during initial training). The lines were white against a black background and were 0.3 cm in width. The side keys could be lit green. Reinforcement consisted of 2.5-s access to a grain hopper located 15 cm below the center key. A fan provided ventilation and some masking noise. A 386sx IBM®-compatible computer and MEDPC® software and interface controlled experimental events and recorded the data.

Procedure

General procedure. An experimental trial began with a center key lit red and unlit side keys. Following a single response to the center key, one of the line orientations (S_1 or S_2) replaced the red light, and the side keys were lit green. The sample stimuli, S_1 and S_2 , appeared with equal probability. The sample stimulus and the green side keys remained lit until the subject responded to a side key. If S_1 was present on the center key, a response to the left key was correct and occasionally produced a reinforcer. If S_2 was present on the center key, a response to the right key was correct and occasionally produced a reinforcer. Unreinforced correct responses and incorrect responses produced a 2.5-s blackout. There was then a 3.5-s intertrial interval in blackout before the next trial began.

The reinforcers for correct responses were arranged using a controlled reinforcer-ratio procedure (McCarthy & Davison, 1979, 1984). The computer determined every 1 s whether a reinforcer would be made available

with a probability of .05 (i.e., every 20 s on average). This reinforcer was probabilistically assigned for the next correct left or right response. Once that reinforcer was obtained, timing of the next reinforcer interval began. The distribution of reinforcers obtained for correct left- and right-key responses was varied across conditions (1:4, 1:1, 4:1) by varying the probability with which the reinforcers were allocated between the two keys.

Each daily session lasted for either 300 trials or 50 min, whichever occurred first. The computer recorded the data from each trial; that is, the stimulus presented, the response made, whether a reinforcer was obtained, and the time taken to make the choice (measured from sample stimulus onset to a side-key response).

Initial training. No shaping of the key-peck response was necessary because the pigeons had been used in a previous study. The subjects were trained initially with the 90° line as S_1 and the 30° line as S_2 , and correct responses on the left and right keys were reinforced with equal probability. After 30 sessions, the subjects were making few errors. The 60° line then replaced the 30° line as S_2 , and the first experimental condition began.

Experimental Conditions

Condition 1 (33 sessions) arranged equal probabilities of reinforcement for correct responses following S_1 and S_2 presentations. In Condition 2 (41 sessions), the ratio of reinforcers for S_1 and S_2 presentations was 1:4. Condition 3 (69 sessions) was a reversal of Condition 2 with a 4:1 reinforcer ratio for S_1 and S_2 presentations. These three conditions were then replicated in the same order. Condition 1R lasted for 89 sessions, Condition 2R lasted for 80 sessions, and Condition 3R lasted for 47 sessions.

Birds were trained in each condition until the data appeared to be stable. To assess stability, measures of discriminability and bias were calculated for three-session blocks for each subject. When there was no apparent trend in any subject's performance, the condition was changed. The data analysis used the last 10 sessions (approximately 3,000 trials) in each condition (Table 1).

RESULTS

There were two parts to the analysis. In the first part, estimates of discriminability ($\log d$,

Equation 4) and bias ($\log b$, Equation 5) were obtained for each subject's performance in each condition. These measures were calculated using the data from all trials in a condition, and also separately for trials with "fast" and "slow" choice responses within each condition (as determined by a median split on the basis of choice latency). The second part of the analysis focused more directly on the choice-latency data. Mean choice latency and the distribution of choice latencies were compared across conditions, and local changes in choice latency following reinforcement were analyzed within conditions.

Signal-Detection Analysis of All, Fast, and Slow Responses

Table 1 shows the number of left-key and right-key responses following S_1 or S_2 presentations, the number of reinforcers obtained for left-key and right-key responses, and the mean time taken to make the choice (choice latency) over the last 10 sessions in each condition. Estimates of discriminability between the sample stimuli ($\log d$, Equation 4) were calculated using the response data shown in Table 1. These are plotted as a function of the reinforcer ratio in Figure 1. There was some evidence that estimates of discriminability were greater for conditions that arranged unequal reinforcer distributions, similar to the result found in studies of the differential outcome effect (e.g., DeLong & Wasserman, 1981). However, this pattern was not consistent across or within all subjects, and the mean data favor a more conservative interpretation, especially those from Conditions 1R to 3R (Figure 1). The only systematic difference in discriminability across subjects and conditions was found in Condition 3 ($\log R_w/R_z = 0.60$), where all subjects showed markedly higher discriminability than in the first two conditions. This result suggested that performance had improved with extended training across Conditions 1 to 3, but the replications indicated that this was not the case.

Each subject's responding in each condition was then divided into slow and fast groups of trials by performing a median split of the data on the basis of choice latency. Estimates of discriminability between the sample stimuli were calculated separately for the fast and slow sets of trials. These values are also shown in Figure 1. Discriminability did

Table 1

The number of left and right responses following S_1 (B_w and B_x) and S_2 (B_y and B_z) presentations, the number of reinforced correct responses following S_1 and S_2 presentations (R_w and R_z), the mean choice latency in milliseconds, and the number of trials completed over the last 10 sessions are shown for each subject in each experimental condition.

Subject	Condition	B_w	B_x	B_y	B_z	R_w	R_z	Latency	Trials
C1	1	967	563	532	930	272	275	1,510	2,992
	2	878	622	130	1,370	137	531	1,520	3,000
	3	1,419	47	321	1,213	542	124	1,310	3,000
	1R	1,109	363	539	919	289	281	1,250	2,930
	2R	1,138	351	33	1,478	135	551	1,090	3,000
	3R	1,474	49	472	939	559	127	1,180	2,934
C3	1	1,178	341	238	1,243	321	326	1,230	3,000
	2	1,298	288	112	1,302	152	540	1,210	3,000
	3	1,541	9	393	1,057	545	136	1,060	3,000
	1R	1,117	438	63	1,382	302	334	950	3,000
	2R	867	608	28	1,497	128	545	1,040	3,000
	3R	1,516	18	59	1,407	533	113	1,000	3,000
C4	1	1,067	415	193	1,325	326	277	960	3,000
	2	384	1,137	82	1,397	110	436	830	3,000
	3	1,476	13	16	1,495	550	132	880	3,000
	1R	1,305	125	43	1,385	321	357	740	2,858
	2R	262	1,210	115	1,413	92	396	770	3,000
	3R	1,285	178	1,116	421	434	109	590	3,000
C5	1	1,258	263	193	1,286	334	323	1,190	3,000
	2	1,295	245	48	1,412	112	533	1,100	3,000
	3	1,536	14	26	1,424	569	130	1,090	3,000
	1R	1,305	153	42	1,500	337	315	1,080	3,000
	2R	1,090	403	21	1,486	118	543	1,120	3,000
	3R	1,481	3	46	1,470	520	138	1,060	3,000

not vary systematically as a function of speed of response, as can be seen most clearly in the plot of the mean data. It should be noted that there were three instances in which this subdivision of the data produced an error cell with no responses (i.e., B_x or $B_y = 0$), and Equations 2 to 5 could not be solved. In this case, the problem was overcome by replacing the 0 with 1 in order to perform the calculations (Subject C5, fast responses in Conditions 2, 3, and 3R). This correction produced measures of discriminability in keeping with those obtained in other conditions and realistic measures of bias in later analyses.

The analysis of response bias was performed in a manner similar to that of the discriminability analysis. First, estimates of response bias ($\log b$, Equation 5) were calculated using all trials for each subject in each condition (Table 1). Response bias changed as a function of the reinforcer ratio in the manner predicted by behavioral models of signal detection (Figure 2); that is, there was a response bias for the right key

(i.e., $\log b < 0$) when more reinforcers were obtained for correct S_2 identifications (i.e., $\log R_w/R_z = -0.60$), there was a response bias for the left key (i.e., $\log b > 0$) when more reinforcers were obtained for correct S_1 identifications (i.e., $\log R_w/R_z = 0.60$), and intermediate measures of response bias were found when equal numbers of reinforcers were arranged on each key. As the mean performance indicates, the initial conditions and the replication conditions produced similar results. Only Subject C3 showed any systematic change in performance across the first and second halves of the experiment.

There was further evidence that the estimates of response bias calculated across all trials were consistent with the behavioral model of signal detection. Equation 5 predicts that changes in the reinforcer ratio should change response bias in a manner similar to that typically obtained in concurrent-schedule procedures. The results of fitting a linear regression to each subject's estimates of response bias are shown in Table

Fig. 1. The measures of discriminability ($\log d$) derived from all responses, fast responses only, and slow responses only are plotted as a function the log reinforcer ratio for each subject. The left panels show the results from Conditions 1 to 3, and the right panels show the results from Conditions 1R to 3R (the replication conditions).

Fig. 2. The measures of response bias ($\log b$) derived from all responses, fast responses only, and slow responses only are plotted as a function the log reinforcer ratio for each subject. The left panels show the results from Conditions 1 to 3, and the right panels show the results from Conditions 1R to 3R.

Table 2

The results of linear regressions fit to the response bias data shown in Figure 2. The slope of the fitted line (reinforcer sensitivity), the intercept, and a measure of the goodness of fit (R^2) are shown for each subject in each set of conditions.

Subject	Condition	Reinforcer sensitivity			Intercept			R^2		
		All	Fast	Slow	All	Fast	Slow	All	Fast	Slow
C1	1-3	0.74	1.07	0.48	0.00	0.19	-0.17	1.00	.96	.94
	1R-3R	0.97	1.96	0.23	0.05	0.14	-0.09	.99	.86	.24
C3	1-3	0.93	1.57	0.43	0.20	-0.24	0.55	.83	.87	.78
	1R-3R	0.88	1.02	0.77	-0.33	-0.72	-0.20	.95	.98	.88
C4	1-3	0.74	0.76	0.67	-0.34	-0.34	-0.40	.94	.68	.97
	1R-3R	1.27	1.73	0.97	-0.16	-0.03	-0.23	.99	.97	1.00
C5	1-3	0.43	1.82	-0.30	-0.10	-0.53	0.21	.99	.98	.40
	1R-3R	1.08	1.62	0.63	-0.14	-0.34	-0.08	.95	.99	.86
M	1-3	0.82	1.23	0.46	-0.01	-0.10	0.06	.99	.98	1.00
	1R-3R	0.97	1.48	0.54	-0.10	-0.06	-0.17	1.00	.98	.86

2, but note that this analysis warrants some caution because only three data points contribute to each fit. The effect of the changes in the reinforcer ratio (sensitivity to reinforcement) ranged from 0.43 to 1.27 across subjects, with a median of 0.9. Overall, these values are in the range expected from concurrent-schedule procedures (e.g., Davison & McCarthy, 1988).

There was a systematic difference between the estimates of response bias obtained from the fast and slow groups of trials. Figure 2 shows that the effect of arranging unequal reinforcer ratios was typically more extreme on the fast trials than on the slow trials. This interpretation was also supported by the linear regression analyses of these data (Table 2). For the fast responses, estimates of sensitivity to changes in the reinforcer ratio ranged

from 0.76 to 1.96 across subjects, with a median of 1.60. For the slow responses, these estimates ranged from -0.30 to 0.97 across subjects, with a median of 0.55. Furthermore, the sensitivity to reinforcement for fast responses was always greater than the corresponding estimate obtained from the slow responses.

Choice Latencies

Figure 3 and Table 1 show the mean choice latencies across subjects in each condition. Choice latencies decreased across Conditions 1 to 3 for each subject, but were more consistent across Conditions 1R to 3R. Only Subject C4 appeared to show a continued decrease in choice latency over the entire experiment (Table 1, Figure 3).

The choice-latency data were then collapsed into successive bins of 0.1 s, and frequency distributions were examined for each subject in each condition. This allowed a comparison of the shape of the frequency distributions in the present study with the bimodal distributions obtained by Blough (1978). The shape of the frequency distributions changed across conditions, but these changes seemed to be more a function of the general decrease in mean choice latencies (Figure 3, Table 1) rather than the effects of the arranged reinforcer ratios in specific conditions. Therefore to simplify the presentation of results, Figure 4 shows only the mean frequency distributions across Conditions 1 to 3 and across Conditions 1R to 3R for each subject. These mean distributions were rea-

Fig. 3. The mean time to make a choice response following stimulus onset are plotted for each subject in each condition.

Fig. 4. The mean choice-latency frequency distributions (in 0.1-s bins) across Conditions 1 to 3 (left panels) and Conditions 1R to 3R (right panels) are plotted for each subject.

sonably representative of each subject's performance in these groups of conditions. The results of the frequency distribution analysis of choice latencies were somewhat mixed. Subject C4 produced a unimodal distribution of choice latencies throughout, and

the peak of the distributions became more pronounced as the experiment progressed (Figure 4). On the other hand, Subject C5's frequency distributions showed evidence of bimodality throughout the study. This subject's frequency distribution in each condi-

tion is shown separately in Figure 5 to demonstrate that the bimodality was not simply a by-product of combining performances across groups of conditions. Subjects C1 and C3 produced bimodal distributions of choice latencies in early conditions (Conditions 1 to 3, Figure 4), but their later frequency distributions were more unimodal (Conditions 1R to 3R, Figure 4). This change seemed to be related to the overall speed of response (Figure 3, Table 1), with the locus of the second peak merging with the first peak as the experiment progressed. It is worth noting that for Conditions 1 to 3, the first mode of the bimodal distributions accounts for most of the group of faster responses in the preceding analysis following the median split.

Sequential Effects Within Sessions

The final set of analyses focused on sequences of trials following reinforcers. This analysis was conducted on the data from each of the three replication conditions because the performance was more consistent across these conditions (e.g., Figures 1, 2, and 3). First, sequences of trials were sorted following R_w and R_z reinforcers. These trials were grouped as a function of their position since the last reinforced trial (trial n); that is, trials immediately following a reinforcer formed Group 1 (i.e., $n + 1$), the next two trials formed Group 2-3 (i.e., $n + 2$ and $n + 3$), and all later trials formed Group >3 . Other divisions of the trials following reinforcement were analyzed (e.g., greater number of groups), but these divisions did not yield substantive differences from the results shown below, and these analyses were more frequently marred by incalculable bias and discriminability measures (i.e., zero error responses in some groups).

Figure 6 plots the mean choice latencies for each of the three groups of trials in Conditions 1R to 3R. The mean choice latencies for trials immediately following a reinforcer (Group 1) were typically faster than those of the trials later in the sequence (Groups 2-3 and >3), but there was no consistent difference between the mean choice latencies of the two latter groups. Furthermore, there were no systematic differences in the results from sequences following R_w and R_z reinforcers in each condition. In other words, reinforcement decreased choice latencies on tri-

Fig. 5. The choice-latency frequency distributions (in 0.1-s bins) across Conditions 1 to 3R are plotted for Subject C5.

Fig. 6. The mean choice latency following S₁ reinforcers (left panels) and S₂ reinforcers (right panels) for trials immediately following a reinforcer (Group 1), two or three trials since the last reinforcer (Group 2-3), and more than three trials since the last reinforcer (Group >3) for each subject in Conditions 1R to 3R.

als immediately following a reinforcer, irrespective of the source of that reinforcer.

Figure 7 plots the estimates of response bias for Groups 1, 2-3, and >3 in the two con-

ditions that arranged unequal reinforcer distributions. There are missing data points for some subjects in Figure 7 because this division of the data into groups produced several

Fig. 7. The estimates of response bias for choice responses following S_1 reinforcers (left panels) and S_2 reinforcers (right panels) for trials immediately following a reinforcer (Group 1), two or three trials since the last reinforcer (Group 2-3), and more than three trials since the last reinforcer (Group >3) for each subject in Conditions 1R to 3R.

instances of subjects making no errors of a particular type. Despite this shortcoming, the remaining data show a reasonably clear pattern; that is, response bias toward the more reinforced alternative was typically more extreme in Group 1 trials than in Group 2-3 and Group >3 trials. Of the 12 instances in which

bias was calculable for Group 1, this measure was the most extreme on 11 occasions. This was the case irrespective of which reinforcer (i.e., following R_w or R_z) had been obtained; in other words, response bias for the alternative with the higher rate of reinforcement was greater following a reinforcer regardless of which alternative produced that reinforcer. There was no systematic difference between measures of response bias from Groups 2-3 and >3.

The corresponding analysis of discriminability measures was unrewarding, and the results are not shown. There were no systematic changes in discriminability across the three groups of trials, or following R_w or R_z reinforcers. This result was consistent with the previous analyses (e.g., Figure 1) that also showed no changes in discriminability within the session.

These analyses offer a possible explanation of the results of the fast and slow responses presented earlier (Figure 2, Table 2). Faster responses showed greater effects of the reinforcer ratio because responses immediately following reinforcers would be disproportionately represented in the fast response group (Figure 6), and these responses also show more extreme measures of response bias (Figure 7). This would be an erroneous conclusion, however. This hypothesis was tested by removing the Group 1 trials from the data set and analyzing the combined trials from Groups 2-3 and >3. These trials were ordered according to choice latency and divided into fast and slow trials by a median split for each of the two conditions that arranged unequal reinforcer ratios (Conditions 2R and 3R). Figure 8 shows the measures of response bias obtained for each subject in each condition for these fast and slow responses. For six of seven comparisons (the bias measure was in-calculable for the fast responses of C2), response bias was more extreme for the fast responses. Therefore, even when the effect of reinforcers on the immediately subsequent trial was removed, the relation between fast and slow responses and response bias was still maintained.

On the other hand, there was evidence that the effects of fast responding and reinforcer presentations on response bias were additive; that is, measures of response bias from the trials immediately following a reinforcer

Fig. 8. The measures of response bias ($\log b$) are plotted separately for the fast and slow responses as a function of the log reinforcer ratio for each subject in Conditions 2R and 3R. This analysis excluded responses on trials immediately following reinforcers.

(Group 1) were greater than those obtained from other trials (Groups 2-3 and >3) with similar choice latencies. This was assessed by calculating estimates of response bias for those trials with choice latencies less than 1 s for the data from Group 1 trials and from Groups 2-3 and >3 trials combined. Of the five comparisons that could be made without zeros on error cells, response biases were greater for the Group 1 responses on all five occasions (Table 3).

DISCUSSION

At a molar level, the results of the present study were similar to those found in previous studies (e.g., Alsop & Davison, 1991; McCarthy & Davison, 1979). Changes in the ratio of obtained reinforcers produced orderly changes in response bias, but there were no systematic changes in discriminability be-

Table 3

Estimates of response bias ($\log b$) are shown for each subject in Conditions 2R and 3R calculated for those trials with choice latencies less than 1 s only, and for Group 1 trials and for Groups 2-3 and >3 trials separately.

Subject	Reinforcer ratio	Group 1 trials	Groups 2-3 and >3 trials
C1	1:4	-1.70	-1.18
	4:1	— ^a	1.38
C3	1:4	-1.34	-0.99
	4:1	— ^a	— ^a
C4	1:4	-1.43	-0.81
	4:1	0.95	0.73
C5	1:4	-1.14	-0.81
	4:1	— ^a	— ^a

^a Estimates could not be calculated due to zero responses in error cells.

tween the sample stimuli. However, subdividing each condition's performance on the basis of choice latency showed that there were local differences in control by the reinforcer distributions; that is, faster choice responses in the session showed a greater effect of the unequal reinforcer distributions (R_w/R_z) than did slower responses. Furthermore, choice responses on trials following reinforcement were generally quicker than responses on other trials, and they were associated with more extreme estimates of response bias. However, there was still a greater effect of the unequal reinforcer ratios on trials with faster response times even when trials immediately following a reinforcer were removed from the analysis, and there was evidence that the effect of a reinforcer on an immediately previous trial added to this effect.

There were no corresponding systematic changes in discriminability ($\log d$) across any of the within-session analyses in the present study (e.g., Figure 1). This finding was somewhat unexpected on the basis of other research (Blough, 1978; Lappin, 1978) that has indicated that discriminability can change as a function of choice latency. It is unclear why we obtained our result. Although there are various differences in procedure or in the species of subject between the present study and previous research, there seems to be nothing particularly unusual about the preparation used here. One possibility is that there was a ceiling effect; that is, the stimuli were too easily discriminated to show any changes in discriminability as a function of

choice latency. This explanation seems unlikely, however. For example, although the levels of discriminability for some subjects were high, the average discriminability shown by other subjects (e.g., C1 and C2) was well below that which can be obtained in this type of procedure (e.g., Alsop & Davison, 1991).

The absence of a relation between choice latency and discriminability presents problems for any simple explanation of the relation between choice latency and response bias. For example, the changes in response bias cannot be accounted for as a by-product of "fast guess" or "speed-accuracy tradeoff" phenomena. In each case, discriminability should be lower for faster responses than for slower responses. Similarly, the discriminability results forestall attempts to link these findings to recent animal memory studies (Jones & White, 1992; McCarthy & Davison, 1991) that have found covariations in discriminability and response bias when choice is delayed. It also should be noted that size of the differences between the choice latencies here (tenths of seconds) and the delay of choice periods in animal memory research (several seconds) also weakens any such comparison.

One obvious question concerning the relation between response bias and choice latency is the manner in which the reinforcer ratio affected response bias across trials within a session. Was there a discrete group of fast-response trials that the reinforcer distribution affected with equal strength and an overlapping group of slow-response trials that showed little effect of the reinforcer distribution?¹ Or, did the reinforcer distribution affect performance on all trials, but the strength of this effect progressively decreased as choice latency increased? The present experiment does not provide a clear answer to this question, but we favor the former idea; that is, there might be two relatively discrete types of responding within a session. The choice-latency distributions show some, albeit weak, evidence of bimodality, with the first peak corresponding to those trials that showed the greatest effect of the reinforcer distribution and the second peak correspond-

ing to those trials that showed little effect of the reinforcer distribution. We acknowledge that this evidence is equivocal. One subject consistently produced a unimodal choice-latency distribution, and bimodality was less evident in choice-latency distributions during the set of replication conditions. We suggest that the two underlying distributions might still be present in these cases, however, and that the means were simply too close to allow a clear separation of the two distributions. This would be consistent with the overall decrease in choice latency in the replication conditions.

A number of empirical and theoretical issues arise from the present study, irrespective of whether the variation in response bias represents continuous or discrete changes in behavior. First, local differences in the effects of reinforcer distributions could contribute to the interaction between discriminability and response bias shown in some previous studies (e.g., Alsop & Davison, 1991). Consider the following hypothetical example: Conditions in which two sample stimuli were indiscriminable might promote a relatively larger number of fast responses than conditions in which the stimuli were moderately discriminable. In this scenario, changes in the reinforcer ratio (R_w/R_c) would produce more extreme response biases in conditions with indiscriminable stimuli than in conditions with moderately discriminable stimuli. Although such an argument cannot account for the results of all studies in this area (Nevin et al., 1992), a systematic investigation of the effects of variations in stimulus discriminability on both bias and choice latency might shed light on some of the difficulties found in current behavioral models of signal-detection performance.

A second line of inquiry concerns the effects of constraining the time in which the subject makes a choice response. In other words, can we control the extent to which the reinforcers influence performance by allowing only faster or slower responses? Or, for example, would delaying the opportunity to respond simply delay the choice-latency distributions, and their correlated bias effects, without changing the basic pattern of results? Such research will be important in order to understand what factors produce the type of relation between choice latency and response

¹ In this scenario, the effect of the reinforcer ratio on slow responses shown in Figures 4 and 5 is seen as the product of the overlap between the two underlying distributions (i.e., the tail of the faster distribution).

bias found here. Furthermore, there are practical grounds for wanting such information. Signal-detection paradigms are common tools in animal and human psychophysics, and variants of the procedure are used in areas such as animal memory research. These experiments are often looking for designs that will minimize the effects of factors such as reinforcer distributions, and the results of the present study point to ways in which this might be achieved.

More generally, moving the level of analysis of signal-detection performance from response totals to an analysis that focuses on different classes of responses within sessions seems to be a useful development in this area. The breakdown of session performance by choice latency produced orderly and systematic changes in behavior that points directly to directions for future research. Similar approaches might provide the key to a better understanding of behavior on these tasks.

REFERENCES

- Alsop, B., & Davison, M. C. (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.
- 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. (1978). Reaction times of pigeons on a wavelength discrimination task. *Journal of the Experimental Analysis of Behavior*, 30, 163–167.
- Davison, M. C., & McCarthy, D. C. (1988). *The matching law: A research review*. Hillsdale, NJ: Erlbaum.
- Davison, M. C., & Tustin, R. D. (1978). The relation between the generalized matching law and signal-detection theory. *Journal of the Experimental Analysis of Behavior*, 29, 331–336.
- DeLong, R. E., & Wasserman, E. A. (1981). Effects of differential reinforcement expectancies on successive matching-to-sample performance in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 7, 394–412.
- Jones, B. M., & White, K. G. (1992). Sample-stimulus discriminability and sensitivity to reinforcement in delayed matching to sample. *Journal of the Experimental Analysis of Behavior*, 58, 159–172.
- Lappin, J. S. (1978). The relativity of choice behavior and the effect of prior knowledge on the speed and accuracy of recognition. In N. S. Castellan, Jr. & F. Restle (Eds.), *Cognitive theory* (Vol. 3, pp. 139–168). Hillsdale, NJ: Erlbaum.
- 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. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.
- McCarthy, D. C., & Davison, M. C. (1979). Signal probability, reinforcement and signal detection. *Journal of the Experimental Analysis of Behavior*, 32, 373–386.
- McCarthy, D. C., & Davison, M. C. (1980a). Independence of sensitivity to relative reinforcement rate and discriminability in signal detection. *Journal of the Experimental Analysis of Behavior*, 34, 273–284.
- McCarthy, D. C., & Davison, M. C. (1980b). On the discriminability of stimulus duration. *Journal of the Experimental Analysis of Behavior*, 33, 187–211.
- McCarthy, D. C., & Davison, M. C. (1984). Isobias and alloibias functions in animal psychophysics. *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 390–409.
- McCarthy, D. C., & Davison, M. (1991). The interaction between stimulus and reinforcer control on remembering. *Journal of the Experimental Analysis of Behavior*, 56, 51–66.
- 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., 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.
- Taylor, R., & Davison, M. C. (1983). Sensitivity to reinforcement in concurrent arithmetic and exponential schedules. *Journal of the Experimental Analysis of Behavior*, 39, 75–91.
- 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.

Received March 21, 1995

Final acceptance January 8, 1996