

*DISCRIMINABILITY OF FIXED-RATIO SCHEDULES
FOR PIGEONS: EFFECTS OF PAYOFF VALUES*

SALLY L. HOBSON¹

COLUMBIA UNIVERSITY

Three pigeons, previously trained to discriminate different numbers of responses (fixed ratios), were tested under different reinforcement contingencies (payoff matrices) at two levels of sensitivity. For one subject, relative reinforcement magnitude was varied—at first, across sessions and then, at midsession by reversing values—without exteroceptive cues. For another, relative reinforcement magnitude and/or probability was varied every 50 trials with cues by correlating different payoff matrices with different key colors. For the third subject, relative reinforcement probability was varied more frequently with cues—in the limit, at random—to demonstrate stimulus control of response bias on a trial-by-trial basis. A signal-detection analysis showed that bias changed with payoffs, for as many as seven different matrices, while sensitivity remained unchanged. The obtained functions (receiver operating characteristics) were similar under different payoff conditions, which suggests that a single mechanism controls bias. However, they differed enough in slope to require a relatively complex account (e.g., the general Gaussian model of detection theory).

Key words: psychophysics, fixed-ratio schedules of reinforcement, reinforcement access, signal-detection theory, key peck, pigeon

Signal-detection theory was originally developed to provide a relatively pure measure of sensitivity in vision in order to partial out response bias effects (*i.e.*, any tendency to report one stimulus more frequently than another) among human observers. A general model of decision-making for several areas of perception and cognition has evolved along with some specialized applications (see Swets, 1973). Usually, a relatively difficult task is examined. For example, subjects may be asked to distinguish between two similar events (S_1 and S_2), using two responses (R_1 and R_2), with different payoffs assigned to the possible decision outcomes (e.g., reinforcement for R_1/S_1 and punishment for R_2/S_1).

In animal psychophysics, several straightforward applications of detection methods have appeared. The approach has been used

to supplement classical psychophysical measures when thresholds are determined. For example, Irwin and Terman (1970) examined the types of response bias that developed in an auditory detection task, using rats. In addition, a number of researchers (e.g., Nevin, Olsson, Mandell, and Yarensky, 1975; Wright, 1972) have shown that a sensitivity estimate (d'), derived from detection theory varies with the physical characteristics of a stimulus in other situations. Moreover, with stimulus conditions held constant and motivational factors varied, systematic changes in response bias have been obtained without changes in sensitivity (also, Clopton, 1972; Elsmore, 1972; Hume, 1974; Hume and Irwin, 1974; Nevin, 1970; Stubbs, 1976; Terman and Terman, 1972). The techniques involved variations in either stimulus probabilities or some aspect of the payoff situation—primarily, reinforcement probabilities—with changes occurring between sessions after long-term training at each set of values. In work with human subjects, changes are typically programmed in this way, but with regular reinforcement such as money. As for within-session changes, the technique Kinchla, Townsend, Yellott and Atkinson (1966) developed for varying stimulus probabilities with cues has not been widely used.

¹This paper is based on Experiment 2 of a dissertation submitted to the Faculty of Pure Science of Columbia University. The research was supported in part by a predoctoral fellowship from the National Institutes of Health. The author wishes to thank John A. Nevin for his advice and encouragement and Pamela G. Real for helping analyze the results shown in Figure 5. Reprints may be obtained from the author, Department of Psychology, Adelphi University, Garden City, New York 11530.

In the animal laboratory, Nevin's (1970) cyclical reinforcement schedule represents the closest parallel: like Kinchla's technique, it manipulates bias across trials within individual sessions (see also Stubbs, 1968). In any case, the empirical receiver-operating-characteristic curves are remarkably similar to those obtained from human subjects. All told, more than 20 studies using animals are now in print.

However, most of the research has been confined to discrimination tasks involving visual or auditory cues or differences in stimulus duration. Discriminability of different numbers of responses has received less attention, despite Rilling and McDiarmid's (1965) original work along this dimension. The small number of studies that followed (Hobson, 1975; Pliskoff and Goldiamond, 1966; Rilling, 1967) emphasized stimulus factors. In each case, a single-stimulus ("yes-no") discrimination task was used: on a given trial, completion of the programmed ratio of responses (FR) was followed by a choice response, which, if correct, was reinforced with grain. The most recent study (Hobson, 1975), which formed the basis for the present one, showed that changes in discriminability were similar to those reported for better-known dimensions. In particular, with payoffs symmetric and ratio sizes varied, the Weber fraction was found to decrease at least up to FR 30, and beyond, when Rilling's results were considered. Signal-detection methods were used to obtain a fine-grained analysis of choice behavior. The results appeared consistent with the "no-threshold" position of detection theory. Consequently, it was assumed that response bias could be similarly controlled once payoffs were varied.

Of course, in the human laboratory, elegant demonstrations of this kind already exist. For example, in audition, Galanter and Holman (1967) showed that the receiver operating characteristics for individual subjects remain invariant across *different* biasing methods, including not only variations in stimulus probabilities and verbal instructions, but also, a five-fold change in payoff values. A comparable demonstration has not been reported using animals. In fact, series of studies on response bias are still relatively rare within the field, though not in related areas—*i.e.*, in free-operant situations, where choice has been parametrically studied, using concurrent reinforcement schedules. For example, Todorov (1973)

varied reinforcement probabilities and magnitudes, both singly and in combination. Taken together, these results suggest the use of several types of payoff variables, including variants of the common methods, to extend the detection approach to fixed-ratio tasks and psychophysics in general.

Specifically, the purpose of the present study was to obtain at different ratio sizes, using a within-subjects design, systematic changes in response bias under a variety of payoff conditions without affecting sensitivity. In the simplest case, changes in reinforcement magnitudes were scheduled after several sessions at one set of values, much as Nevin *et al.* (1975) or Stubbs (1976) had done with reinforcement probabilities. Next, mid-session changes were studied. Then, to obtain more frequent changes, two subjects were trained from the beginning under more complex, within-session procedures, using exteroceptive cueing. Both reinforcement probabilities and magnitudes were varied. At first, different colors were correlated with different payoffs across blocks of trials, so that each appeared several times a session. Structurally, the procedure is not unlike the verbal-instruction methods of human psychophysics, where subjects are requested to adopt different biases at different times. However, in the limit, the procedure reduces to a probabilistic payoff schedule, with cues. In the present case, the limit was gradually approached. Thus, what evolved was a variant of Kinchla's random technique to provide, along with the other procedures, a strong demonstration of the isolability of bias effects, which is basic to detection theory.

METHOD

Subjects

Three White Carneaux pigeons from an earlier study on fixed-ratio discriminability served (see Hobson, 1975). Because long test sessions were anticipated, the subjects had been switched to probabilistic reinforcement before the present research began. Previous weight levels were maintained.²

²Birds 363 and 366 were maintained within 3% of the usual criterion (about 80% of their free-feeding weights) despite daily testing. Bird 323 was tested in long, overnight sessions under a more flexible criterion (no upper weight limit), with no signs of satiation.

Apparatus

The same three-key pigeon chambers were used here as in the earlier study, with one modification: for Birds 363 and 323, the center-keylight matrix was adapted to project different colors across trials, depending on payoffs—*i.e.*, to provide cueing. The side keys, when lit, were white under all conditions. Experimental contingencies were controlled by electromechanical programming equipment plus a tape reader, all located in another room.

Procedure

Except for differences in payoff conditions, the task facing each subject was basically the same. When the center key was lit, one of two different sized ratios was scheduled at random, with one restriction: run lengths over five could not occur. Completion of a ratio darkened the center key and illuminated the two side keys; a peck on either darkened both, with consequences depending on ratio size. On small-ratio trials (*e.g.*, after FR 6), a left-key peck was considered "correct" and either occasionally produced grain, or else lengthened the intertrial interval by an equivalent period; a right-key peck was "incorrect" and always turned off the houselight for 3 sec, which distinguished it from other possible outcomes. Contingencies were reversed on large-ratio trials (*e.g.*, after FR 10). The next trial did not begin until at least 11 sec had elapsed without center-key responding. On the average, reinforcements were arranged on about half the trials. The access time to grain averaged 1.5 to 2.5 sec per trial, depending on the subject. Table 1 shows the specific payoff matrix for correct responses, session length, *etc.*

Bird 366 was first tested at FR 14 *versus* FR 20, under four payoff matrices that varied in magnitude across sessions, without exteroceptive cueing (*i.e.*, with the center key white regardless of payoffs). For example, in Condition 1-1, when payoffs favored detections of the smaller ratio by a 3:1 margin, 3 sec of grain was scheduled for its detection and only 1 sec for the larger ratio. The subject was also tested at FR 16, using a mid-session reversal of matrix values, without cueing (see Condition 2, Table 1).

The other two subjects were tested at different ratio sizes, using relatively frequent, within-session changes in payoffs and cueing.

When payoffs favored detections of the smaller ratio, the center key was green during either ratio. When payoffs were reversed, it was red. The white key was retained to signal the symmetrical condition. For the most part, Bird 363 was tested at FR 7 *versus* FR 10, with payoffs that covered about the same range of values as Bird 366 had encountered (see Conditions 3 to 7, Table 1). Initially, Bird 323 was also tested in this fashion: specifically, changes in payoffs occurred with cues once every 50 trials, in permutations of three (compare Conditions 3 and 8A in the table). Then, payoff changes were scheduled more frequently, with periods of random alternation included in each session. Finally, the subject was switched to a new level (FR 26 *versus* FR 30) and tested there under the original set of matrices, with one further modification: in Condition 9B, payoffs alternated *entirely* at random.

In general, a subject was tested for several weeks in each condition, until its data appeared asymptotic. In the end (Conditions 2, 6, and 9B), however, only 10 to 12 sessions were scheduled, because the subjects were well-trained and already had been extensively tested in other, similar conditions. In each case, the results are based on a large number of trials—more than 300 per session in most conditions and at least 140 trials per matrix in every session (see Table 1). To examine performance in signal-detection terms, conditional response probabilities were estimated for each matrix at a given level (*i.e.*, at a particular ratio difference), following an earlier convention (see Hobson, 1975)—namely, right-key pecks given the larger ratio were termed "hits" and right-key pecks given the smaller ratio, "false alarms".

RESULTS

Figure 1 shows the receiver operating characteristics (ROCs) obtained from Bird 366 at two levels (FR 14 and FR 16 *versus* FR 20) by varying relative reinforcement magnitudes without cueing (Conditions 1 and 2). The data have been plotted on normalized coordinates, with straight lines fit by eye, as is customary for other dimensions (see Green and Swets, 1966). Consider first the subject's performance at FR 14, where it had been trained under each matrix for weeks at a time. Separate sym-

Table 1
Summary of Test Conditions for Each Subject

Condition ^a	Payoffs (relative values)				Color Cue	Matrix Change in Session	Trials per Session	Reinf. Magn. (sec) ^b	Ratios Tested	
	Reinf. Prob.		Reinf. Magn.							
	Sm FR	Lg FR	Sm FR	Lg FR						
BIRD 366										
1-1 (3)			0.75	0.25			300	2.0		
1-2 (2)	equal ^c		0.50	0.50	none	none	300	2.0	FR 14	
1-3 (4)			0.25	0.75			300	2.0	vs. 20	
1-4 (1)			0.22	0.78			300	2.5		
2 (5)	equal ^c		0.75	0.25	none	at mid-session	400	2.0	FR 16	
			0.25	0.75					vs. 20	
BIRD 363										
3 (1)	0.70	0.30	equal		green	every	50	450	2.5	FR 6
	0.50	0.50								
	0.30	0.70								
4 (2)	0.75	0.25	equal		green	every	50	450	2.0	FR 7
	0.50	0.50								
	0.25	0.75								
5 (4)	equal		0.75	0.25	green	every	50	450	2.0	FR 7
			0.50	0.50						
			0.25	0.75						
6 (5)	equal		0.81	0.19	green	every	50	450	2.0	FR 7
			0.50	0.50						
			0.19	0.81						
7 (3)	0.75	0.25	equal		green	every	50	450	1.5	FR 7
	0.50	0.50								
	0.25	0.75								
			0.33	0.67	white	50			2.0	FR 7
			0.33	0.67	red	trials			1.5	vs. 10
BIRD 323										
8A (1)	0.70	0.30	equal		green	every	50	600	2.5	FR 24
	0.50	0.50								
	0.30	0.70								
8B (2)	payoff values and cues same as 8A					every 25	600	2.5	FR 24	vs. 30
8C (3)	payoff values and cues same as 8A					25/random	600	2.5	FR 24	vs. 30
8D (4)	0.80	0.20	equal		green	20/random	560	2.5	FR 24	vs. 30
	0.20	0.80								
9A (5)	0.70	0.30	equal		green	every	50	600	2.5	FR 26
	0.50	0.50								
	0.30	0.70								
9B (6)	0.80	0.20	equal		green	random	560	2.5	FR 26	vs. 30
	0.20	0.80								

^aTest order is in parentheses.

^bAverage access-to-reinforcement time, per trial.

^cReinforcement probability = 0.40; otherwise, the average was 0.50 across trials.

bols have been used to represent session-by-session estimates of its hits and false alarms for each payoff matrix: open squares indicate the results obtained when hits were "poorly" reinforced; filled triangles, when payoffs were symmetrical, circles for the remaining phases, when hits were favored by a 3:1 margin or better. There is virtually no overlap between the

phases. Also, the obtained function is located at some distance from the major diagonal or "chance" line in the square, which indicates a relatively high level of sensitivity overall. In short, the clustering of data along different portions of the function reflects a systematic change in response bias with payoffs—one that yields a linear ROC with unit slope, when

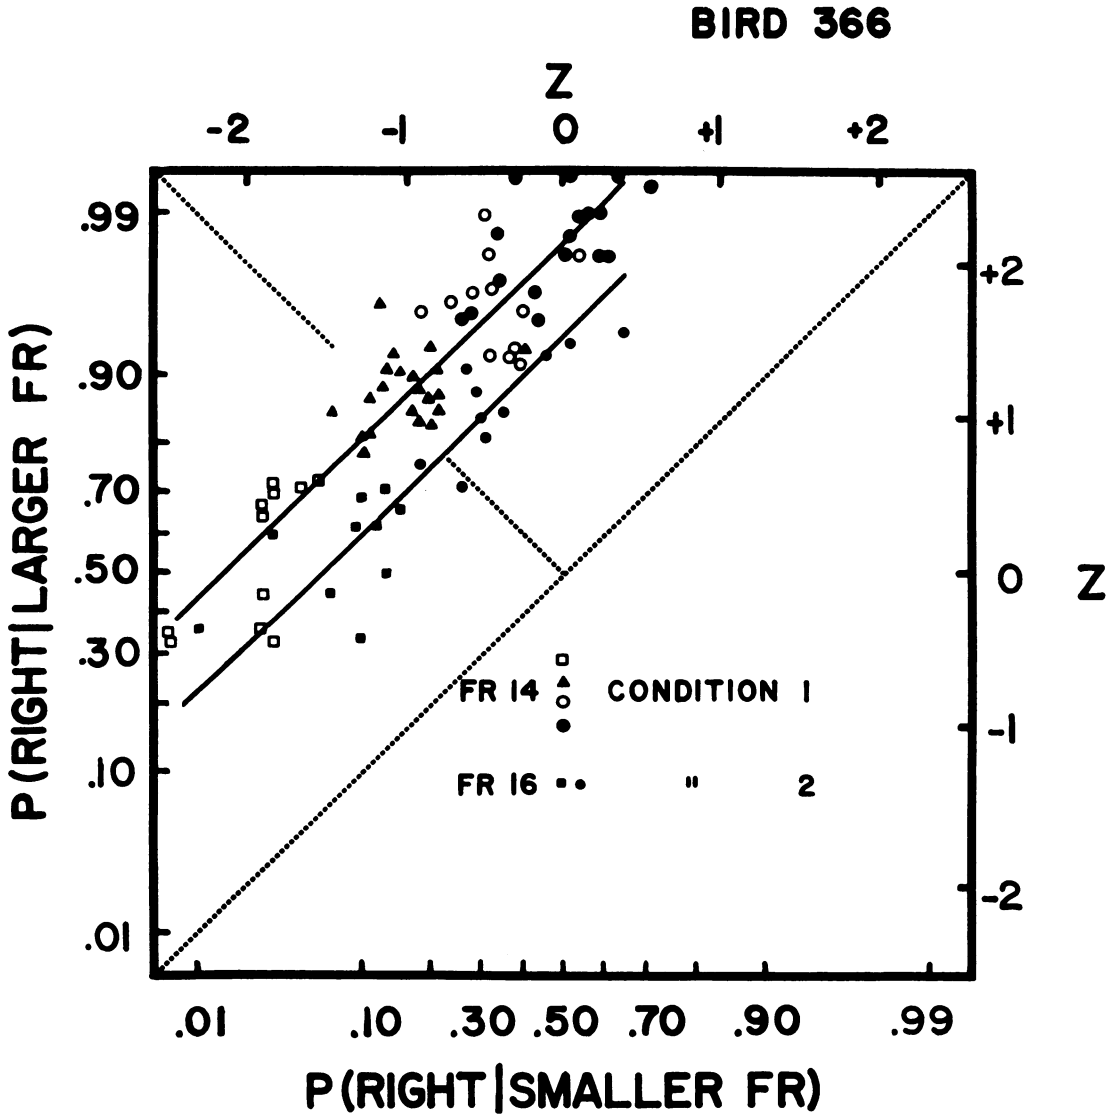


Fig. 1. Receiver operating characteristics for Bird 366 on normalized coordinates, with corresponding z-scores indicated on the top and right-hand margins of the square. Session data are shown. See text and Table 2 for further details.

normalized. In other words, the subject's performance may be quite simply described, by the equal-variance, Gaussian model of signal-detection theory (or some equivalent). Quantitatively, this means that sensitivity (d') may be estimated in standard deviation units from any point on the function (see Elliott's tables in Swets, 1964). Here, however, d_e , a more general statistic, has been adopted; it is equivalent to the d' of an ROC function at the equal-bias line (minor diagonal) of the square (Egan and Clarke, 1966). The obtained value appears in

Table 2, with the function's slope included, to characterize the subject's overall performance.

Similar results were obtained from Bird 366 at a lower level when relative reinforcement magnitudes were reversed at mid-session (see Figure 1, Condition 2). As the figure shows, performance consistently fell below the level attained at FR 14, an easier discrimination. Response bias varied enough across sessions to define the subject's ROC for a range of values, though only two matrices were used. Thus, a function was obtained based on several points

Table 2
Parameters of the normalized ROC curves
for individual subjects (Conditions 1 to 9)

Condition	Largest Payoff Ratio	Sample Size (sessions)	Curve			
			d_s	Slope	Ratios Tested	
BIRD 366						
1	4:1	11-20	2.16	1.0	FR 14	FR 20
2	3:1	10	1.57	1.0	FR 16	FR 20
BIRD 363						
3	7:3	10	1.72	1.1	FR 6	FR 10
4	3:1	10			FR 7	FR 10
5	3:1	15	1.82	1.5	FR 7	FR 10
6	4:1	10			FR 7	FR 10
7	6:1	10			FR 7	FR 10
BIRD 323						
8A	7:3	21	1.72	1.4	FR 24	FR 30
8B	7:3	21	1.68	1.2	FR 24	FR 30
8D			1.72	1.2	FR 24	FR 30
8D	4:1	10	1.79	1.2	FR 24	FR 30
9A	7:3	15	1.00	0.9	FR 26	FR 30
9B	4:1	10	1.08	1.0	FR 26	FR 30

Note: For Bird 363 at FR 7, the ROC was determined from the medians of Conditions 4 to 7, because all the curves were similar at that level. In all other cases, session data were used. For Condition 3, only the early results (Sessions 18 to 27) are shown. For 8D, the first entry refers to those parts of the sessions where payoffs alternated in blocks, and the second to where they alternated at random. To demonstrate that relative sensitivity increased with absolute ratio size, as reported earlier (Hobson, 1975), compare d_s values at FR 10, 20, and 30 for equal ratio differences—e.g., note that Bird 323 did “better” at FR 24 versus FR 30 than Bird 366 at FR 16 versus FR 20.

(session data), just as in the first condition. As for bias separation, a clear difference emerged by the third test day and was maintained, with no appreciable effect on sensitivity.

At first, the results for Bird 363 at FR 10, with cues, were similar to those for Bird 366 at FR 20, in the uncued conditions (see Table 2 for the ROC parameters). In fact, by the second week of testing, the subject was able to maintain separate hit and false-alarm rates by matrix within individual sessions. These data, and the medians from the first month of testing, appear in Table 3. Later, however, performance differences emerged. The final ROC function, not shown here, fell above and substantially to the left of the earlier one, which indicates that while sensitivity increased, a left-key bias developed overall. In addition, hit and false-alarm rates partially overlapped in most sessions, though not on the average: in the long run, a rank-ordering of response

bias according to payoff prevailed across but not within sessions (see Table 3). To gain better control of bias on a day-to-day basis at a lower level, the subject was shifted from FR 6 to FR 7 versus FR 10 and tested there with values more extreme than those it first encountered.³

The ROC function for Bird 363 at the new level appears in Figure 2, which plots the medians obtained from each set of payoff matrices with separate symbols, to simplify the figure. Open symbols indicate that relative reinforcement magnitude varied across cue color; filled ones refer to changes in relative reinforcement probability. Half-tone symbols have been used for Condition 7 because both magnitude and probability were varied. Plotted in this fashion, the results clearly indicate that, both within and across conditions, sensitivity remained constant as response bias varied with payoffs. They also show that virtually the same performance was obtained from a 3:1 distribution of reinforcement probabilities as magnitudes (i.e., under the green and red matrices in Conditions 4 and 5). The medians from Condition 7 indicate that bias may change—in this case, increase—during testing without affecting sensitivity. As for session data, the results showed that separate biases appeared relatively early and were maintained in most, if not all, of the individual sessions (see Table 3). Overall, then, similar performances were obtained when payoffs varied across blocks of trials with color cues as in the uncued conditions. However, they differed in one detail: the ROC for Bird 363 in the FR 7 conditions was relatively steep (slope > 1.0) on normalized coordinates (see Table 2). This result indicates that while the normal transform may be retained for curve-fitting purposes, the underlying model requires some modification to accommodate individual differences in slope.

Figure 3 shows the ROC functions obtained from Bird 323 at two levels, using more-frequent changes in payoffs than Bird 363 encountered. Again, medians have been plotted rather than session data (but see Table 3; also,

³Another subject, which was initially trained like Bird 363 at a different ratio size (FR 24 versus 30), failed to maintain a separate set of biases in most sessions until shifted to Condition 7 at FR 25. Its median hit and false-alarm percentages for that condition were: 41, 3 (green); 74, 13 (white); and 97, 64 (red).

Table 3

Analysis of performance according to payoff matrix in each of the color-cued conditions. The top section presents session-by-session results from the second week of testing. The middle section shows the median hit and false-alarm percentages for the final sessions in each condition. The last row indicates the percentage of final sessions within which hits and false alarms were rank-ordered according to payoffs.

Condition	3	4	5	6	7	8A	8B	8D ^a	9A	9B
<i>Matrix</i>	<i>Percentage of hits and false alarms for Sessions 8-12</i>									
Green	66, 12 70, 14 66, 18 58, 12 52, 08	59, 09 57, 04 72, 08 80, 24 71, 24	64, 12 54, 09 76, 05 61, 04 71, 11	69, 09 80, 15 87, 07 73, 05	49, 12 40, 05 40, 11 48, 03 53, 05	71, 12 77, 17 84, 23 66, 16 73, 19	70, 14 59, 07 54, 07 69, 18 77, 15	60, 13 (71, 07) 47, 06 (66, 10) 64, 12 (67, 10) 67, 07 (53, 03) 63, 12 (66, 06)	56, 14 53, 23 37, 12 54, 13 40, 13	44, 08 42, 11 31, 11
White	82, 34 78, 24 80, 32 89, 32 70, 22	79, 09 75, 15 84, 16 95, 41 91, 37	85, 09 76, 11 92, 27 72, 11 83, 19	76, 12 88, 12 89, 23 88, 17	89, 21 83, 28 71, 12 80, 20 80, 27	92, 23 88, 17 61, 29 83, 25 91, 19	88, 27 90, 41 80, 29 85, 20 85, 37		74, 38 77, 29 74, 37 80, 38 79, 43	
Red	88, 32 94, 52 94, 68 99, 51 92, 57	97, 33 95, 40 93, 35 99, 44 91, 53	95, 32 96, 32 100, 52 95, 37 88, 31	92, 25 100, 40 99, 48 97, 41	95, 36 97, 49 97, 47 91, 36 96, 44	93, 33 95, 36 92, 39 99, 41 96, 29	89, 39 93, 32 91, 40 95, 47 94, 48	86, 54 (96, 60) 94, 44 (93, 35) 97, 56 (99, 43) 97, 64 (82, 39) 98, 68 (97, 63)	89, 70 89, 63 77, 51 87, 61 85, 64	93, 61 97, 73 96, 68
<i>Session N</i> (trials/matrix)	100 ^b	150	150	150	150	150 ^b	200	140 140	200	280
<i>Matrix</i>	<i>Percentage of hits and false alarms, medians (final sessions)</i>									
Green	61, 05 ^c	71, 13	72, 11	47, 05	34, 04	65, 11	69, 14	45, 04 (52, 04)	54, 17	48, 12
White	85, 10	91, 23	83, 17	84, 14	77, 16	87, 25	89, 29		75, 38	
Red	89, 18	94, 33	93, 32	93, 35	99, 49	95, 39	93, 40	98, 55 (97, 47)	88, 63	92, 61
<i>Median N</i> (Sessions)	10	10	15	10	10	21	21	10	15	10
<i>Total</i> <i>Sessions</i>	53	17	21	11	23	63	27	21	23	10
<i>% Sessions</i> <i>with bias</i> <i>rank ordered</i>	40	60	60	80	100	90	71	100	100	100

^aFor Condition 8D, results for the blocked sections are listed first, followed by those from the random sections in parentheses.

^bFor Conditions 3 and 8A, trials per matrix were later increased by 50% over the listed values.

^cFor Condition 3, the medians obtained earlier in testing were: 67, 11; 84, 25; 95, 43.

Figure 4). For Condition 8D, which included random alternation, data from an earlier 10-day period are plotted along with the final set to show that the same kind of changes in bias over time occurred there as in other situations, such as Condition 7. The results for 8C are not shown because the subject's performance deteriorated during its first exposure to random conditions. They constitute the only exception. Otherwise, similar results were obtained at both levels, whether payoffs varied every 50 trials or much more often. Sensitivity estimates derived from session data support this conclusion; also, slope differences were small (see Table 2). Still, it should be noted that at FR 24, the ROC slopes, initially quite

steep, tended to decrease in the later conditions—a trend that did not appear in the other subjects' data. In any case, even in the beginning, when payoff differences were moderate, the subject was able to maintain a much clearer bias separation than Bird 363 within individual sessions (*cf.* Conditions 8A and 8B to 3 and 4 in Table 3). In fact, as Figure 4 indicates, the separation was maintained from one day to the next, across many sessions.

DISCUSSION

The present study showed that, in a ratio-difference task, response bias may be systematically varied, using differences in relative rein-

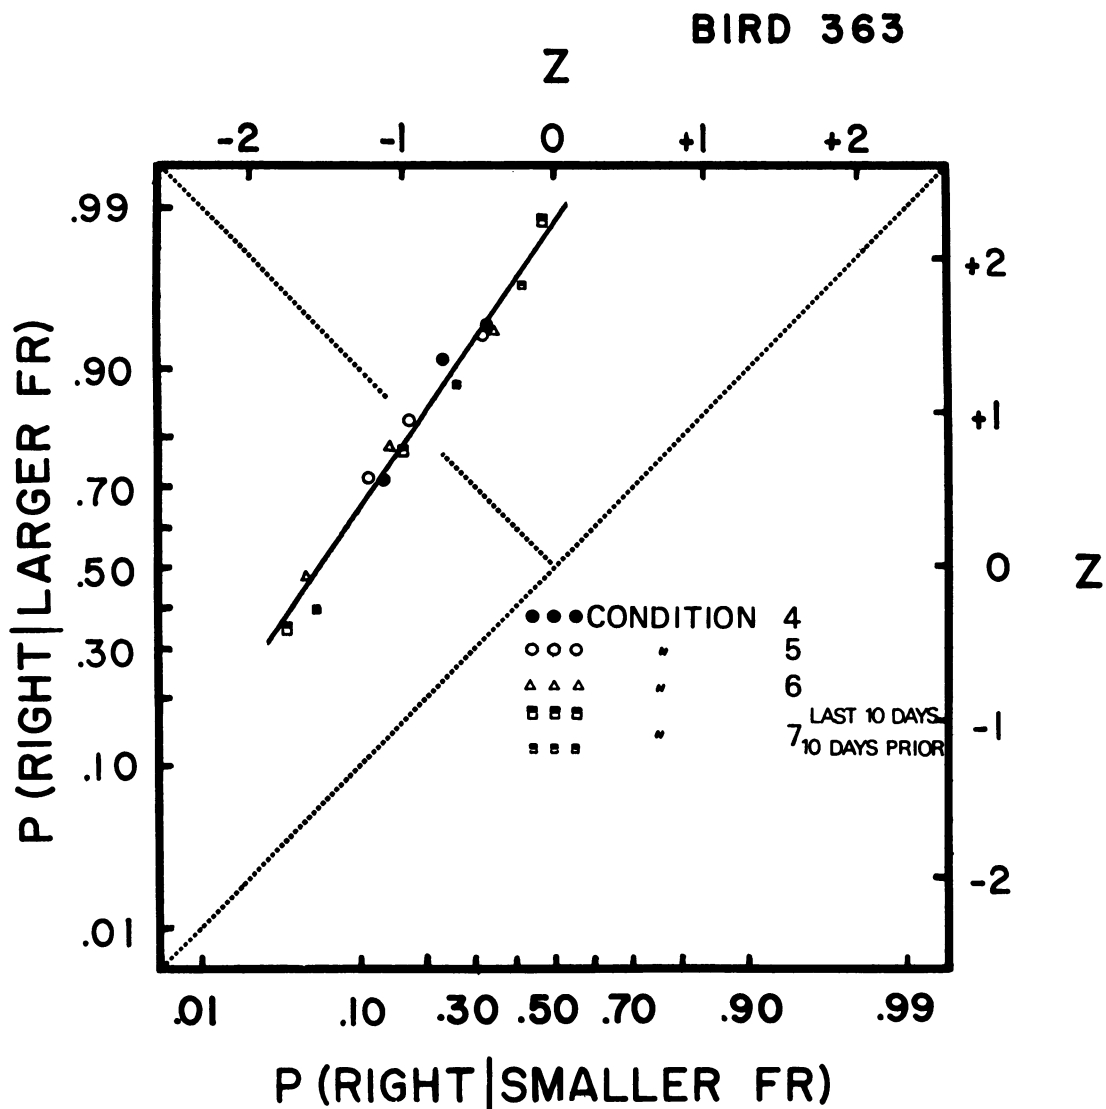


Fig. 2. Receiver operating characteristic for Bird 363 at FR 7 versus FR 10, on normalized coordinates. The results (medians) are for the green, white, and red matrices in each condition. Some of the session data appear in Table 3.

forcement probabilities and/or magnitudes, without appreciably affecting sensitivity. Overall, the resulting functions (ROCs) were like those routinely obtained for other species and dimensions (Clopton, 1972; Nevin *et al.*, 1975; Stubbs, 1976). Further, within subjects, ROCs were similar across different payoff conditions, which suggests that a single response mechanism exists, however sensory information is processed. The demonstration is on the order of Galanter and Holman's (1967) work with human subjects, although admittedly less ex-

tensive (see also Green and Swets, 1966, pp. 87-91). Thus, the study provides a strong demonstration of the applicability of signal-detection methods and theory to animal psychophysics in general and fixed-ratio discriminations in particular.

To pursue the parallel further, Green and Swets (1966, p. 93) noted that, typically, human subjects can perfectly rank-order their biases according to payoffs for at least five different matrices and come close to that for 10. The same sort of correspondence was obtained

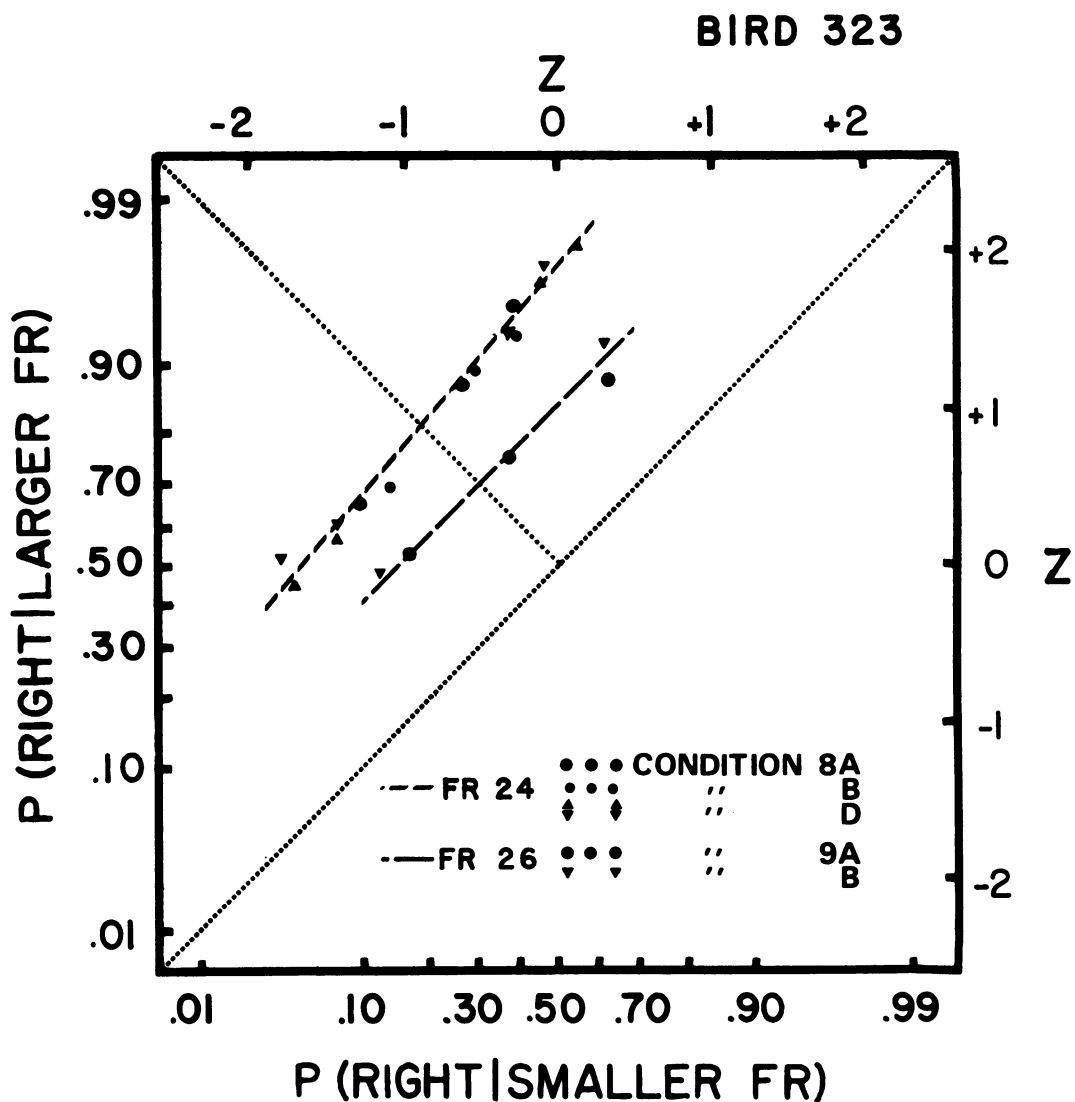


Fig. 3. Receiver operating characteristics for Bird 323 on normalized coordinates. The results (medians) are plotted separately for each of the blocked- and random-cue conditions. For session data, see Figure 4 and Table 3.

here. The largest number of matrices presented at a single level was seven (Bird 363, Conditions 4 to 7). The results for all subjects are shown in Figure 5 in terms of the proportion of right-key responses overall as a function of relative access to reinforcement for right-key responses—that is, as a function of relative payoff for hits, assuming a multiplicative relation between reinforcement probabilities and magnitudes. The connected lines indicate optimal response proportions (the performance) required to maximize total access to reinforcement), given the ROC functions for individual

subjects. The separate points represent obtained proportions (actual performance), as calculated from the final medians (for details, see Appendix; also, Elmsore, 1972). Consider first the results for Bird 363 in the 3:1 conditions. The similarity suggests that reinforcement probabilities and magnitudes are interchangeable, at least when they are separately varied. Although preliminary, the finding is provocative because exceptions recently have been reported for concurrent reinforcement schedules. For example, using a cueing combination similar to the present one, Todorov

BIRD 323

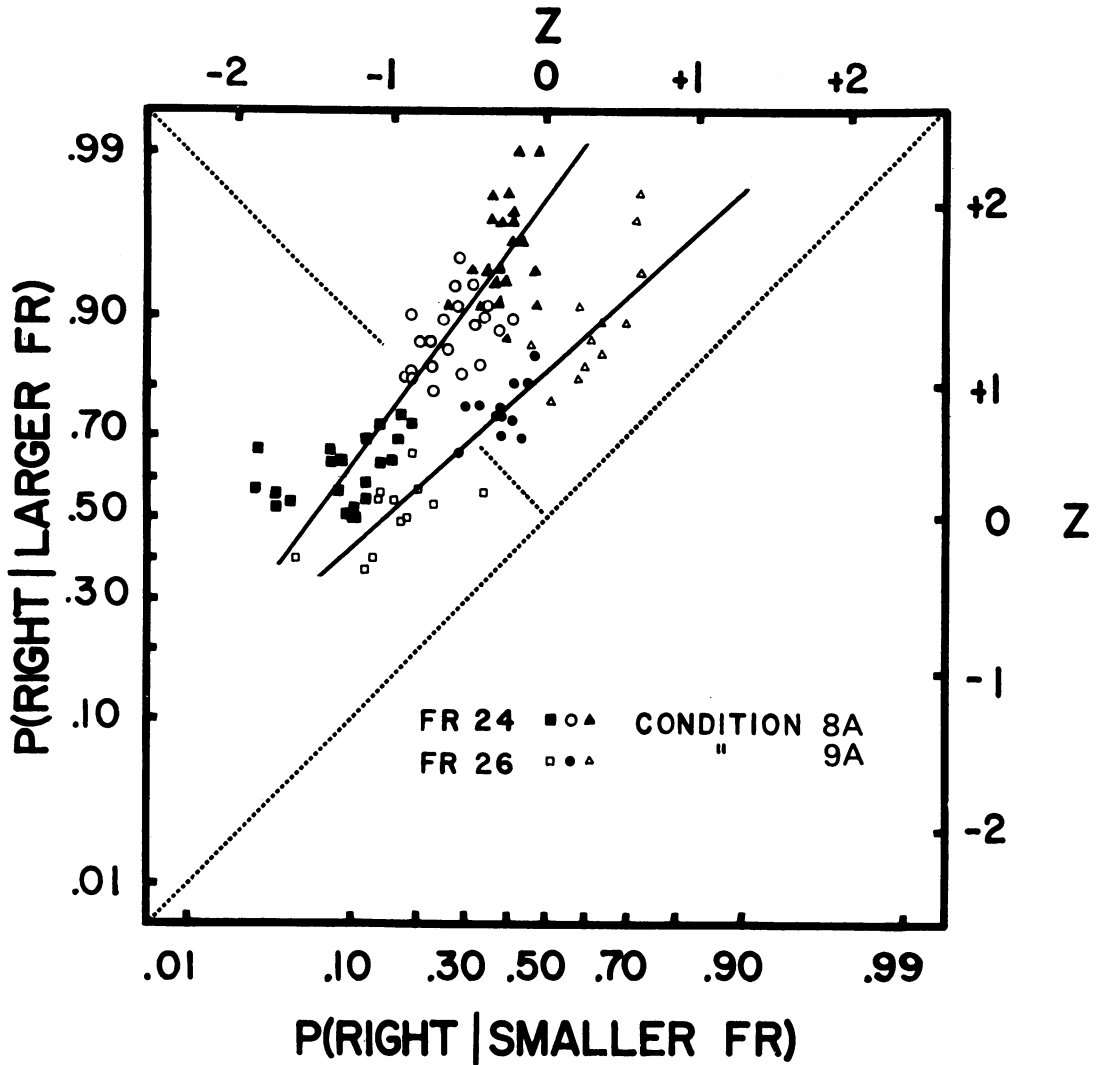


Fig. 4. Receiver operating characteristics for Bird 323, based on session data from the green (squares), white (circles), and red (triangles) in two conditions.

(1973) found that reinforcement probability was the more important determinant of choice. The comparison is mentioned here to emphasize the difficulty one encounters extrapolating across different tasks and procedures, which should underscore the need for additional research in both animal psychophysics and concurrent schedules. In any case, the match between obtained and optimal values is a close one, overall (see Figure 5), and thus similar to what Elsmore and others (especially Hume, 1974; Hume and Irwin, 1974) have reported. In other words, in a variety of psychophysical

tasks, subjects appear to maximize total reinforcements or some closely related quantity. Contrary to Stubbs' (1976) recent suggestion, the correspondence in terms of relative error rates appeared much poorer, which means neither relation is entirely general. Unfortunately, the source of this discrepancy is unclear, no doubt because in psychophysics the focus has always been on sensory processes rather than bias mechanisms. In a recent review, Duso (1975) reported that he failed to find "any adequate experimental foundation for current treatments of bias", for all the varied

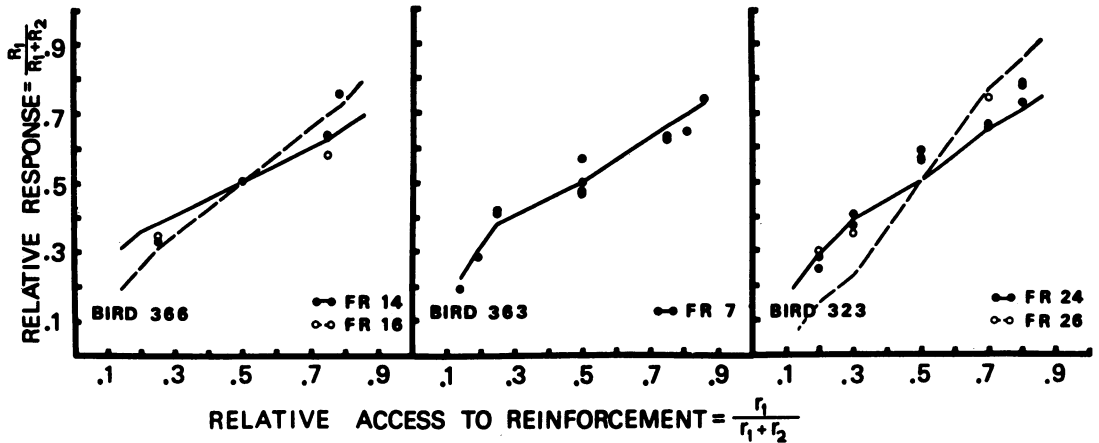


Fig. 5. Comparison of optimal and obtained bias for different payoffs. The relation is expressed in terms of right-key responses and right-key reinforcers. See text and Appendix for further details.

proposals. At this stage, then, what we have are ordinal relations: like humans, animals are responsive to differences in relative payoff across a wide range of values.

Thus, the present results indicate that infra-human subjects can be readily induced to adopt and maintain different biases at different times by any one of several methods. However, the results of Condition 8C suggest one possible limitation to cueing: even in the "best" of subjects (Bird 323) performance may deteriorate unless the random-cue procedure is simplified in some way. In the present study, recovery occurred once the symmetrical matrix was eliminated and payoff differences were increased. Additional research is required to determine exactly the factors influencing recovery. With this one exception, alternation in payoffs had relatively little effect. The results from the modified procedures (8D, 9B) provide a clear instance of stimulus control of response bias by payoffs on a trial-by-trial basis, with sensitivity invariant. Hence, response bias appears to be a kind of complex operant subject to differential reinforcement under cueing, in much the same way that individual responses are controlled by multiple schedules in simpler, free-operant situations. The results of Condition 2, where payoffs were reversed at mid-session without cues, suggest that response bias may be brought under mixed-schedule control as well.

Several parallels have been drawn between animal and human work, in a variety of tasks, to emphasize the isolability of bias effects, a basic premise of detection theory. However,

the results also point up a limitation in the overall approach. In particular, the variation in ROC slopes (see Table 2) is difficult to interpret. In human work, ROCs are known to vary, not only across subjects, but also within subjects across sessions and procedures, not to mention levels. Typically, within levels, the change consists of a slight rotation of the function about the equal-bias line in the ROC space, with no overall change in location (Egan and Clarke, 1966; Markowitz and Swets, 1967). That a slope change of this kind occurred here as well (Condition 8) may reflect a common source. Certainly, the result argues strongly for the use of $d_{0.5}$, or some equivalent statistic, to index sensitivity at equal bias, the most stable region of the function. However, in addition, it means that a more complex account is required. For example, according to the general Gaussian model of signal-detection theory, the slope of the normalized ROC function reflects the ratio of the variances of the underlying distributions, which may not be equal, while its location reflects the distribution means scaled in terms of the average variance. An interpretation of this kind suggests, of course, that future research focus on response variability during ratio execution. For other possible accounts, see Egan (1975), Green and Swets (1966), and Pastore and Scheirer (1974). However, none of these models specifically addresses slope variability within levels. Rather, they were developed to predict changes of a different sort: the progressive flattening of the normalized ROCs at higher sensitivity levels, a frequent occur-

rence in vision and audition, but not apparent here—the fixed-ratio ROCs are relatively steep (slopes $\cong 1.0$) by comparison. The results of Condition 3, where an overall shift in the ROC occurred, may reflect a different phenomenon. Here again, it is not clear how the change should be construed—as a stimulus effect or a consequence of the subject's response criterion. The deterioration in bias, which also occurred, does not readily suggest a learning interpretation.

In any case, while generally small in size, the slope changes appear complex, which recommends the inclusion of within-session procedures, at least for further study. As Kinchla *et al.* (1966) noted, an "ROC generated in this fashion has the advantage of not being affected by session-to-session changes in sensitivity since each sensitivity level is equally represented at each point on the curve." In addition, the techniques may provide a relatively efficient means of estimating sensitivity in well-trained subjects, because replication is built in.

Meantime, the results should encourage experimenters to intervene when response biases develop during psychophysical testing, by counteracting them with appropriate changes in payoffs (or stimulus probabilities) so that a relatively pure measure of sensitivity, like d_s , may be obtained. In practice, this has not been done. Instead, psychometric functions are obtained with stimulus probabilities and payoffs equal throughout, even though, invariably, response biases differ across subjects and within subjects across sessions as stimulus differences are reduced (Hobson, 1975; Irwin and Terman, 1970). Subsequently, the researcher must either discount bias, assume that d' is the appropriate measure, or adopt its nonparametric equivalent, with some loss in power. In my own case (Hobson, 1975), the stimulus effects seemed large enough to override the problem, an argument that the present ROCs confirm. The d_s values derived from them show that, in relative terms, sensitivity increased with ratio size (see Table 2).

REFERENCES

Clopton, B. M. Detection of increments in noise intensity by monkeys. *Journal of the Experimental Analysis of Behavior*, 1972, 17, 473-481.

Dusoir, A. E. Treatments of bias in detection and

recognition models: a review. *Perception and Psychophysics*, 1975, 17, 167-178.

Egan, J. P. *Signal detection theory and ROC analysis*. New York: Academic Press, 1975.

Egan, J. P. and Clarke, F. R. Psychophysics and signal detection. In J. B. Sidowski (Ed), *Experimental methods and instrumentation in psychology*. New York: McGraw-Hill, 1966. Pp. 211-241.

Elsmore, T. F. Duration discrimination: effects of probability of stimulus presentation. *Journal of the Experimental Analysis of Behavior*, 1972, 18, 465-469.

Galanter, E. and Holman, G. L. Some invariances of the isosensitivity function and their implications for the utility function of money. *Journal of Experimental Psychology*, 1967, 73, 333-339.

Green, D. M. and Swets, J. A. *Signal detection theory and psychophysics*. New York: Wiley, 1966.

Hobson, S. L. Discriminability of fixed-ratio schedules for pigeons: effects of absolute ratio size. *Journal of the Experimental Analysis of Behavior*, 1975, 23, 25-35.

Hume, A. L. Auditory detection and optimal bias. *Perception and Psychophysics*, 1974, 15, 425-433.

Hume, A. L. and Irwin, R. J. Bias functions and operating characteristics of rats discriminating auditory stimuli. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 285-295.

Irwin, R. J. and Terman, M. Detection of brief tones in noise by rats. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 135-143.

Kinchla, R. A., Townsend, J., Yellott, J. I. Jr., and Atkinson, R. C. Influence of correlated visual cues on auditory signal detection. *Perception and Psychophysics*, 1966, 1, 67-73.

Markowitz, J. and Swets, J. A. Factors affecting the slope of empirical ROC curves: comparison of binary and rating responses. *Perception and Psychophysics*, 1967, 2, 91-100.

Nevin, J. A. On differential stimulation and differential reinforcement. In W. C. Stebbins (Ed), *Animal psychophysics*. New York: Appleton-Century-Crofts, 1970. Pp. 401-423.

Nevin, J. A., Olson, K., Mandell, C., and Yarensky, P. Differential reinforcement and signal detection. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 355-367.

Pliskoff, S. S. and Goldiamond, I. Some discriminative properties of fixed-ratio performance in the pigeon. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 1-9.

Rilling, M. Number of responses as a stimulus in fixed interval and fixed ratio schedules. *Journal of Comparative and Physiological Psychology*, 1967, 63, 60-65.

Rilling, M. and McDiarmid, C. G. Signal detection in fixed-ratio schedules. *Science*, 1965, 148, 526-527.

Stubbs, D. A. The discrimination of stimulus duration by pigeons. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 223-238.

Stubbs, D. A. Response bias and the discrimination of stimulus duration. *Journal of the Experimental Analysis of Behavior*, 1976, 25, 243-250.

Swets, J. A. *Signal detection and recognition by human observers: Contemporary readings*. New York: Wiley, 1964, Appendix 1.

- Swets, J. A. The relative operating characteristic in psychology. *Science*, 1973, 182, 990-1000.
- Ten Eyck, R. L. Jr. Effects of rate of reinforcement upon concurrent operant performance. *Journal of the Experimental Analysis of Behavior*, 1970, 14, 269-274.
- Terman, M. and Terman, J. S. Concurrent variation of response bias and sensitivity in an operant-psycho-physical test. *Perception and Psychophysics*, 1972, 11, 428-432.
- Todorov, J. C. Interaction of frequency and magnitude of reinforcement on concurrent performances. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 451-458.
- Wright, A. A. Psychometric and psycho-physical hue discrimination functions for the pigeon. *Vision Research*, 1972, 12, 1447-1464.

Received 26 January 1977.

(Final acceptance 8 February 1978.)

APPENDIX

An iterative procedure similar to Elsmore's (1972) was used to determine optimal bias values. First, for each payoff matrix, total access

to reinforcement (r) was computed for each point on a subject's ROC function to locate the hit (H) and false-alarm (FA) percentages required to maximize its value, assuming a multiplicative relation between relative reinforcement probability (f) and magnitude (m). The following formula was used:

$$r = H(p_{r|H}) + 1 - FA(p_{r|1-FA})$$

$$\text{where } p_{r|H} = f_1(m_1) = r_1 \\ \text{and } p_{r|1-FA} = f_2(m_2) = r_2.$$

Once the optimal values were determined, they were converted to a relative response measure, using the following formula for bias:

$$\frac{R_1}{R_1 + R_2} = \frac{p_H + p_{FA}}{2}.$$

To determine obtained bias, median hit and false-alarm rates were substituted into the formula. Both sets of values are plotted in Figure 5 as a function of relative access to reinforcement ($r_1/r_1 + r_2$).