

REINFORCEMENT CONTINGENCIES AND SIGNAL DETECTION

JOHN A. NEVIN, PETER JENKINS, STEPHEN WHITTAKER,
AND PETER YARENSKY

UNIVERSITY OF NEW HAMPSHIRE

Pigeons were trained to discriminate temporal stimuli in a discrete-trial signal-detection procedure. Pecks to one side key were reinforced intermittently after exposure to one duration, and pecks to the other side key were reinforced intermittently after exposure to a different duration. In Experiment I, the allocation of reinforcers was varied systematically for correct responses and for errors, using a procedure that controlled the obtained numbers of reinforcers. When reinforcers were allocated symmetrically, the level of discrimination decreased as the proportion of reinforcers for errors increased. When reinforcers were allocated asymmetrically, the decrease in discrimination was less systematic. Bias toward one or the other side key roughly matched the ratio of reinforcers obtained by pecks at those keys, independent of the level of discrimination. In Experiment II, the overall rate of reinforcement for correct responses was varied both within and between experimental conditions. The level of discrimination was positively related to the overall rate of reinforcement. The discrimination data of both experiments were interpreted in relation to the contingencies of reinforcement and nonreinforcement, characterized by the average difference in reinforcement probability for correct responses and errors.

Key words: signal detection, contingencies, reinforcement, nonreinforcement, discrimination, bias, key peck, pigeons

An adequate formulation of the interaction between an organism and its environment must always specify three things: (1) the occasion upon which a response occurs, (2) the response itself, and (3) the reinforcing consequences. The interrelations among them are the "contingencies of reinforcement." (Skinner, 1969, p. 7)

The discriminated operant defined by Skinner's three-term contingency is fundamental to the analysis of behavior. A simple formal characterization of the three-term contingency requires a well-defined set of mutually exclusive and exhaustive events: the presentation of a

particular state of the environment (S_1) or the alternative state (S_2); the occurrence of a particular response (B_1) or its alternative (B_2); and the consequent presentation of a reinforcer (R_1) or the state of affairs defined by its absence (R_2). Taken together, the possible joint events may be represented as a 2 by 2 by 2 cube, depicted in Figure 1. Each plane of the cube represents a 2 by 2 contingency table of the sort used by Gibbon, Berryman, and Thompson (1974) to represent simple classical contingencies, where the response is irrelevant to the stimulus-reinforcer contingency, or simple operant contingencies, where there are no explicit environmental stimuli correlated with response-reinforcer contingencies.

A side view of the cube collapses the response dimension and shows how R_1 and R_2 depend upon S_1 and S_2 . This aspect of the cube represents stimulus-reinforcer contingencies, as in classical conditioning. A top view of the cube collapses the stimulus dimension and shows how R_1 and R_2 depend on B_1 and B_2 . This aspect of the cube represents response-reinforcer dependencies, as in operant choice experiments. Considered separately, the top and bottom planes represent response-reinforcer dependencies signaled by distinctive en-

This research was based on a series of discussions in graduate seminars at the University of New Hampshire, and was supported in part by grant 76-11028 from the National Science Foundation. We are indebted to Dr. Michael Davison for constructive commentary on earlier drafts of this paper and for stimulating correspondence on alternative approaches to this problem area. We are also indebted to Dr. Dianne McCarthy for providing us with preprints of her research and for preliminary accounts of her current work. Reprints are available from John A. Nevin at Department of Psychology, University of New Hampshire, Durham, New Hampshire 03824. Dr. Peter Jenkins is now at Vassar College, Poughkeepsie, New York, and Dr. Stephen Whittaker is now at the College of Optometry, Philadelphia, Pennsylvania.




Fig. 1. The contingency cube relating the three terms of the discriminated operant: Stimuli S_1 or S_2 , Responses B_1 or B_2 , and Consequences R_1 or R_2 . See text for explanation.

environmental conditions, and thus correspond to components of a multiple schedule, where each component is itself a concurrent schedule.

A front view of the cube shows how different kinds of outcomes depend on stimuli and responses. The standard yes-no signal-detection experiment is a special but well-known case. On each trial, either signal (S_1) or noise (S_2) is presented. A positive response, "yes" (B_1) or its complement, "no" (B_2) is emitted, and either a positive (R_1) or negative (R_2) outcome is presented, depending on the stimulus and response. R_1 typically follows correct responses ($B_1|S_2$ or $B_2|S_2$), whereas R_2 typically follows errors ($B_1|S_2$ or $B_2|S_1$). The relations between stimuli, responses, and outcomes may be summarized in the form of a payoff matrix by collapsing the front and back planes of the contingency cube (Figure 1).

In the standard signal-detection experiment, R_1 and R_2 are maximally contingent upon stimuli and responses. That is, R_1 is presented if, and only if, S_1 has been presented and B_1 occurs, or if S_2 has been presented and B_2 occurs. Likewise, R_2 is presented if, and only if, S_2 has been presented and B_1 occurs, or if S_1 has been presented and B_2 occurs. One way to characterize the contingency between stimuli, responses, and outcomes is to consider the differences between conditional probabilities

of the joint events defined by the payoff matrix. In the standard detection experiment, the probability of R_1 given S_1 and B_1 , and the probability of R_1 given S_2 and B_2 , are both 1.0. At the same time, the probability of R_1 given S_1 and B_2 , and the probability of R_1 given S_2 and B_1 , are 0. (The event probabilities in the back plane of the cube are simply the complements of those in the front plane.) The difference, 1.0, defines the maximal contingency.

The contingency may be weakened in either of two ways. 1) If R_1 is sometimes presented given S_1 and B_2 or S_2 and B_1 —that is, if errors are sometimes reinforced—the difference between probabilities of R_1 in the diagonally opposed cells of the matrix is reduced. 2) If R_2 is sometimes presented given S_1 and B_1 , or S_2 and B_2 —that is, if some correct responses are not reinforced—the difference between probabilities of R_1 in the diagonally opposed cells of the matrix is likewise reduced. The following two experiments explore these ways of weakening the contingency between stimuli, responses, and their consequences in a signal-detection procedure.¹

NOTATION AND MEASURES

To avoid cumbersome conditional probability notation, and to permit ready comparison with related work by Davison and McCarthy (1980), the numbers of events in the cells of a plane of the contingency cube will be denoted by the subscripts w , x , y , and z , as shown in Figure 1. Thus, for example, in the R_1 plane, B_w represents the number of occurrences of B_1 on S_1 trials, and R_w represents the number of reinforcers obtained by B_1 on S_1 trials.

Variation of R_w relative to R_z with the $S_1 - S_2$ difference held constant generates the empirical isosensitivity curve, the locus of points reflecting constant effects of the stimuli while response bias changes. The constant effects of the stimuli may be characterized by a measure of discrimination related to one proposed by Luce (1959, 1963), which is defined as:

$$1/\eta = \sqrt{\frac{B_w}{B_x} \cdot \frac{B_z}{B_y}} \quad (1)$$

¹A closely related approach by Catania (1971) identified the discriminated operant within a larger set of procedures and proposed a similar characterization of the reinforcement contingency.

In Luce's work, η was interpreted as a measure of the similarity or confusability of the stimuli, which should remain invariant when the stimuli are unchanged but bias toward B_1 or B_2 is varied. Luce was not concerned with possible effects of the reinforcement contingencies. We will refer to $1/\eta$ as a measure of discrimination, which may or may not vary when the stimuli are constant, depending on the contingencies of reinforcement. As a measure of response allocation, $1/\eta$ has the same metric properties as response-ratio measures commonly employed in studies of concurrent reinforcement schedules (e.g., Baum, 1974). In fact, $1/\eta$ is simply the geometric mean of the ratio of correct to incorrect responses in the presence of S_1 and S_2 , when S_1 and S_2 are presented equally often. Assuming that correct responses are always at least as frequent as errors, the value of η ranges from 0, signifying perfect discrimination, to 1.0, signifying that responding is not differentiated with respect to the stimuli. The logarithm of $1/\eta$ ranges from ∞ when discrimination is perfect to 0 when responses are random with respect to the stimuli. $\log 1/\eta$ is linearly related to d' , a commonly used sensitivity parameter in detection research (cf. Green & Swets, 1966; Luce, 1963, p. 129).

Variation of the physical difference between S_1 and S_2 with R_w and R_z constant generates the empirical isobias curve, the locus of points representing constant effects of the outcomes while discrimination varies. A measure of response bias that has properties similar to the discrimination measure (Luce, 1963) is:

$$b = \sqrt{\frac{B_w}{B_x} \cdot \frac{B_y}{B_z}} \quad (2)$$

The parameter b is simply the geometric mean of the ratios of responses to alternatives 1 and 2 in the presence of S_1 and S_2 when S_1 and S_2 are presented equally often. It is analogous to the ratio of responses in conventional free-operant choice experiments (cf. Baum, 1974).

EXPERIMENT I

It is known that $1/\eta$ depends systematically on the S_1 - S_2 difference, and that b depends systematically on the R_w/R_z ratio in conventional signal-detection experiments with animal subjects (e.g., McCarthy & Davison, 1980a,

b; Stubbs, 1976). It is not clear how these measures will change when the contingency between stimuli, responses, and reinforcers is varied. One way to weaken the contingency between reinforcement and the stimuli controlling performance is to introduce reinforcers for responses conventionally scored as errors: false alarms (reporting "yes" when the signal is absent, denoted B_x) and misses (reporting "no" when the signal is present, denoted B_y). Although the introduction of reinforcement for these response classes makes it doubtful whether they should be termed errors, we will use this term to denote the response producing less frequent reinforcement in a particular stimulus condition.

Previous work from this laboratory (Nevin, 1970; Nevin, Olson, Mandell, & Yarensky, 1975) has shown that reinforcement for errors reduces discrimination relative to a baseline condition in which only correct responses are reinforced. Davison and McCarthy (1980) have recently replicated and extended the findings of Nevin et al. (1975), and have proposed a model based on the generalized matching law (Baum, 1974) as extended to detection procedures by Davison and Tustin (1978) to account for their results. In these studies, the rates of reinforcement for correct responses and errors were not controlled, being dependent upon the subject's performance in a fashion analogous to ratio schedules of reinforcement. Nevin, Jenkins, Whittaker, and Yarensky (Note 3) reported part of a study in which rates of reinforcement for both correct responses and errors were controlled by the use of interdependent concurrent schedules analogous to those employed by Shimp (1966), Stubbs (1976), and Stubbs and Pliskoff (1969). Experiment I is a full report of the Nevin et al. (Note 3) study.

METHOD

Subjects

Three White Carneaux pigeons with extensive experimental histories in a variety of stimulus control and reinforcement schedule procedures served as subjects. They were maintained at a minimum of 80% of their free-feeding weights. During the experiment, their body weights stabilized between 80 and 90% of their free-feeding weights without supplemental feedings.

Apparatus

The experiment was conducted in a three-key Lehigh Valley pigeon chamber, with a houselight mounted above the center key and a grain feeder below it. A Sonalert and speaker were located behind the front panel, on the right side of the chamber. The experiment was programmed by conventional electro-mechanical equipment in an adjacent room.

Procedure

The basic task required discrimination of 2-sec and 3-sec lights. On each trial, the center key was lighted orange. A single peck changed the center key from orange to white. The white light went off after either 2.0 or 3.0 sec, where the durations were scheduled randomly with probability .5, and the two side keys were lighted green. A single peck at either green side key terminated the trial and occasionally produced 3.5-sec access to mixed grain. A new trial began after 6.0 sec, regardless of whether grain was delivered.

The 2-sec light was designated S_1 , and the 3-sec light S_2 , and right-key pecks were designated B_1 and left-key pecks B_2 . On each trial, one probability generator determined whether a reinforcer would be scheduled or not; a second determined whether it would be assigned to B_1 or B_2 ; a third determined whether it would be assigned to the correct or error cell for that response. Once assigned to a cell, the reinforcer remained available until obtained, and no further assignments were made. Note that with this procedure a reinforcer might be assigned to B_w or B_x on a given trial, even though the independent stimulus schedule arranged for S_2 to be presented on that trial and on several successive trials. As a result, several trials might elapse before the subject had an opportunity to obtain a scheduled reinforcer. This procedure controls the numbers of reinforcers assigned to each cell of the payoff matrix, within sampling error, but the number of responses per reinforcer depends in part on the subject's performance.

Sessions ended when 60 reinforcers had been obtained. To minimize contamination by warmup efforts, no response data were recorded until the subject had obtained 10 reinforcers. All reinforcers were recorded, however.

In view of their extensive histories, the subjects were introduced to this procedure with no special training. Initially, S_1 was 2 sec and S_2 was 10 sec. The value of S_2 was gradually reduced to 3 sec over the course of 60 sessions of preliminary training. These values of S_1 and S_2 were chosen to insure a moderate rate of errors, and hence rapid contact with reinforcement for errors when it was programmed. Then, the probabilities that a reinforcer would be assigned to each cell were varied over successive conditions, as listed in Table 1. Each condition remained in effect for 15 sessions, except for Condition 4, which was terminated inadvertently after 11 sessions, and Condition 13 (see below).

Table 1 shows how the proportion of reinforcers allocated to B_1 and the proportion of reinforcers allocated to the error cells varied over the course of the experiment. The first 12 conditions were devoted to independent variation of these proportions, with two determinations of performance in the conventional detection procedure with $R_x = R_y = 0$ at each of the side-key bias conditions. The scheduled ratios $R_w:R_z$ and $R_x:R_y$ were constant and equal within blocks of four conditions, except when $R_x = R_y = 0$. Condition 13 was run briefly to eliminate side-key bias, and then Conditions 14 to 16 were conducted to

Table 1
Order of Experimental Conditions

Condition	Probability of reinforcement assignment per trial				Proportion of reinforcers	
	R_w	R_z	R_y	R_x	for B_1	for errors
1	.20	.05	.05	.20	.50	.20
2	.25	0	0	.25	.50	0
3	.1625	.0875	.0875	.1625	.50	.35
4	.25	0	0	.25	.50	0
5	.40	0	0	.10	.80	0
6	.32	.02	.08	.08	.80	.20
7	.26	.035	.14	.065	.80	.35
8	.40	0	0	.10	.80	0
9	.10	0	0	.40	.20	0
10	.08	.08	.02	.32	.20	.20
11	.065	.14	.035	.26	.20	.35
12	.10	0	0	.40	.20	0
13	.25	0	0	.25	.50	0
14	.25	.25	0	.25	.33	.33
15	.25	.40	0	.40	.24	.38
16	.25	.075	0	.075	.63	.19
17	.25	0	0	.25	.50	0

assess the effects of equal values of R_y and R_z while R_x remained 0, to assess the effects of reinforcement for only one type of error.

RESULTS

Figure 2 presents discrimination and bias indices, averaged across subjects, for the last five sessions of Condition 2 and for all sessions of Conditions 3, 4, 5, and 6. These conditions were chosen to illustrate the separate effects of varying the proportion of reinforcers assigned

to errors and of varying the proportion of reinforcers assigned to B_1 versus B_2 . The lower panel shows that the introduction of reinforcement for errors produced systematic decrements in discrimination, with rapid recovery when error reinforcement was discontinued. The upper panel shows that bias was essentially constant near 1.0 when the proportion of reinforcers allocated to B_1 was .5, and that it rose rapidly to about 3.0 when the proportion of reinforcers allocated to B_1 was changed

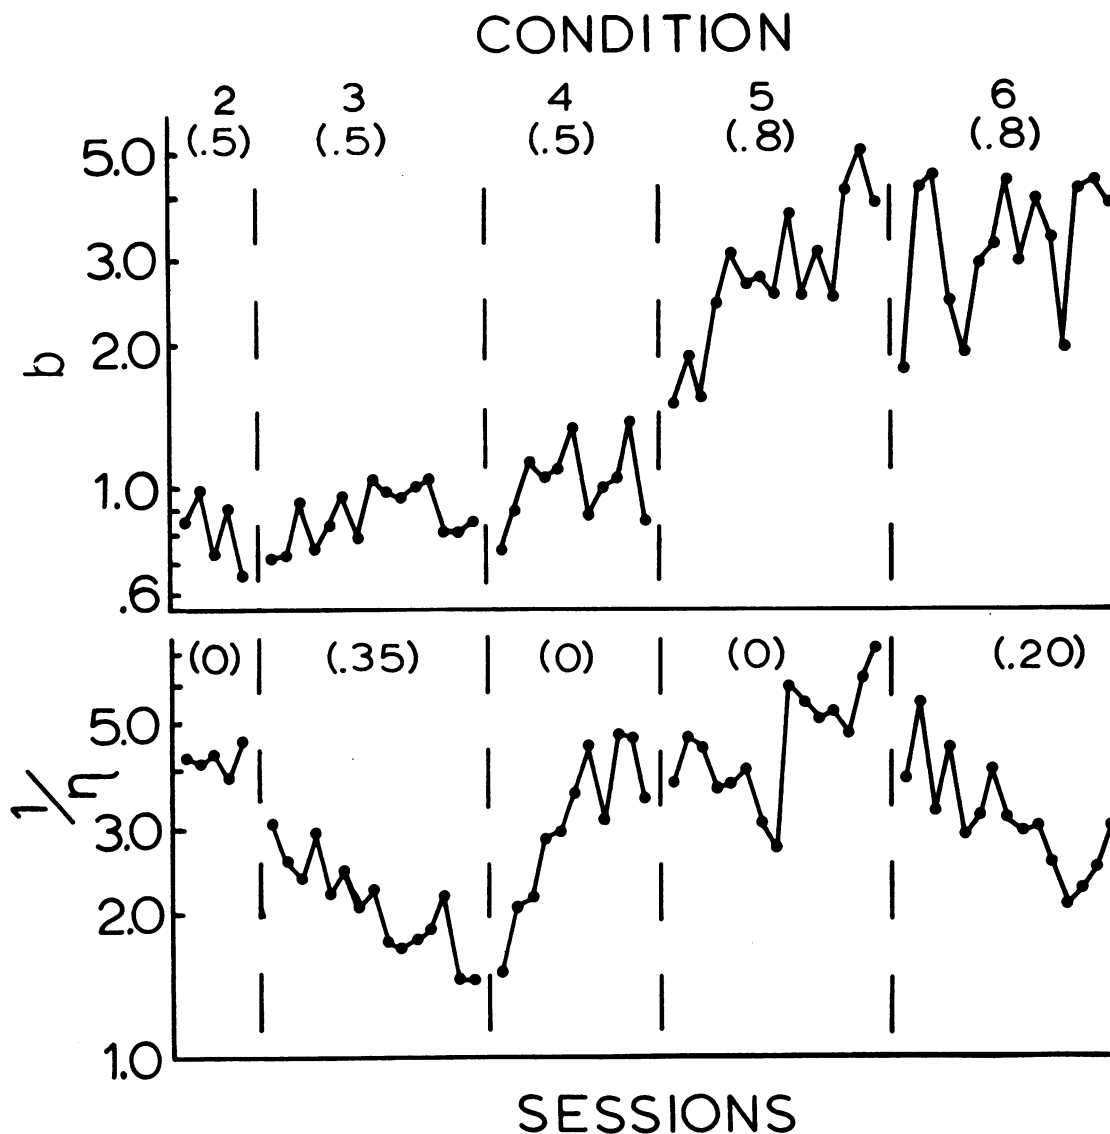


Fig. 2. The index of discrimination, $1/\eta$, and the index of bias, b , during successive sessions of Conditions 2 to 6. The data are geometric means for the three subjects and illustrate the independent determination of discrimination and bias by the allocation of reinforcers to errors and to the left and right responses.

to .8, regardless of changes in discrimination produced by reinforcing errors. Thus, our procedure succeeded in producing independent control over discrimination and response bias.

Examination of Figure 2 suggests that although 15 (or 11) sessions may not have been sufficient to establish stable asymptotic performance in each condition, the direction and approximate magnitude of the experimental effects are clear. The following results are based on data pooled for the final five sessions of each condition. Complete data for individual subjects are presented in Table 2.

The relation between the discrimination index $1/\eta$ and the proportion of reinforcers obtained for errors is displayed in Figure 3, for Conditions 1 through 12. Discrimination was fairly stable across conditions when errors were not reinforced, and recovery data did not differ systematically from those obtained during initial exposure, suggesting that response biases induced by the reinforcement conditions and the order of conditions had little effect. Discrimination clearly decreased as a function of the proportion of reinforcers for errors. Figure 4 presents data for Conditions 14 to 17 and shows that there was no orderly relation between discrimination and proportion of reinforcers for errors,² although discrimination was worse in the error reinforcement conditions (except for one point) than when errors were not reinforced. Thus, it appears that arranging reinforcers for errors has a consistent decremental effect on discrimination, but the magnitude of the effect is not always directly related to the proportion of reinforcers obtained for errors.

Figure 5 shows how response bias (b) depends on the ratio of reinforcers obtained for responding to the two alternatives, summed across stimuli. In this figure, the solid line represents exact equality of these variables:

$$b = \frac{R_w + R_y}{R_x + R_z} \quad (3)$$

The data conform quite well to Equation 3, except that the data for Bird 59 fall below the line. In double-logarithmic coordinates

²Inspection of Table 2 shows that the numbers of reinforcers obtained by R_x given S_A or S_B were not equal in Conditions 14 to 16, although the programmed assignment probabilities were equal. This discrepancy resulted from the failure of a probability generator to operate at its programmed value.

Table 2

Responses emitted and reinforcers obtained, summed for the last five sessions of the conditions of Experiment I. Condition 13 is omitted because it lasted only five sessions.

Condition	B_w	B_x	B_y	B_z	R_w	R_x	R_y	R_z
<i>Bird 58</i>								
1	423	147	110	420	116	33	30	119
2	341	159	136	345	159	0	0	137
3	348	225	252	347	86	59	55	100
4	367	150	175	309	143	0	0	148
5	391	14	187	170	245	0	0	55
6	600	71	410	208	177	12	40	71
7	585	106	513	82	156	20	87	37
8	389	59	272	147	251	0	0	49
9	261	222	29	404	63	0	0	237
10	226	481	35	580	41	61	14	184
11	213	401	112	411	39	106	15	140
12	264	216	89	365	64	0	0	236
14	514	198	193	451	181	69	0	50
15	217	299	26	438	71	139	0	90
16	609	97	239	440	237	47	0	16
17	280	64	41	298	155	0	0	145
<i>Bird 59</i>								
1	298	246	98	469	131	30	33	106
2	265	175	51	371	136	0	0	164
3	303	321	155	448	90	53	64	93
4	296	222	62	399	158	0	0	142
5	392	53	182	202	239	0	0	61
6	461	72	257	236	200	11	48	41
7	534	136	311	307	139	25	94	42
8	383	26	167	225	248	0	0	52
9	171	310	16	402	55	0	0	245
10	219	396	23	519	39	60	18	183
11	229	483	67	498	36	93	31	140
12	171	273	27	406	50	0	0	250
14	378	200	67	427	180	62	0	58
15	92	395	13	401	52	138	0	110
16	499	204	366	303	241	39	0	20
17	297	110	63	277	167	0	0	133
<i>Bird 60</i>								
1	351	163	270	327	109	30	35	116
2	382	38	49	346	155	0	0	152
3	454	207	230	384	89	57	55	89
4	379	23	58	298	153	0	0	147
5	402	11	90	247	245	0	52	55
6	486	42	265	188	192	9	80	47
7	480	71	348	194	157	26	0	37
8	383	4	96	303	238	0	0	62
9	290	124	10	383	54	0	11	246
10	362	217	60	489	57	58	31	174
11	189	428	64	468	34	86	0	149
12	310	77	12	365	61	0	0	239
14	406	173	129	418	189	69	0	42
15	235	287	13	439	83	132	0	85
16	569	166	93	519	224	53	0	23
17	353	56	46	278	183	0	0	117

such as those of Figure 6, this consistent discrepancy implies an additional source of bias that operates as a constant multiplier on the




Fig. 3. Discrimination between S_1 and S_2 , indexed by $1/\eta$, as a function of the proportion of reinforcers obtained for errors in Conditions 1 to 12. Unconnected data points are for replications of conditions in which errors were not reinforced.

right side of Equation 3. It is worth noting that the data for Conditions 1 to 12 do not depend on whether errors were reinforced or not. However, the data for Conditions 14 to 16 appear to fall below the majority of the data for Conditions 1 to 12, suggesting an additional

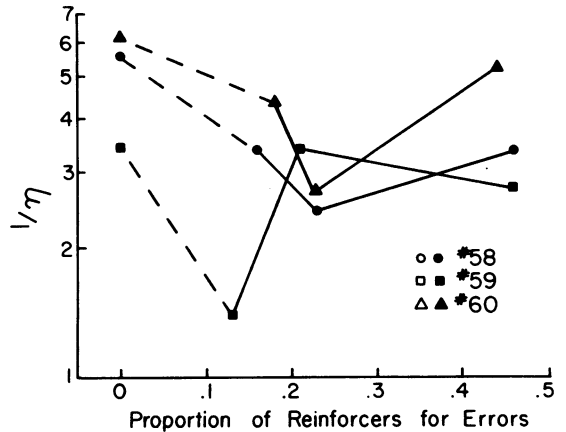


Fig. 4. Discrimination between S_1 and S_2 , indexed by $1/\eta$, as a function of the proportion of reinforcers obtained for errors in Conditions 14 to 17.

source of bias in these three conditions with asymmetrical error-reinforcement conditions.

In summary, weakening the signal-detection contingency by introducing reinforcers for errors generally decreased discrimination, although the effects were less clear under asymmetrical error-reinforcement conditions. Orthogonal variation of the allocation of reinforcers to B_1 or B_2 produced shifts in response bias that were similar at all values of the signal-detection contingency.

EXPERIMENT II

As suggested in the introduction, the contingencies of a signal-detection experiment may be weakened either by reinforcing errors, as in Experiment I, or by withholding reinforcers for correct responses. In Experiment II, the effects of the latter method were explored by varying the probability or rate of reinforcement for correct responses while errors were never reinforced. Nevin (Note 2) demonstrated that discrimination depended on the number of reinforcers obtained per trial in a discrete-trial analog to multiple schedules of reinforcement, in which the effects of differential reinforcement should be especially clear (e.g., Shimp & Wheatley, 1971; Todorov, 1972; see de Villiers, 1977, for review). Experiment II is a full report of Nevin's (Note 2) study.

METHOD

Subjects and Apparatus

As in Experiment I.




Fig. 5. Bias toward B_1 or B_2 , indexed by b , as a function of the ratio of reinforcers obtained by B_1 to reinforcers obtained by B_2 , for all conditions of Experiment I.

Procedure

The basic duration-discrimination procedure was identical to that employed in Experiment I. On 50% of the trials, unrelated to the duration of the white key light, a tone came on with onset of the center key and remained on until the end of the trial, whereas a noise came on during the other 50% of the trials. Over the course of the experiment, the method of scheduling reinforcers and the overall rate of reinforcement varied, but the scheduled rate or probability of reinforcement was always five times greater on tone than on noise trials. In Condition 1, every correct response was reinforced on tone trials, whereas correct responses on noise trials were reinforced with probability .2. In Condition 2, correct responses on tone trials were reinforced according to a VI 45-sec schedule, whereas correct responses on noise trials were reinforced according to a VI 225-sec schedule. In Condition 3, correct responses on tone trials were reinforced with probability .2, whereas correct responses on noise trials were reinforced with probability .04. These latter values were chosen to provide approximately the same rates of reinforcement as the VI schedules in Condition 2, to determine whether the method of scheduling reinforcers affected discrimination independently of rate of reinforcement. Finally, in Condition 4, correct responses on tone trials were reinforced with probability .83, whereas correct responses on noise trials were reinforced with probability .17. Both errors and unreinforced correct responses simply initiated the intertrial interval, as in Experiment I. No attempt was made to control the allocation of reinforcers for the two classes of correct responses, B_w and B_x .

Each of the four conditions remained in effect for 30 sessions, with 60 reinforcers per session.

RESULTS

Inspection of the data for the first condition revealed no systematic changes in discrimination after the first 15 to 20 sessions. Figure 6 shows the means and ranges of discrimination performance, averaged across subjects, for the final 10 sessions. It also shows how discrimination changed throughout the second condition, again averaged across subjects. There was a large initial decrement in discrimination

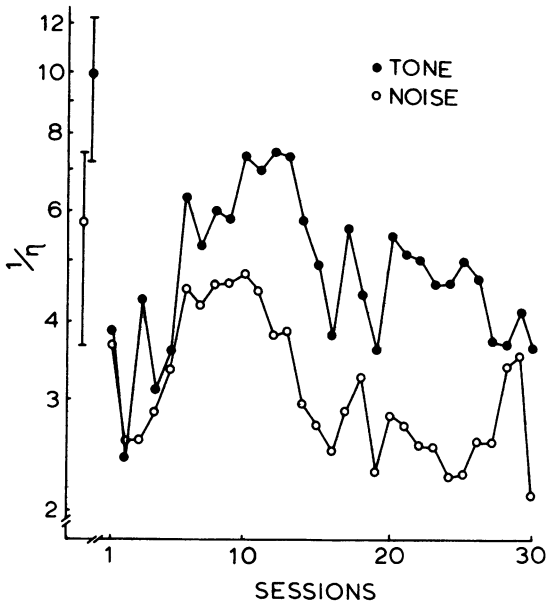


Fig. 6. Discrimination between S_1 and S_2 , indexed by $1/\eta$, for all 30 sessions of Condition 2 of Experiment II. The data are geometric means for the three subjects. Medians and ranges of the average data for the final ten sessions of Condition I are presented at the left. Reinforcement was always five times more frequent on tone than on noise trials. Reinforcement rate on both tone and noise trials in Condition 2 was about one-fifth of that in Condition 1.

with little difference in tone and noise trials during the first two sessions, followed by transient partial recovery. No trends are evident over the final 10 to 15 sessions. Accordingly, the results for all conditions are based on data pooled for the final 10 sessions. Complete data for individual subjects are presented in Table 3.

The value of $1/\eta$ was computed for each subject and condition, separately for tone and noise trials. Figure 7 shows that all subjects exhibited a rough positive relation between $1/\eta$ and the number of reinforcers obtained per trial (which is analogous to obtained reinforcement rate). There was no consistent difference between the effects of the VI schedules in Condition 2, in which reinforcement rate was largely independent of discrimination, and the effects of the probabilistic schedules in Conditions 1, 3, and 4, in which reinforcement rate depended directly on discrimination performance. Within conditions, discrimination accuracy was higher on tone than on noise trials in 11 of 12 comparisons (one exception for Bird 59). Thus, accuracy of discrimination

Table 3

Responses emitted, and reinforcers obtained, summed separately for tone and noise trials for the last 10 sessions of Experiment II.

Condition		B_w	B_s	B_y	B_z	R_w	R_s
Bird 58	1 Tone	277	8	56	217	277	217
	1 Noise	240	23	57	195	72	35
	2 Tone	1063	151	263	1034	254	249
	2 Noise	1042	46	637	628	64	33
	3 Tone	822	232	48	992	243	290
	3 Noise	646	272	164	823	38	29
	4 Tone	264	25	28	345	231	279
	4 Noise	253	104	57	317	29	61
Bird 59	1 Tone	250	48	61	243	250	243
	1 Noise	306	37	102	297	68	36
	2 Tone	759	440	220	1186	173	336
	2 Noise	1022	159	883	448	64	27
	3 Tone	738	314	281	977	217	311
	3 Noise	1030	18	1183	127	67	5
	4 Tone	204	160	48	397	175	327
	4 Noise	313	83	148	314	44	54
Bird 60	1 Tone	249	13	22	239	249	239
	1 Noise	248	43	42	206	56	37
	2 Tone	945	284	229	1074	254	260
	2 Noise	982	202	806	486	57	30
	3 Tone	859	125	267	920	226	293
	3 Noise	669	301	359	745	39	42
	4 Tone	267	14	43	337	224	275
	4 Noise	286	78	70	303	47	54

was positively related to reinforcement rate both across conditions and across signaled components within conditions.

DISCUSSION

Experiment I demonstrated that discrimination was a decreasing function of the proportion of reinforcers for responses conventionally scored as errors, regardless of bias, except that discrimination did not decrease systematically with the proportion of reinforcers for errors when only one class of errors was reinforced. In addition, bias toward one or the other side key was proportional to the ratio of reinforcers obtained by pecks at those keys, regardless of the level of discrimination. Experiment II demonstrated that discrimination was an increasing function of overall rate of reinforcement, both within and between experimental conditions. We now attempt a unified account of the discrimination data in relation to the contingencies of reinforcement.




Fig. 7. Discrimination between S_1 and S_2 , indexed by $1/\eta$, as a function of the number of reinforcers obtained per trial in Experiment II. The data are coded according to trial type and method of reinforcement scheduling. Lines connect pairs of points determined within each condition.

To achieve such an account, we will characterize the contingency by the average difference in obtained probabilities of reinforcement per response in the pairs of diagonally opposed cells of the payoff matrix, which represent correct responses and errors. Specifically, this measure is computed as:

$$\text{Contingency} = \frac{\left(\frac{R_w}{B_w} + \frac{R_z}{B_z}\right) - \left(\frac{R_x}{B_x} + \frac{R_y}{B_y}\right)}{2} \quad (4)$$

This contingency measure may be specified independently of the subject's behavior whenever reinforcement probabilities per response are arranged by the experimenter, as in Experiment II, Conditions 1, 3, and 4. When reinforcers are arranged intermittently on an

interval-like basis, as in Experiment II, Condition 2, and all conditions of Experiment I, the subject's behavior affects the experienced asymptotic probabilities. For example, if reinforcers are assigned for correct responses only, improvements in discrimination decrease the obtained probability of reinforcement per response because the interval program sets an upper limit on the number of reinforcers that can be obtained. If discrimination is a positive function of reinforcement probability for correct responses, a negative feedback relation would operate to stabilize discrimination at some level that would depend on the confusability of the stimuli and the properties of the function relating discrimination to reinforcement probability.

The varying discrimination levels shown in Figure 6 may exemplify the operation of the negative feedback relation. The increase in discrimination performance during early sessions would lower obtained probability of reinforcement, which may in turn be responsible for later decreases to intermediate levels.

A still more complicated case arises when reinforcers are assigned to all four cells, as in most conditions of Experiment I. If the stimuli are perfectly distinguishable, matching to reinforcers in all four cells may occur. For example, in Condition 1, reinforcers were assigned in the following ratios:

$$\text{Given } S_1, \frac{R_w}{R_x} = \frac{4}{1};$$

$$\text{Given } S_2, \frac{R_y}{R_z} = \frac{1}{4}$$

If the subjects allocate responses so as to match reinforcement ratios, the resulting measure of discrimination, $1/\eta$, would be 4.0. However, the obtained probabilities of reinforcement in each cell would be identical, so the measure of contingency would be 0 (cf. Gibbon *et al.*, 1974; Revusky, 1963). By comparison, if the stimuli were indistinguishable, responses would be equally frequent in each cell and $1/\eta$ would be 1.0. However, in view of the reinforcement assignment ratios, the obtained probabilities of reinforcement for B_w and B_z would be four times greater than those for B_x and B_y , and the contingency measure would be positive. Variations in discrimination between these limits would lead to a negative feedback relation between discrimination and

contingency, as in the simpler case described above.

In view of the forced negative relation between discrimination and contingency that must exist within conditions when reinforcers are scheduled on an interval basis, it is of more than passing interest that the results exhibit a clear positive relationship across conditions. Moreover, the relationship incorporates the data for conditions in which contingency was experimentally determined. Figure 8 shows how the average value of $1/\eta$ varied with the difference between obtained probabilities of reinforcement for correct responses, B_w and B_z , and errors, B_x and B_y , computed according to Equation 4. The obtained probabilities were determined by pooling responses and reinforcers across subjects for each experimental condition. The data for the first 12 conditions of Experiment I, and for all conditions of Experiment II, exhibit nonsystematic variation around an increasing, negatively accelerated function. The data of Conditions 14 to 16 of Experiment I, in which discrimination did not vary systematically with the proportion of reinforcers for errors, converge on an essentially constant difference in obtained reinforcement probabilities. It is noteworthy that discrimination remains above chance ($1/\eta > 1$) even when contingency is zero or negative. This persistence of better-than-chance performance may be attributed to the fact that reinforcers are always obtained more often per

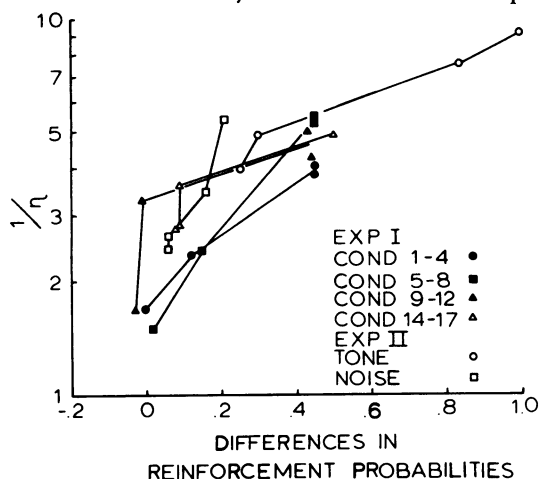


Fig. 8. Discrimination between S_1 and S_2 , indexed by $1/\eta$, as a function of the average difference between the obtained probabilities of reinforcement for correct responses and errors. The data are geometric means across subjects for all conditions of Experiments I and II.

trial for correct responses than for errors, even though the obtained probabilities of reinforcement per response differ only trivially.

The results shown in Figure 8 are entirely consistent with previous data on the effects of reinforcement for errors. Using rats as subjects in a luminance discrimination task, Nevin et al. (1975) varied the probabilities of water reinforcement per response for correct responses and errors in the three series of determinations. Discrimination was a decreasing function of the ratio of reinforcement probabilities, remaining above chance when the ratio was 1.0 (i.e. the probability difference was 0). Using pigeons as subjects in a temporal discrimination task, Davison and McCarthy (1980) reinforced correct responses with probability .7, and showed that discrimination decreased as the probability of reinforcement for errors increased from 0 to .9. All subjects performed above chance when the reinforcement probabilities were the same. The data from these studies are replotted in Figure 9, in a form that parallels Figure 8. The similarities between the functions for Experiment I in Figure 8 and those in Figure 9 suggest that the effects of reinforcement for errors are similar across stimuli, reinforcers, and species, regardless of whether reinforcers are arranged probabilistically or by interdependent interval-like schedules.

Another relevant study was reported by Nevin, Mandell, and Whittaker (1978). They trained pigeons on a discrete-trial choice pro-

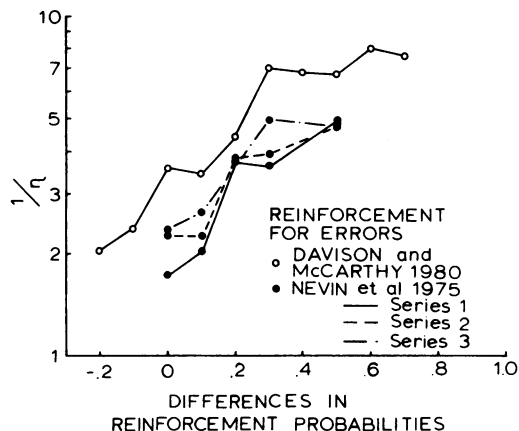


Fig. 9. Discrimination between S_1 and S_2 , indexed by $1/\eta$, as a function of the difference in probabilities of reinforcement for correct responses and errors. The data are geometric means for six pigeons (Davison & McCarthy, 1980) or three rats (Nevin et al., 1975).

cedure that is formally equivalent to a signal-detection procedure and is readily described by reference to the contingency cube (Figure 1). On a random half of the trials, both keys were lighted red (S_1) and left-key pecks (B_1) were reinforced three times as often as right-key pecks (B_2). On the other half of the trials, both keys were lighted white (S_2), and reinforcement frequencies were varied across conditions. In one condition, for example, in 100 red-key trials, left pecks (B_w) were reinforced 12 times, and right pecks (B_x) were reinforced four times. In 100 white-key trials, left pecks (B_y) were never reinforced, whereas right pecks (B_z) were reinforced eight times. When white-key reinforcement conditions were changed so that B_z was always reinforced, B_z/B_y increased. At the same time, B_x/B_w decreased even though the red-key schedules remained constant. Nevin *et al.* (1978) interpreted this result as an instance of behavioral contrast. However, it can also be seen as an instance of control by the contingencies of reinforcement, defined as the difference in reinforcement probabilities.

If B_w and B_z are construed as correct responses, and B_x and B_y as errors in a signal-detection paradigm, the overall procedure is much like that of Experiment 1, Conditions 14 to 16, in which only one class of errors was reinforced. In the conditions described above, increasing the probability of reinforcement for B_z increases the contingency measure, and discrimination improves accordingly: that is, both B_z/B_y and B_w/B_x increase, with the consequence that $1/\eta$ for the red-white discrimination increases markedly. Thus, our so-called contrast effect is readily interpreted as a contingency effect on discrimination. The same analysis may be applied to all other conditions of the Nevin *et al.* (1978) experiment for which $R_w + R_z > R_x + R_y$, and the results are all consistent with the contingency interpretation. Thus, the contingency between stimuli, responses, and reinforcers characterized by the reinforcement probability difference provides a unified account of a number of findings.

Figure 8 also shows that the difference in reinforcement probabilities for correct responses and errors accounts for the data when only correct responses were reinforced in Experiment II. Here, however, there is less agreement across experiments. McCarthy (Note 1) reported that discrimination performance in a signal-detection

task was essentially unaffected by moderate changes in the obtained probability of reinforcement that resulted from the use of variable-interval schedules. Nevin (1967) observed only small decrements in simultaneous-discrimination performance when reinforcement probabilities were reduced to low levels, including extinction. It is clear that empirical study of variables affecting the relation between reinforcement probability and discrimination is required, and that quantitative characterization of the functions in Figures 8 and 9 is premature.

Signal Detection and Matching

We turn now to a model of detection performance based on the matching law, which has provided a satisfactory account of signal-detection performance to date. However, the model cannot encompass some of our data as it is presently formulated.

Davison and Tustin (1978) and Nevin *et al.* (Note 3) independently proposed quite similar accounts of signal-detection performance based on the matching principle. Davison and McCarthy (1980) have extended the Davison-Tustin approach to the case in which errors are reinforced, so we will describe their treatment here. For review and discussion of matching-law approaches to signal-detection performance, see McCarthy and Davison (1981) and Nevin (1981).

Davison and Tustin (1978) proposed that the allocation of responses in a signal-detection experiment obeyed the generalized matching law (Baum, 1974) with parameters reflecting inherent bias and stimulus bias. Using the notation above, they proposed two equations, in logarithmic form, for S_1 and S_2 .

Given S_1 :

$$\log \left(\frac{B_w}{B_x} \right) = a \log \left(\frac{R_w}{R_z} \right) + \log d + \log c. \quad (5a)$$

Given S_2 :

$$\log \left(\frac{B_y}{B_z} \right) = a \log \left(\frac{R_w}{R_z} \right) - \log d + \log c \quad (5b)$$

where a represents the sensitivity of response allocation to reinforcer allocation, and $\log c$ represents inherent bias toward B_1 or B_2 . Their major innovation was to treat stimulus effects as equivalent to bias: presentation of S_1 biases

responding toward B_1 , and presentation of S_2 biases responding toward B_2 . Stimulus bias is represented by $\log d$, which operates to shift response ratios in opposite directions depending on the stimulus. Its value depends on the physical difference between S_1 and S_2 , going to 0 ($d = 1$) when the stimuli are indistinguishable.

Subtracting Equation 5b from 5a, rearranging, antilogging, and taking square roots, we obtain

$$d = \sqrt{\frac{B_w}{B_x} \cdot \frac{B_z}{B_y}} \quad (6)$$

which is exactly equivalent to Equation 1, where $d = 1/\eta$. Adding Equations 5a and 5b, rearranging, antilogging, and taking square roots, we obtain

$$c \left(\frac{R_w}{R_z} \right)^a = \sqrt{\frac{B_w}{B_x} \cdot \frac{B_y}{B_z}} \quad (7)$$

which is exactly equivalent to Equation 2, where

$$b = c \left(\frac{R_w}{R_z} \right)^a.$$

Thus, the sensitivity and bias indices proposed by Luce (1959, 1963) are derivable from the generalized matching law.

Davison and McCarthy (1980) incorporated the effects of error reinforcement in their formulation by proposing that $\log d$, the stimulus bias term, be weighted by the proportion of reinforcers that were differentially allocated with respect to the stimuli. Thus,

Given S_1 :

$$\begin{aligned} \log \left(\frac{B_w}{B_x} \right) &= a \log \left(\frac{R_w + R_y}{R_x + R_z} \right) \\ &+ \left(\frac{R_w + R_z - R_x - R_y}{R_w + R_z + R_x + R_y} \right) \log d + \log c. \end{aligned} \quad (8a)$$

Given S_2 :

$$\begin{aligned} \log \left(\frac{B_y}{B_z} \right) &= a \log \left(\frac{R_w + R_y}{R_x + R_z} \right) \\ &- \left(\frac{R_w + R_z - R_x - R_y}{R_w + R_z + R_x + R_y} \right) \log d + \log c. \end{aligned} \quad (8b)$$

These equations suggest that the influence of $\log d$ goes to zero (i.e., discrimination deteriorates) as R_x and R_y , the rates of reinforcement for errors, approach R_w and R_z , the rates of

reinforcement for correct responses. All extant data on this question accord with this prediction except for those of Conditions 14 to 16 of the present Experiment I, which were not available to Davison and McCarthy when they prepared their 1980 report. In those three conditions, where only one class of errors was reinforced, discrimination did not decrease systematically with increases in the proportion of reinforcers for errors as their model would predict. This difficulty cannot be remedied by replacing their term $(R_w + R_z - R_x - R_y/R_w + R_x + R_x + R_y)$ in Equations 8a and 8b by the contingency measure based on reinforcement probabilities, computed as above, because discrimination remains well above chance even when the contingency is equal to or less than zero. In order to accommodate these data into a model of detection performance based on matching, terms referring to both reinforcement rate and reinforcement-probability differences will be required.

The Davison-McCarthy model for error reinforcement also has difficulty with the findings of Nevin et al. (1978), whose procedure resembled that of Conditions 14 to 16 of Experiment I (see discussion above). Although their model provides a rough description of the data (see Davison and McCarthy, 1980, pp. 44, 46), it predicts that the ratio B_w/B_x will decrease when B_y/B_z decreases because of a large increase in R_z , which is the reverse of the finding described above. Moreover, the value of $\log d$ estimated by their model for the red-white discrimination studied by Nevin et al. (1978) is not much greater than for the temporal discrimination employed in the present experiments, and in fact it underestimates obtained red-white discrimination performance in some conditions of that study by a factor of two or more. These difficulties might be resolved by using different weighting terms in Equations 8a and 8b, but the alternatives are not yet well-defined.

As presently formulated, the Davison-Tustin and Davison-McCarthy models predict that variation in the absolute rate of reinforcement will have no effect on discrimination so long as reinforcement ratios remain constant. Experiment II demonstrated a systematic effect of absolute rate of reinforcement on discrimination, a result which can be incorporated into their model in either of two ways. First, the contingency measure employed here can be

applied as a weighting term on $\log d$ in their Equations 5a and 5b. The contingency term would itself have to be weighted by a free parameter, because reductions in the probability of reinforcement for correct responses do not have uniform effects across experiments, as discussed above. Alternatively, a term for extraneous reinforcement could be added to each cell of the payoff matrix, to represent the value of terminating a trial and/or advancing the reinforcement program. Such a term would constitute a free parameter analogous to the value of reinforcement for other behavior in Herrnstein's (1970) formulation of the law of effect. However, the addition of a free parameter of this sort greatly complicates the analysis, and without the identification of experimental variables that determine its value, there would be little gain in understanding. Many more data are required to ascertain how best to incorporate variations in reinforcement rate or probability into these generalized matching law equations, or whether an alternative formulation explicitly based on the notion of contingency may prove more parsimonious.

Summary

The analysis of behavior in the signal-detection paradigm suggests that two independent parameters suffice to characterize performance. One parameter, identified with discrimination, depends on the physical differences between stimuli (e.g. McCarthy & Davison, 1980a; Wright, 1972) and on the extent to which reinforcers for correct responses and errors are scheduled differentially, as shown here and in earlier related work. The degree of differential reinforcement may be characterized adequately by a simple measure of contingency.

The second parameter, identified with response bias, depends on the ratio of reinforcers for the two alternative responses. The matching relation between response bias and the reinforcement ratio is independent of the level of discrimination, regardless of whether discrimination is controlled by physical stimulus differences (e.g., McCarthy & Davison, 1980b) or by the contingency between stimuli, responses, and reinforcers as here. An account of discrimination in relation to the contingency suggests an alternative or a supplement to accounts based on the generalized matching law, and provides a broad conceptual framework for the analysis of related problem areas,

as summarized by the contingency cube relating the three terms of the discriminated operant.

REFERENCE NOTES

1. McCarthy, D. Research presentation, University of New Hampshire, July 1980.
2. Nevin, J. A. Accuracy of discrimination: Effects of reinforcement rate and distraction. Paper presented at the meetings of the Eastern Psychological Association, Philadelphia, Pa., April 1979.
3. Nevin, J. A., Jenkins, P., Whittaker, S., & Yarensky, P. Signal detection and matching. Paper presented at the meetings of the Psychonomic Society, Washington, D. C., November 1977.

REFERENCES

- Baum, W. M. On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, 1974, **22**, 231-242.
- Catania, A. C. Elicitation, reinforcement, and stimulus control. In R. Glaser (Ed.), *The nature of reinforcement*. New York: Academic Press, 1971.
- Davison, M. C., & McCarthy, D. Reinforcement for errors in a signal-detection procedure. *Journal of the Experimental Analysis of Behavior*, 1980, **34**, 35-47.
- Davison, M. C., & Tustin, R. D. The relation between the generalized matching law and signal-detection theory. *Journal of the Experimental Analysis of Behavior*, 1978, **29**, 331-336.
- de Villiers, P. A. Choice in concurrent schedules and a quantitative formulation of the law of effect. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior*. New York: Prentice-Hall, 1977.
- Gibbon, J., Berryman, R., & Thompson, R. L. Contingency spaces and measures in classical and instrumental conditioning. *Journal of the Experimental Analysis of Behavior*, 1974, **21**, 585-605.
- Green, D. M., & Swets, J. A. *Signal detection theory and psychophysics*. New York: Wiley, 1966.
- Herrnstein, R. J. On the law of effect. *Journal of the Experimental Analysis of Behavior*, 1970, **13**, 243-266.
- Luce, R. D. *Individual choice behavior*. New York: Wiley, 1959.
- Luce, R. D. Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 1). New York: Wiley, 1963.
- McCarthy, D., & Davison, M. C. On the discriminability of stimulus duration. *Journal of the Experimental Analysis of Behavior*, 1980, **33**, 187-211. (a)
- McCarthy, D., & Davison, M. C. Independence of sensitivity to relative reinforcement rate and discriminability in signal detection. *Journal of the Experimental Analysis of Behavior*, 1980, **34**, 273-284. (b)
- McCarthy, D., & Davison, M. C. Matching and signal detection. In M. L. Commons & J. A. Nevin (Eds.), *Quantitative analyses of behavior* (Vol. 1). *Discrimi-*

- native properties of reinforcement schedules.* Cambridge: Ballinger, 1981.
- Nevin, J. A. Effects of reinforcement scheduling on simultaneous discrimination performance. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 251-260.
- Nevin, J. A. On differential stimulation and differential reinforcement. In W. C. Stebbins (Ed.), *Animal psychophysics*. New York: Appleton-Century-Crofts, 1970.
- Nevin, J. A. Psychophysics and reinforcement schedules: An integration. In M. L. Commons & J. A. Nevin (Eds.), *Quantitative analyses of behavior* (Vol. 1). *Discriminative properties of reinforcement schedules.* Cambridge: Ballinger, 1981.
- Nevin, J. A., Mandell, C., & Whittaker, S. Contrast and induction in multiple schedules of discrete-trial concurrent reinforcement. *Journal of the Experimental Analysis of Behavior*, 1978, **30**, 53-61.
- Nevin, J. A., Olson, K., Mandell, C., & Yarensky, P. Differential reinforcement and signal detection. *Journal of the Experimental Analysis of Behavior*, 1975, **24**, 355-367.
- Revusky, S. H. A relationship between responses per reinforcement and preference during concurrent VI VI. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 518.
- Shimp, C. P. Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 443-455.
- Shimp, C. P., & Wheatley, K. L. Matching to relative reinforcement frequency in multiple schedules with a short component duration. *Journal of the Experimental Analysis of Behavior*, 1971, **15**, 205-210.
- Skinner, B. F. *Contingencies of reinforcement: A theoretical analysis.* New York: Appleton-Century-Crofts, 1969.
- Stubbs, D. A. Response bias and the discrimination of stimulus duration. *Journal of the Experimental Analysis of Behavior*, 1976, **25**, 243-250.
- Stubbs, D. A., & Pliskoff, S. S. Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 887-895.
- Todorov, J. C. Component duration and relative response rates in multiple schedules. *Journal of the Experimental Analysis of Behavior*, 1972, **17**, 45-49.
- Wright, A. A. Psychometric and psychophysical hue discrimination functions for the pigeon. *Vision Research*, 1972, **12**, 1447-1464.

Received October 3, 1980

Final acceptance July 9, 1981