

MAGNITUDE AND FREQUENCY OF REINFORCEMENT AND FREQUENCIES OF INTERRESPONSE TIMES¹

CHARLES P. SHIMP

UNIVERSITY OF UTAH

The relative magnitude and relative frequency of reinforcement for two concurrent interresponse times (1.5 to 2.5 sec and 3.5 to 4.5 sec) were simultaneously varied in an experiment in which pigeons obtained grain by pecking on a single key. Visual discriminative stimuli accompanied the two time intervals in which reinforcements were arranged by a one-minute variable-interval schedule. The resulting interresponse times of each of three pigeons fell into two groups; "short" (1.0 to 2.5 sec) and "long" (3.0 to 4.5 sec). Steady-state relative frequencies of these interresponse times were orderly functions of both reinforcement variables. The combined effects of both independent variables were well summarized by a linear function of one variable, relative access to food. Unlike corresponding two-key concurrent variable-interval schedules, the present schedule did not produce an equality between the relative frequency of an operant and either the relative magnitude or the relative frequency of reinforcement of that operant. A tentative account is provided for this difference between one-key and two-key functions.

A basic problem in the experimental analysis of behavior is to determine the functions relating relative frequencies of concurrent operants to the relative magnitudes and relative frequencies of reinforcement for those operants. Here the problem was attacked with an experiment using variable-interval reinforcement and pigeons. Two different approaches to the solution of the problem correspond to the nature of the concurrent operants. First, an operant may be defined as a peck on one of several keys. Second, an operant may be defined as one of several classes of interresponse times (IRTs) between successive pecks on a single key. The first approach has been remarkably successful; it has led to the discovery that the function in question is a matching function. That is, in concurrent variable-interval schedules of reinforcement, the relative frequency of pecks on one of two or three keys approximately equals the relative frequency, or relative magnitude, of reinforcement for pecks on that key (Herrnstein, 1961; Catania, 1963; Reynolds, 1963; Neuringer, 1967).

The second approach, that in which an operant is a class of IRTs, has been less successful. Anger (1954) asked if the relative frequencies of occurrence of the different IRTs in a variable-interval schedule depended on the corresponding relative frequencies of reinforcement. To answer this question he devised a "synthetic" variable-interval schedule; a concurrent schedule of reinforcement for different IRTs. His results suggested that this dependency may exist sometimes but also that sometimes it may seem too small to account for the development of normal variable-interval behavior. Specifically, Anger showed that a decrease in reinforcement frequency for one band of IRTs decreased the frequency of occurrence of responses terminating in that band, but that a corresponding increase in reinforcement frequency for another band had no effect. Thus, it is not clear how much the relative frequencies of IRTs depend on the relative frequencies of reinforcement for them. But in any case, this dependency is not the same as the matching function in two-key concurrent variable-interval schedules. Before now this second approach has not been used with reinforcement magnitude, although Hendry (1962) has shown in a different context

The second approach, that in which an operant is a class of IRTs, has been less successful. Anger (1954) asked if the relative frequencies of occurrence of the different IRTs in a variable-interval schedule depended on the corresponding relative frequencies of reinforcement. To answer this question he devised a "synthetic" variable-interval schedule; a concurrent schedule of reinforcement for different IRTs. His results suggested that this dependency may exist sometimes but also that sometimes it may seem too small to account for the development of normal variable-interval behavior. Specifically, Anger showed that a decrease in reinforcement frequency for one band of IRTs decreased the frequency of occurrence of responses terminating in that band, but that a corresponding increase in reinforcement frequency for another band had no effect. Thus, it is not clear how much the relative frequencies of IRTs depend on the relative frequencies of reinforcement for them. But in any case, this dependency is not the same as the matching function in two-key concurrent variable-interval schedules. Before now this second approach has not been used with reinforcement magnitude, although Hendry (1962) has shown in a different context

¹This work was done at Stanford University and was supported by the author's PHS Postdoctoral Fellowship and by NICHD grant HD-00954, Gordon H. Bower, principal investigator. The author would like to thank Dr. Bower for his loan of facilities, Mr. James McLean for helping to set up the equipment, and Linda A. Shimp for her constructive editorial criticism. Reprints may be obtained from the author, Dept. of Psychology, Annex 150, The University of Utah, Salt Lake City, Utah 84112.

that differential reinforcement magnitude can affect an IRT distribution.

In the present study, IRTs were selected as the concurrent operants to receive different relative frequencies and relative magnitudes of reinforcement. The present schedule was similar to Anger's synthetic variable-interval schedule, but the interpretation of the results was made easier by removing the effects of induction and by using only two classes of reinforced IRTs. Since the number of concurrent operants was two, the present one-key data could be compared with the two-key data described above.

METHOD

Subjects

Three male White Carneaux pigeons were maintained at 80% of their free-feeding body weights. The birds had served previously in an experiment on the differential reinforcement of short IRTs (Shimp, 1967).

Apparatus

Interresponse times were recorded with a Foringer Multiple Class Time Analyzer stepped by an electronic timer. A Lehigh Valley Electronics pigeon key, operated by a minimal force of 22 g, was mounted in a standard experimental space.

Procedure

Stimuli. Three different stimuli could appear on the translucent response key, depending on the time since the preceding response. Specifically, the key was dark with a white, vertical bar (S_1) when at least 1.5 sec but not more than 2.5 sec had elapsed since the preceding response. The key was dark with a white, horizontal bar (S_2) between 3.5 and 4.5 sec after a response. The key appeared yellow at all other times. The purpose of the visual stimuli was to reduce the frequency of responses outside either reinforced band. The sequence of visual stimuli appears in Fig. 1. It was thought that visual feedback might help to stabilize response topography and to reduce the frequency of very short IRTs. Thus, a peck at a yellow key darkened the key for approximately 0.20 sec, and, of course, a peck at either S_1 or S_2 produced a yellow key. The key was dark during reinforcement and then, starting from the end of each reinforcement, the

stimuli were presented as if a response had just occurred, that is, as shown in Fig. 1.




Fig. 1. Temporal sequence of visual stimuli on the response key. A response reset the sequence to the beginning of band 1. Reinforcements were available only in the presence of S_1 or S_2 .

Reinforcing contingencies. The variable-interval schedule of reinforcement used here was one for which the "reinforcements per opportunity" (Catania and Reynolds, 1968) were approximately constant at 0.1 every 8 sec, except at long times after reinforcement. The average interreinforcement interval was 1.0 min. When a reinforcement was assigned by the punched tape, it stopped until the reinforcement was delivered. If a reinforcement became available for a peck in, for example, S_1 , then pecks in S_2 went unreinforced until the reinforcement of a peck in S_1 enabled the tape to move again. Only responses terminating IRTs in S_1 or S_2 were reinforced. In summary, after a reinforcement became available, it was delivered for a response in the presence of either S_1 or S_2 , it was never simultaneously available for both stimuli at once, and it had to be collected before the next assignment could be made. The two reinforcing events followed a random sequence. The relative frequency of reinforcement for responses in S_1 equals the number of reinforcements for responses in S_1 divided by the sum of the number of reinforcements for responses in S_1 plus the number of reinforcements for responses in S_2 . This relative frequency will be denoted by π . It equals the relative time rate of reinforcement for responses in S_1 . As a consequence of the present method of scheduling reinforcements, the obtained relative frequency of reinforcement equalled the scheduled relative frequency, so long as a bird responded at least occasionally to both stimuli. That is, the relative frequency of reinforcement was controlled mainly by the experimenter: within broad limits, a bird's behavior did not affect it.

In addition to different reinforcement frequencies, the experiment employed different reinforcement magnitudes. Responses in S_1 and S_2 could be reinforced with different feeder durations. The relative feeder duration for responses in S_1 equals the feeder duration for responses in S_1 divided by the sum of the feeder duration for responses in S_1 plus the feeder duration for responses in S_2 . This relative duration will be denoted by F .

Recording. Each response terminating an IRT was recorded in one of 11 bands of IRTs. The bandwidth was 0.5 sec for all but the eleventh band, which includes all IRTs longer than 5 sec. A response terminating a latency after reinforcement was not recorded. Also, every response occurring before the tenth reinforcement was discarded to avoid recording warm-up effects.

Training and experimental conditions. The experimental conditions listed in Table 1 were preceded by six weeks of preliminary training, during which the bandwidth was selected and the birds learned to respond mostly in the presence of S_1 or S_2 . Experimental sessions lasted 1 hr and were conducted an average of six days a week. Bird 3 ceased to respond more than a few dozen times per session during the fifth experimental condition. Instead of rejecting the bird, the experimenter made various attempts to reinstate responding. The successful attempt consisted of replacing the usual reinforcement food (wheat, kafir, vetch) with

home-cage food (corn, milo maize, kafir, Austrian peas). Thus, while Birds 1 and 2 continued to receive the usual reinforcement food, Bird 3, beginning with the eighth experimental condition, received only home cage food. Once Bird 3 began to respond again, it behaved the same as before it stopped responding.

RESULTS

Table 2 gives for each bird the frequencies, summed over the last two days of each condition, of the IRTs in each of the 11 recorded bands. As can be seen from Fig. 2, which shows some representative IRT distributions, the distributions are bimodal, with each mode near one of the two bands of reinforced IRTs. The figure suggests that the pecking behavior of a bird may be summarized by the relative frequencies of two groups of IRTs; aside from band 1, the only bands with non-negligible frequencies form two distinct groups: 3, 4, 5 and 7, 8, 9. These two groups will be denoted for the sake of brevity by "short" and "long". Responses in band 1, *i.e.*, responses terminating IRTs less than 0.5 sec, are omitted from the discussion below because they apparently enter into different functional relationships with reinforcement and other stimulus variables (Blough, 1963, 1966). Observe in Fig. 2 that the relative frequencies of responses in the unreinforced bands 3 and 7 are typically

Table 1
Sequence of Experimental Conditions

Condition No.	Feeder Times (Time for Short IRT, Time for Long IRT in Seconds)	Relative Feeder Duration for Short IRT (F)	Relative Freq. of Reinforcement for Short IRT (π)	Relative Access to Food for Short IRT	No. of Days
1	(1.5, 3.0)	0.33	0.50	0.33	15
2	(1.5, 6.0)	0.20	0.50	0.20	15
3	(1.5, 1.5)	0.50	0.50	0.50	13
4	(6.0, 1.5)	0.80	0.50	0.80	20
5	(6.0, 1.5)	0.80	0.20	0.50	16
6	(1.5, 1.5)	0.50	0.20	0.20	13*
7	(1.5, 6.0)	0.20	0.50	0.20	25*
8	(3.0, 3.0)	0.50	0.50	0.50	14
9	(3.0, 3.0)	0.50	0.85	0.85	10
10	(3.0, 3.0)	0.50	0.70	0.70	12
11	(3.0, 3.0)	0.50	0.20	0.20	10
12	(1.5, 6.0)	0.20	0.20	0.06	10
13	(3.0, 3.0)	0.50	0.50	0.50	11

*Data not available for Bird 3.

Table 2
Frequencies of IRTs¹ over Last Two Days of Each Condition for Each of Three Birds

Experimental Conditions													
1	2	3	4	5	6	7	8	9	10	11	12	13	
BIRD 1													
1	1853	1482	1563	1876	1678	858	1122	1537	1522	1282	520	1573	1782
2	18	6	22	20	14	8	42	2	14	12	4	21	18
3	229	120	130	123	36	30	113	19	69	41	17	100	149
4	821	644	850	1163	928	274	379	726	818	654	160	348	596
5	130	70	81	211	178	81	83	179	165	247	45	131	245
6	26	20	15	31	15	21	52	18	10	10	2	17	24
7	139	174	89	16	12	74	206	52	52	44	44	163	113
8	451	400	322	247	424	267	181	465	338	233	206	639	527
9	21	18	25	20	33	39	13	58	31	23	33	113	91
10	2	0	3	3	3	1	1	4	3	0	0	6	1
11	5	14	8	24	23	51	20	13	20	15	7	26	8
BIRD 2													
1	2038	1980	1826	2749	2769	1332	1712	2225	2908	2809	2145	1564	2238
2	13	15	3	22	25	11	10	13	29	5	3	3	1
3	329	304	183	270	473	74	159	198	407	112	77	18	122
4	756	700	709	1482	1133	363	609	968	1427	1511	704	401	945
5	160	145	236	140	128	94	93	93	144	145	134	135	168
6	17	7	8	11	14	13	17	16	8	22	23	30	17
7	187	59	117	55	173	206	167	85	45	28	87	56	46
8	338	534	331	189	352	296	442	465	179	282	603	563	474
9	117	115	79	68	70	147	72	69	24	79	169	115	75
10	5	5	8	1	3	2	2	3	2	3	3	6	3
11	40	48	40	12	10	51	12	11	21	24	31	16	14
BIRD 3													
1	1979	2166	1817	2016	462			555	3214	2811	2449	1766	2426
2	25	98	36	39	2			7	53	55	59	44	34
3	296	625	501	649	6			119	1114	1122	528	419	667
4	627	693	809	810	37			175	1331	1143	730	296	953
5	169	108	120	186	51			41	185	226	125	50	183
6	29	37	14	13	4			8	29	21	13	5	5
7	419	264	163	81	13			55	45	70	329	473	265
8	264	199	118	116	90			60	38	58	306	363	243
9	63	49	22	31	108			17	14	20	82	45	26
10	7	4	0	2	10			18	1	0	9	0	2
11	34	26	2	34	140			36	21	16	51	18	16

¹Row 1 for each bird gives total response frequencies, while the other rows give IRT frequencies in 0.5-sec bands, except for row 11, which gives the frequencies of all IRTs longer than 5.0 sec.

less than, and never much greater than, those of responses in the two corresponding groups of reinforced bands; *i.e.*, in 4, 5 and 8, 9, respectively. This relationship is different from that in a recent experiment in which non-discriminated IRTs were differentially reinforced (Shimp, 1967). In the earlier experiment, the relative frequency of responses in a band just before a reinforced band was, for reinforced IRTs about equal to those here, usually much greater than in the reinforced band. In both the earlier and the present experiment, the relative frequency of responses in a band just

following a reinforced band was extremely low, reflecting the tendency of a bird to peck near the shorter boundary of a reinforced interval. Apparently then, the stimuli in the present experiment significantly reduced the frequency of responses preceding reinforced intervals and tended to concentrate the responses in the latter two intervals. Since this concentration was far from perfect, however, it seems most reasonable in the following analysis to view an IRT distribution as being composed of the bands with non-negligible frequencies of responses that are controlled by

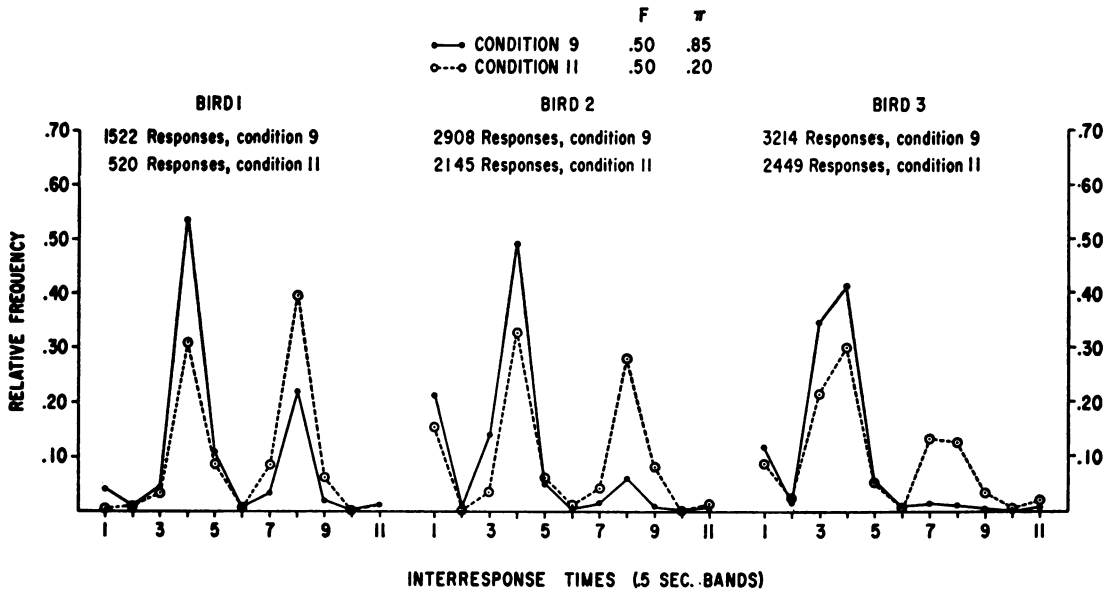


Fig. 2. Representative IRT distributions. The distributions have two distinct, non-overlapping sub-distributions over bands 3, 4, 5, and 7, 8, 9.

external stimuli and by reinforcing contingencies. In short, from Fig. 2 and from the above discussion, it follows that IRTs terminated by responses in bands 3, 4, or 5 may be viewed as instances of one operant, and those terminated by responses in bands 7, 8, or 9 as instances of a second operant. The dependent variable in this experiment is the frequency of key pecks in bands 3, 4, 5, 7, 8, or 9. For brevity, this variable will be denoted by R. The numerical values reported here for R are ratios of the frequencies over the last two days of each experimental condition.

The experiment provided two separate sets of functions. Figure 3 shows the first set, which gives R as a function of F, with π equal to 0.2 and 0.5, in the left and right panels respectively. The value of R increased when the relative magnitude of reinforcement for a short IRT was increased. For Birds 1 and 2, the functions for $\pi = 0.5$ are roughly similar in shape to those for $\pi = 0.2$. The missing data point (see Procedure) prevented a similar assessment for Bird 3. The value of R for this bird was lower in the first experimental condition than the values for either Bird 1 or Bird 2. But the value of R for Bird 3 then increased dramatically (apparently due to a considerable increase in responses in band 3) and

thereafter was consistently greater than the values for either of the other birds. This shift in R can be seen as the atypically low value for Bird 3 at F = 0.33 in the right-hand panel of Fig. 3. The two most important conclusions from Fig. 3 are, then, that reinforcement magnitude controlled R but that when π was equal to 0.5, R did not even approximately equal F.

Figure 4 shows the second set of functions. They reveal that R increased when the relative frequency of reinforcement for a short IRT was increased. The three curves are roughly linear, with the single exception of the curve for Bird 1 at $\pi = 0.85$, and have about the same shape. The different y-intercepts show that Bird 3, when compared to Bird 1 and Bird 2, had a pronounced response bias in favor of short IRTs. Of course, Bird 3's bias is also observable in Fig. 3. The most important conclusion to be drawn from Fig. 4 is that reinforcement frequency controlled R but that R did not even approximately equal π .

Whenever more than one experimental condition gave the same value for the independent variable (*i.e.*, for F or π), the data from these conditions were averaged. Figures 3 and 4 show the averages. Comparisons among various conditions suggest that the averages do not distort the impression gained from plotting individual conditions. For example, two con-

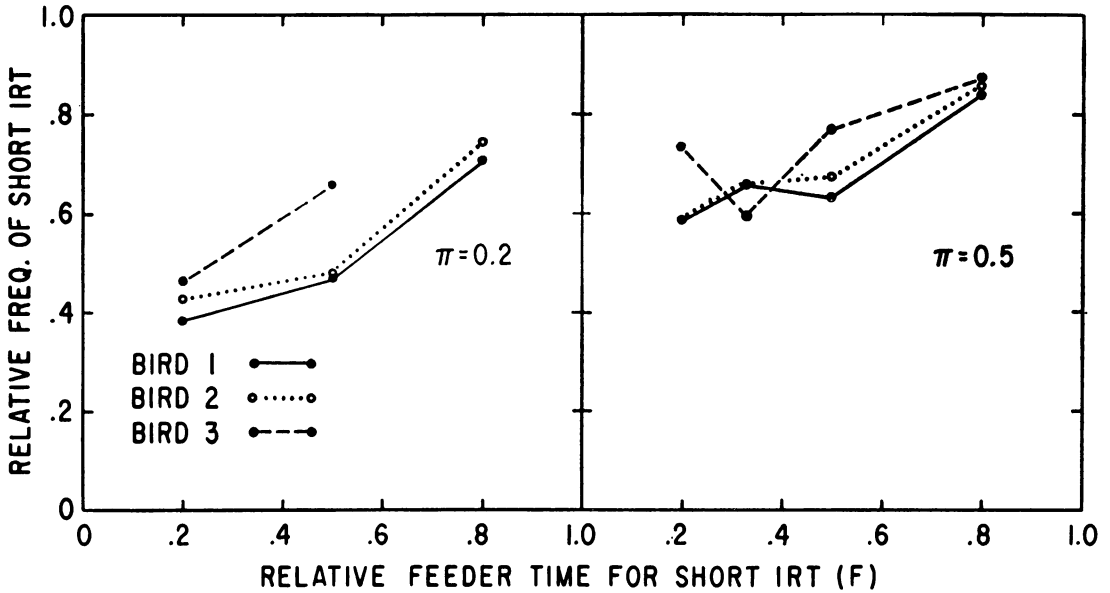


Fig. 3. The relative frequency of a short IRT (shortened to R in the text) as a function of the relative feeder time for that IRT time (F). The relative frequency of reinforcement of this IRT is 0.2 and 0.5 in the left and right panels, respectively. Steady-state values of R systematically increased as F increased.

ditions were replicated to observe the importance of the ordering of experimental conditions. Conditions 7 and 13 replicated conditions 2 and 8, respectively. Complete data from only one of these pairs of conditions were collected from Bird 3, as explained above. The values of R in condition 2 and its replication,

condition 7, were 0.585 and 0.590 for Bird 1, and 0.619 and 0.558 for Bird 2. The values of R in condition 8 and its replication, condition 13, were 0.616 and 0.575 for Bird 1, 0.670 and 0.675 for Bird 2, and 0.717 and 0.771 for Bird 3. The average of the five differences is only 0.033. Therefore the ordering of conditions had little importance, or in other words, the experimental procedures generated recoverable behavior.

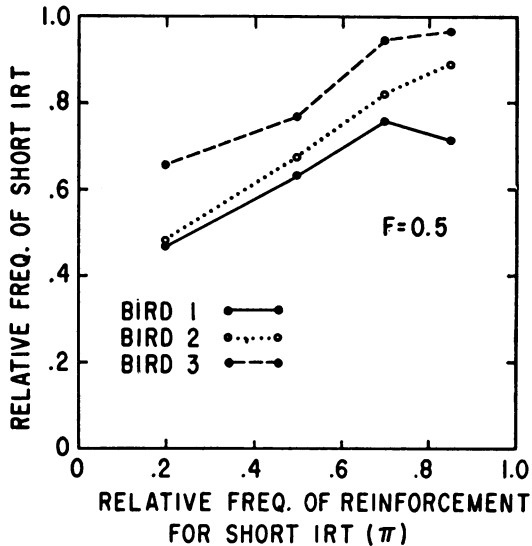


Fig. 4. The relative frequency of a short IRT (R) as a function of the relative frequency of reinforcement for that IRT (π). The relative feeder time of this IRT is 0.5. Steady-state values of R systematically increased as π increased.

A possible complication for an analysis in terms of F, as in Fig. 3, is that different pairs of absolute feeder times, while giving the same value of F, may produce different behavior. In particular, one might expect that, e.g., the pair (3.0 sec, 3.0 sec) of feeder times in conditions 8 and 13 would generate higher values of R than the pair (1.5 sec, 1.5 sec) in condition 3 because a bird in condition 8 or 13 received twice as much food in a session as it did in condition 3. Nevertheless, Table 3 shows that doubling the amount of food available, if it had any effect at all, actually reduced R. The present experiment does not, of course, answer unequivocally whether or not all pairs of feeder times with the same ratio F control the same behavior. However, the present data do suggest that if there is an effect of absolute feeder times, it is small over the range of values studied here.

Table 3

The probability of a short IRT as a function of pairs of feeder times having equal F values but different absolute durations.

feeder times (sec) :	$F = 0.5, \pi = 0.5$			$F = 0.5, \pi = 0.2$	
	(1.5, 1.5) Condition 3	(3.0, 3.0) Condition 8	Condition 13	(1.5, 1.5) Condition 6	(3.0, 3.0) Condition 11
Bird 1	0.709	0.616	0.575	0.503	0.440
Bird 2	0.682	0.670	0.675	0.450	0.516
Bird 3	0.825	0.717	0.771	—	0.659

Figure 5 shows how the relative frequency of the short IRT varied as a function of the relative access to food. The way this latter variable was computed is illustrated by the following example. In condition 12, where π was 0.20, about 20 in 100 reinforcements were for short IRTs. (Since the reinforcement schedule was random this number actually would have fluctuated to some small extent.) Reinforcements were 1.5-sec long for short IRTs and 6.0-sec long for long IRTs. Thus, in 100 reinforcements, the feeder would have been presented about $20 \times 1.5 \text{ sec} = 30 \text{ sec}$ for short IRTs and $80 \times 6.0 \text{ sec} = 480 \text{ sec}$ for long IRTs. The relative access to food, for the short

IRT, would then be $\frac{30 \text{ sec}}{30 \text{ sec} + 480 \text{ sec}} \approx 0.06$.

This number is only a rough estimate of the relative amount of food consumed because it ignores time spent by the bird in getting to the feeder. Figure 5 shows the least-squares best-fitting straight line and the resulting standard error of estimate for each bird. The data plotted in the lower right-hand panel are from all the individual bird-by-condition combinations: the meaning of the dashed line in this panel is explained in the Discussion. Figure 5 suggests that the way in which relative frequency of reinforcement and relative magnitude of reinforcement combine to control behavior is fairly well summarized by a linear function of relative access to food. In fact, the data points appear better described by linear functions than those in Fig. 3, which shows one of the two functions that combine to form relative access to food. Perhaps the most systematic deviation from linearity in Fig. 5 results from a too-low relative frequency of the short IRT when the relative access to food for that IRT was 0.06. This deviation is not surprising, since presumably R decreases rapidly to

zero when the relative access to food closely approaches zero.

The results may be summarized as follows. The slopes of the functions relating R, the relative frequency of the short IRT, to relative reinforcement frequency and relative reinforcement magnitude are greater than zero but less than one. That is, an increase of a given magnitude in either relative reinforcement frequency or relative reinforcement magnitude for an IRT produced an increase of smaller magnitude in the relative frequency of that IRT. A linear function of relative access to food may describe how the effects of reinforcement frequency and reinforcement magnitude combine.

DISCUSSION

The principal topics to be discussed are: (1) the relationships between the experiments on reinforcement magnitude which have generated either no effect or a matching effect, and on the other hand, the present experiment, which generated an intermediate effect, and (2) the analogous relationships among experiments on reinforcement frequency.

If it is accepted that different classes of IRTs may be treated as different operants (see, e.g., Catania, 1966), then the present results are in accord with the explanation given by Neuringer (1967) for why some experiments on reinforcement magnitude have shown no effect and others have shown a pronounced effect. According to this explanation, since different IRTs received different reinforcement magnitudes in the present experiment, reinforcement magnitude determined their relative frequencies. That some dependence should have been observed is therefore clear, but the observed degree of dependence does not follow from this explanation. Thus, it re-




Fig. 5. The relative frequency of the short IRT (R) as a function of the relative access to food for the short IRT. Three panels show averaged data for individual birds, and the bottom right panel shows all bird-by-condition combinations. The least-squares best-fitting straight lines provide good descriptions of the data, and suggest that the way relative frequency and relative magnitude of reinforcement combine may be described by a linear function of relative access to food. The dashed line in the bottom right panel is derived partly from earlier two-key data (see Discussion) and gives an adequate first-order description of the averaged data.

mains to explain the difference between matching in the two-key experiments by Catania and Neuringer and the smaller dependence observed here.

One provisional explanation of this difference would describe both matching and the present results, for any given pair of π and F , as only two possible points on a continuum. Consider that pecks on different keys may be said to consume roughly the same lengths of time; however, different IRTs by their very

nature consume different lengths of time. The difference between the matching obtained in two-key procedures and the data obtained in the present one-key experiment may reflect these differences between temporal properties of the two different classes of concurrent operants. Specifically, the present data suggest that operants that are more time-consuming (*e.g.*, longer IRTs) require greater payoffs to be preferred as much as short ones. For example, the left panel in Fig. 3 shows that when

the relative feeder durations for short and long IRTs were equal, a relative frequency of reinforcement of only 0.2 for the short IRT was enough to determine almost equal responding across the two IRT classes. Or, conversely, the relative frequency of reinforcement for the long IRT needed to be as high as 0.8 to make it as attractive as the short IRT. Similarly, the right-hand panel of Fig. 3 shows that when the relative frequencies of reinforcement were equal, no relative feeder duration used here was small enough for the short IRT to reduce its relative frequency of occurrence to 0.5. It is reasonable to expect that if, with the present procedure, the short IRT class were made even shorter, while the long IRT class remained constant, the obtained functions corresponding to those in Fig. 3 and 4 would show an even greater preference for the short IRT. Whether such functions would differ from those obtained here in terms of slope, elevation, shape, or any other feature is to be determined only by further research. But it is likely to be found that, for a given pair of values for π and F, the relative frequency of an operant will depend in addition on its temporal length in comparison with that for the other operant(s). Other things being equal, the shorter operant will be chosen more often. In the case of equally long concurrent IRTs we might expect to find the matching obtained by Catania, who used concurrent pecks on different keys.

It is possible, although as explained below, not likely, that a procedural difference between the present one-key experiment and Catania's and Herrnstein's two-key experiments weakens the analogies drawn in this paper between these one- and two-key experiments. Customarily, a bird partly controls relative frequency of reinforcement in two-key concurrent variable-interval schedules; the separate variable-interval tapes run independently so that a bird conceivably could collect its reinforcements only from a single key. It is possible that this independence makes a difference in the resulting behavior. If so, the analogies drawn in the present paper between the present experiment and two-key experiments may be somewhat weakened. However, it is known that a two-key procedure similar to a concurrent variable-interval schedule generates matching, as does a concurrent variable-interval schedule, when the relative frequency

of reinforcement is controlled by the experimenter as in the present study (Shimp, 1966). Thus, it seems that in concurrent schedules, such as the present one, whether a bird does or does not partly control the relative frequency of reinforcement of an alternative appears not to affect the relative frequency of occurrence of that alternative. Since this datum is the only one under consideration here, the analogies given appear valid.

The controlling relations for reinforcement frequency presumably are much like those briefly outlined above for reinforcement magnitude. That is, reinforcement frequency may be expected to control behavior if different operants are reinforced at different rates. Furthermore, the extent of this control will depend on the different lengths of time required to emit the different operants. Again, there are the two traditionally different cases corresponding to operants defined as pecks on different keys or as different IRTs.

Anger's synthetic variable-interval schedule has provided most of the direct evidence concerning the effects of relative frequencies of reinforcement on relative frequencies of IRTs. As noted above, his data (Anger, 1954) revealed that the relative reinforcement frequency of an IRT could change the relative frequency of that IRT, but his data also appeared rather complicated. That is, the relative frequency of a short IRT (0 sec to 4 sec) went down when its reinforcement frequency was reduced. However, the relative frequency of a long IRT (8 sec to 12 sec, or 12 sec to 16 sec) did not always go up. Certainly its frequency never increased as much as the frequency of the short IRT decreased, when its reinforcement frequency was increased by an amount precisely equal to the reduction in the reinforcement frequency for the short IRT. The present schedule is a special case of Anger's schedule. That is, the relative frequencies of reinforcement for the different IRTs in the present variable-interval schedule were precisely controlled, independently of a bird's behavior, as they were in Anger's schedule; but here, there were just two classes of reinforced IRTs. Since these classes were non-adjacent and were associated with discriminative stimuli, the present schedule avoided complexities such as induction among different classes of reinforced IRTs. The present experiment extends Anger's findings by revealing that the

relative frequency of an IRT may be a linear function of the relative reinforcement frequency of that IRT. The complication discovered by Anger and discussed above was impossible in the present experiment because when there are only two reinforced IRT classes, a decrease in one relative frequency must equal the resulting increase in the other.

While the present schedule is a special case of a synthetic variable-interval schedule, simultaneously it is a concurrent variable-interval schedule of reinforcement of two different operants. Therefore, the present results may be compared with the two-key results obtained by Herrnstein (1961). The difference between the control by relative reinforcement frequency here and the matching in Herrnstein's experiment may be likened to the corresponding difference between the control by relative reinforcement magnitude here and the matching in Catania's experiment. Therefore the discussion above need not be repeated in detail. Suffice it to say that the dependency of the relative frequency of one among several concurrent operants on the lengths of time required to emit the various operants may be the same in the case of reinforcement frequency as described above for reinforcement magnitude.

If, as suggested, one-key and two-key concurrent variable-interval schedules produce different functions partly because of differences in lengths of operants, then perhaps the functions could be made more similar by somehow taking these lengths into account. One way to do this is suggested by an experiment performed by Chung and Herrnstein (1967) on delay of reinforcement in concurrent variable-interval schedules. They found that the relative frequency of responses on a key approximately equalled the relative immediacy of reinforcement on that key, when roughly equal amounts of food were obtained on both keys. Delay of reinforcement, like operant length, is a temporal variable, and it is conceivable that these two variables have similar effects. The simplest assumption would be that their effects are actually equal; that is, that the length of an IRT is interpretable as a delay of reinforcement of precisely that length. (For a discussion of this assumption, and of some of its implications, see Shimp, *in press*.) This assumption predicts in the present experiment, when the relative access to food is equal for

each IRT, that the relative frequency of the short IRT will approximately equal the "relative immediacy of reinforcement" for the short IRT. If the lengths of the two IRTs are described by their midpoints, this latter variable here equals $1 - 2/(2 + 4) = 0.667$ (see Chung and Herrnstein, 1967). This value is indeed approximately equalled by the obtained value of 0.699, the average of the 11 data points in the lower right-hand panel in Fig. 5, for relative access to food equal to 0.5. While a predicted value for this single point is easily obtained by analogy with the two-key experiment by Chung and Herrnstein, the predicted slope of the entire function is not so readily deducible: how delay, frequency, and magnitude of reinforcement combine in two-key concurrent schedules is a question not yet answered. One possible generalization of the matching rule to experiments where magnitude and frequency of reinforcement are simultaneously varied is that the slope of the linear function of relative access to food, such as of the one in Fig. 5, will approximately equal the relative immediacy of reinforcement. In the present case this function would then have a slope of 0.667, and would pass through the points (0.5, 0.667) and (1, 1). The resulting function here appears as the dashed line in Fig. 5 (the presumably more complex nonlinear portion near the origin is omitted). The predicted function gives a moderately good description of the data; the fit is about on a par with that of the matching function itself in some cases (see *e.g.*, Herrnstein, 1964). But several aspects of the curve are clearly in error. First, like the matching function, it neglects individual differences. Second, the predicted slope is too high and the predicted y-intercept is too low. Nevertheless, the average difference between obtained and predicted curves is only 0.043. Therefore, the predicted curve, and the combination rule from which it is derived, is a useful summary of the more prominent tendencies in the data.

The present data are inadequate to distinguish the rule suggested above from another quite different yet equally possible rule. The linear function may pass through the matching point, *i.e.*, the point (0.5, 0.67) in the example above, and through the point (1, 1). The slope would then equal the relative immediacy of the short interresponse time only if the latter equalled 0.67, as in the present ex-

periment. Clearly, more data are needed from both one-key and two-key experiments to establish a general combination rule.

REFERENCES

- Anger, D. *The effect upon simple animal behavior of different frequencies of reinforcement*. Report PLR-33, Office of the Surgeon General. Document No. 7779, ADL, Auxiliary Publications Project, Photoduplication Service, Library of Congress, Washington, D.C., 1954.
- Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 1956, **52**, 145-161.
- Blough, D. S. Interresponse time as a function of continuous variables: a new method and some data. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 237-246.
- Blough, D. S. The reinforcement of least-frequent interresponse times. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 581-591.
- Catania, A. C. Concurrent performances: a baseline for the study of reinforcement magnitude. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 299-300.
- Catania, A. C. Concurrent operants. In W. K. Honig (Ed.), *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. Pp. 213-270.
- Catania, A. C. and Reynolds, G. S. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 327-383.
- Chung, S. and Herrnstein, R. J. Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 67-74.
- Hendry, D. P. The effect of correlated amount of reward on performance on a fixed-interval schedule of reinforcement. *Journal of Comparative and Physiological Psychology*, 1962, **55**, 387-391.
- Herrnstein, R. J. Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1961, **4**, 267-272.
- Herrnstein, R. J. Secondary reinforcement and rate of primary reinforcement. *Journal of the Experimental Analysis of Behavior*, 1964, **7**, 27-36.
- Neuringer, A. J. Effects of reinforcement magnitude on choice and rate of responding. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 417-424.
- Shimp, C. P. Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 443-455.
- Shimp, C. P. The reinforcement of short interresponse times. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 425-434.
- Shimp, C. P. Optimal behavior in free-operant experiments. *Psychological Review*, in press.

Received 17 November 1967.