**CONCURRENT SCHEDULES OF INTERRESPONSE TIME
REINFORCEMENT: PROBABILITY OF REINFORCEMENT
AND THE LOWER BOUNDS OF THE REINFORCED
INTERRESPONSE TIME INTERVALS¹**

RICHARD W. MALOTT² AND WILLIAM W. CUMMING

COLUMBIA UNIVERSITY

Data were obtained with rats on the effects of interresponse time contingent reinforcement of the lever press response using schedules in which interresponse times falling within either of two temporal intervals could be reinforced. Some of the findings were (a) the mode of the interresponse time distribution generally occurred near the first lower bound when the maximum reinforcement rate for the two lower bounds was equal; this also frequently occurred even when the reinforcement rate was less for the first lower bound; (b) as is the case with schedules using a single interval of reinforced interresponse times the values of the lower bounds partially determined the location and spread of the distributions; but the particular pair of values used did not seem to influence the effects of the probabilities of reinforcement; (c) although the modal interresponse time was usually at the lower bound of one of the two intervals of reinforced interresponse times, no simple relation existed between either the probability or rate of reinforcement of interresponse times in these two intervals and the location of this mode.

This is part of a series of experiments (Malott and Cumming, 1964, 1965) aimed at analyzing the effects of schedules of interresponse time (IRT)³ reinforcement. It is hoped that eventually a general formulation will be obtained which will allow the IRT distribution to be described as a function of any arbitrary schedule of IRT reinforcement. Such a formulation should make it easier to determine the importance of the reinforcement of IRTs which occurs as a by-product of classical schedules of reinforcement, *e.g.* ratio and interval schedules. In the present study, IRTs falling within either of two temporal intervals could be reinforced. This concurrent schedule of IRT reinforcement is intermediate in complexity between the usual simple schedules using a single interval of reinforced IRTs and the arbitrary schedules referred to above. It

therefore seems to be a necessary step toward obtaining a general formulation of IRT reinforcement (See Malott and Cumming [1964] for a detailed exposition of this theoretical analysis). By manipulating the lower bound of both the first and second interval of reinforced IRTs and the probability of reinforcement associated with both the intervals, it was also hoped that some information could be obtained about the relative importance of rate and probability of reinforcement of IRTs.

METHOD

Subjects

Twelve naive male albino rats were fed at the end of each session and were maintained at approximately 80% of their extrapolated, normal weight.

Apparatus

A Grason-Stadler test chamber was modified to contain a liquid dipper, the cup of which was normally within the chamber. The dipper was used to deliver a drop of Eagle Brand sweetened condensed milk (mixed 1:1 with water) which served as the reinforcement for the lever press response. A force of about 24 g was required to depress the lever. The microswitch click provided the only auditory feed back for

¹This research was supported by Public Health Service Grant MH-06911-01.

²Reprints may be obtained from Richard W. Malott, Dept. of Psychology, Western Michigan University, Kalamazoo, Michigan.

³The IRT is the time between successive responses and is reciprocally related to the response rate. Throughout, the expression "reinforcing an IRT" refers, of course, to providing a reinforcement for a response terminating the IRT.

the lever press response. The programming and recording circuitry was housed in an adjoining room.

Probability of reinforcement was controlled by a 90 rpm cam timer. Each time an IRT of appropriate length occurred, a fixed duration pulse was fed through the microswitch on the cam; if the switch was closed, a reinforcement was delivered. The cam was adjusted to control the per cent of time the switch was closed and thereby the probability of reinforcement. It was necessary to empirically calibrate the cam. Because the time per revolution of the cam was small relative to the minimal reinforced IRT, and because distribution of IRTs was varied, this method seemed quite satisfactory for generating probabilities. There was no evidence that the animals ever "synchronized" with the cam or that the reinforcement delivery ever deviated from a random pattern.

Procedure

Sessions lasted 50 min. After the subjects had been trained to drink from the dipper as soon as it was presented and conditioned to press the lever, they were randomly divided into groups of two and assigned different pairs of intervals of reinforced IRTs. The interval of reinforced IRTs is a temporal interval specified by its lower and upper bounds; only IRTs falling within one of these intervals can be reinforced. The width of each interval was set equal to its lower bound; this procedure made the intervals behaviorally equal as shown by Malott and Cumming (1964). For each subject the probability of reinforcement

associated with each interval was manipulated. The general design of the experiment is shown in Table I. On the basis of data obtained earlier (Malott and Cumming, 1964, 1965) the duration of each phase was preselected with the hope that the last five sessions would represent asymptotic performance. For reasons presented by Cumming and Schoenfeld (1960), the present data were not analyzed in terms of a stability criterion. It should be kept in mind when interpreting the results that the subjects were exposed to increasing probabilities of reinforcement in the first interval of reinforced IRTs and decreasing probabilities in the second interval. This problem was to some extent counteracted by a procedure described below.

At the beginning of some phases, only a small number of IRTs would fall within the reinforcement interval for which the probability of reinforcement was being manipulated. In these cases it appeared unlikely that a sufficient sample of this new probability would be obtained to affect the animal's behavior. Consequently, the probability of reinforcement in the other interval was reduced to zero for a single session; this resulted in an increase in the frequency of IRTs falling within the interval of concern. This procedure involved a calculated risk; it was possible that a temporary change in the schedule of reinforcement would produce a permanent change in the IRT distribution. The earlier data on recoverability of baseline (Malott and Cumming, 1964) indicated that this would probably not be a problem. This was also indicated

Table I

Design of Experiment

(The columns correspond to groups and the rows to the different phases of the experiment to which each group was exposed. The phases were presented in descending order. The number of sessions for each phase is shown in the body of the table.)

		Lower bounds (sec)				Prob. of reinforcement	
8-16	4-16	2-16	4-8	2-8	2-4	1st interval	2nd interval
60	60	60	45	45	30	.000	1.000
15	15	15	15	15	15	.125	1.000
15	15	15	15	15	15	.250	1.000
15	15	15	15	15	15	.500	1.000
15	15	15	15	15	15	1.000	1.000
15	15	15	15	15	15	.125	.500
15	15	15	15	15	15	.250	.500
			15	15	15	.500	.500
					15	.125	.250

by the present data: frequently on the return to the probabilities used for the remainder of the phase, the distribution would return to its previous form, indicating that the new probabilities were not sufficiently different from those in the previous phase to result in a change in IRT distribution.

Data for the IRT distributions shown in Fig. 1-4 were obtained from the last five sessions of each phase; it was estimated that, by that point, the behavior of all animals had stabilized. Data were not obtained for one of the subjects in the 4- and 8-sec lower bounds group during phases 6-8 because it died from a respiratory infection. The width of the IRT class interval used in recording the IRTs was equal to $\frac{1}{16}$ of the lower bound of the second interval of reinforced IRTs for all but the last group. Due to apparatus limitations, the class width was $\frac{1}{8}$ of the second lower bound for this group. To facilitate comparisons, the data for this group were transformed to approximate the results that would have been obtained if the class widths had been $\frac{1}{16}$ of the second lower bound (*c.f.* Malott, 1963). This transformation slightly decreases the maxima and increases the minima from what would ordinarily be obtained.

RESULTS⁴

The data obtained when the probability of reinforcement was unity for the second rein-

forcement interval are shown in Fig. 1 and 2. See Table II for the total number of IRTs used in computing each distribution. The distributions for the two subjects in the same groups are shown on the same sets of coordinates. All of the data for a given column are obtained from the same two subjects and the same pair of intervals of reinforced IRTs. In a particular column, the distributions drawn with the solid lines are for one subject and those drawn with the dotted line are for the other subject. The vertical dashed lines with the numbers at the top indicate locations of the lower bounds of the reinforced IRT intervals. The numbers at the top of these lines indicate values of the lower bounds. The vertical dashed lines without the numbers at the top indicate locations of the upper bounds; in some cases the upper bound of the first interval served as the lower bound of the second. Each column presents the data for a different pair of lower bounds, and each row presents the data for a different value of the proba-

⁴A 13-page table giving the relative frequency distributions of the interresponse times has been deposited as Document number 8965 with the American Documentation Institute, Auxiliary Publication Project, Photoduplication Service, Library of Congress, Washington 25, D.C. A copy may be secured by citing the Document number and by remitting \$2.50 for photoprints or \$1.75 for 35 mm microfilm. Advance payment is required. Make checks or money orders payable to: Chief, Photoduplication Service, Library of Congress.

Table II

Total number of IRTs used in distributions shown in Fig. 1-4
(The solid curves in those figures are for the distributions of the even numbered subjects, the dotted curves for the odd numbered subjects.)

Lower bounds (sec)												Prob. of reinforcement	
8-16		4-16		2-16		4-8		2-8		2-4		1st interval	2nd interval
Rat number													
66	67	68	69	70	71	72	73	74	75	76	77		
736	901	900	695	921	229	1715	1400	1652	1482	2222	2706	0.000	1.000
766	992	933	859	1213	854	1475	1301	1703	1081	2198	3064	.125	1.000
1493	1888	2228	1808	5909	5460	2415	1469	3762	4788	2631	1503	.250	1.000
1175	2146	2299	2289	5798	6274	1294	1906	2892	5942	5695	5703	.500	1.000
1270	1792	3001	1824	3040	4312	2257	1630	3684	5649	4390	4944	1.000	1.000
541	1125	1302	1253	4768	5610		402	1663	2958	4531	1510	.125	.500
779	1474	372	1308	4987	5860		1179	3625	6059	3403	5183	.250	.500
							1501	3235	6165	6146	5024	.500	.500
										3198	3412	.125	.250

Fig. 1. Interresponse time distributions when the probability of reinforcement for the second reinforcement interval was 1.000.

Fig. 2. Interresponse time distributions when the probability of reinforcement for the second reinforcement interval was 1.000.

bility of reinforcement in the first interval of reinforced IRTs. Because of the large amount of data, it was necessary to use two figures. Figure 1 presents the data for the groups for which the lower bound of the second interval of reinforced IRTs was 16 sec. Figure 2 presents the data for the remaining groups.

The distributions remained fairly constant until the probability of reinforcement in the first interval was increased to .250; then, with three exceptions, the mode of the IRT distribution shifted from the lower bound of the second interval to that of the first.

Two of the exceptional subjects comprise the group with the lower bounds of 2 and 4 sec. It is not clear whether their deviation was due to the values of the lower bounds or simply individual subject differences. The mode shifted to the first lower bound for these two subjects during the next phase when the probability of reinforcement for the first interval was increased to .500.

The third deviant subject is in the 4- and 8-sec lower bounds group. The distribution for this subject became unusually flat during the fourth phase, and a clear mode did not develop near the first lower bound until the fifth phase where probability of reinforcement for the first interval was 1.000. The data for this subject deviated throughout most of the experiment.

The distributions for the other subjects remained constant or increased in peakedness as probability of reinforcement in the first interval increased from .250 to 1.000.

The data obtained when probability of reinforcement for the second interval was .500 is shown in Fig. 3 and 4. In all but one case, the mode remained at the first lower bound. The distributions were generally less peaked than in the previous phase where probability for both intervals was unity, and two of the distributions showed distinct modes at both lower bounds. Again as probability of rein-

Fig. 3. Interresponse time distributions when the probability of reinforcement for the second interval was .500.

forcement in the first interval increased, the peakedness of the mode associated with that interval increased; also, the mode at the second lower bound disappeared.

For the 2- and 4-sec lower bounds group, probability of reinforcement in the second in-

terval was also reduced to .250 and probability associated with the first interval was set at .125. In this case, the mode of the distribution was at the first lower bound.

The effect of probability of reinforcement in the second reinforcement interval may be

Fig. 4. Interresponse time distributions when the probability of reinforcement for the second interval was .500. Note that the probability of reinforcement in the second interval for the data in the last row of this figure was .250.

observed when probability in the first interval is .125. There were no modes at the first lower bound when probability of reinforcement in the second interval was 1.000. When probability in the second interval was reduced to .500, only three of the subjects had modes near the second lower bound. When probability in the second interval was reduced to .250, the two subjects exposed to that schedule no longer showed a mode at the second lower bound. As probability of reinforcement in the second interval decreased, probability in the first interval came to control the IRT distribution.

The values of the lower bounds of the two intervals had little, if any, effect that could not be observed in simple non-concurrent schedules. As is the case with schedules using a single interval of reinforced IRTs, the values of the lower bounds partially determined location and spread of the distributions; but the particular pair of values used did not seem to influence the effects of probabilities of reinforcement. Where a mode was obtained at only one lower bound, the distributions were essentially the same as if the interval without the mode had not been present.

DISCUSSION

In a previous experiment on concurrent schedules (Malott and Cumming, 1964), the mode of the IRT distribution gradually shifted from the second to the first lower bound as probability of reinforcement in the first interval was increased. In other words, for some probabilities of reinforcement the mode was at values intermediate between the extremes represented by the two lower bounds. In the present experiment, the shift of the mode was not gradual; the mode was either at one or the other or both lower bounds; *e.g.* there was no probability of reinforcement which resulted in a mode midway between the two lower bounds. In no case in the earlier experiment was there a mode at both lower bounds. Although the methods of the two experiments differed slightly, it is not clear why the results differ.

The results of this experiment are similar to those of a previous experiment on concurrent schedules in which the width of the first interval of reinforcement was manipulated (Malott and Cumming, 1964). In that experi-

ment, there was also no gradual shift in the mode; however the bimodal distribution was more consistently produced.

It remains to be seen whether probabilities of reinforcement can be found which will reliably generate modes at both lower bounds. In the present experiment, that probability for the first interval might have been between .125 and .250 for most of the subjects when the probability in the second interval was 1.000. The reason that the bimodal distributions have been so scarce when probability was manipulated may be that only a very narrow range of probabilities will generate such distributions. It may be that the range of interval widths that will have such an effect is not so narrow.

Attempts have been made to describe the results of IRT reinforcement schedules and schedules in general in terms of rate of reinforcement, *i.e.* number of reinforcements per unit time (*c.f.* Newman and Anger, 1954; and Millenson, 1964). The notion is that the rate of occurrence of the various IRTs or responses is a function of the reinforcement rate or relative reinforcement rate for those IRTs or responses. A straightforward application of such a notion would be to determine for each set of lower bounds and reinforcement probabilities the reinforcement rate associated with each lower bound. If the reinforcement rate is important, then the mode should occur at the lower bound with the highest rate of reinforcement. Consider for example, the data in the first column and third row of Fig. 1: the maximum reinforcement rate associated with the first lower bound was

$$.250 \text{ reinf} / 8 \text{ sec} = 1/32 \text{ reinf/sec}$$

and for the second lower bound was

$$1.000 \text{ reinf} / 16 \text{ sec} = 1/16 \text{ reinf/sec.}$$

The maximum reinforcement rate was twice as great for the second lower bound as for the first, yet the mode occurred at the first lower bound. In this experiment, there were eight combinations of the parameters which produced a maximum reinforcement rate for the first lower bound which was equal to half the rate for the second lower bound. Out of the resulting 15 distributions, two were bimodal, seven had a mode only at the second lower bound, and six had a mode only at the first lower bound. Incidentally, these data bring

into question the utility of a "least effort" principle. This result is opposite to that predicted by the notion of maximum reinforcement rate. It might be argued, to account for these results, that although reinforcement rate controls the IRT distribution, there is a bias in favor of the short IRTs; *i.e.*, the animal will "prefer" to emit short IRTs. If that is the case, it would be predicted that where the maximum reinforcement rate is the same for both lower bounds, the mode should occur at the first. There were 12 combinations of parameters which produced equal maximum reinforcement rates for the two lower bounds. Out of the resulting 23 distributions, two were flat, two were bimodal, one had a mode only at the second lower bound, and 18 had a mode only at the first lower bound. The data clearly support the notion that there was a response bias in favor of the short IRTs.

Instead of the *maximum* reinforcement rate, the *obtained* reinforcement rate might be the important variable. Consider, again, the data in the first column and third row of Fig. 1. The technique of manipulating the probabilities during the first session of the phase so the behavior might be affected by the changed contingency was not used. Nonetheless the mode of the distributions shifted to the lower bound of the first interval. Such a change would not be predicted on the basis of the obtained rate of reinforcement since the obtained rate was initially much greater in the vicinity of the second interval.

A similar argument might be made for probability of reinforcement (reinforcements per response) as the controlling variable. But again the data are contradictory, for it would be predicted that the mode should be near the second lower bound for all of the groups in the second and third rows of Fig. 1. This is clearly not the case. It should be noted that unlike obtained reinforcement rate, probability of reinforcement may be specified independently of behavior.

Either probability or rate of reinforcement can be manipulated and a change in behavior may be observed. It is possible to find two situations such that, in one, probability of reinforcement is the higher and, in the other, rate of reinforcement is the higher. However, for any single situation, a change in probability

of reinforcement always results in a change in rate of reinforcement and *vice versa*. Since the two variables cannot be manipulated independently, the argument has arisen as to which variable "really" controls the behavior. This seems to be a search for the independent variable or way of expressing the independent variable which yields the simplest functional relation with the behavior. Although reinforcement rate and probability of reinforcement may both be useful in describing data of the sort presented in this paper, the description will not be a simple one.

In conclusion, it was found that (a) as probability of reinforcement in the first reinforcement interval increased, the mode of the IRT distribution shifted from the lower bound of the second to that of the first interval; (b) contrary to previous findings, this shift was discontinuous; (c) there was a positive correlation between probability of reinforcement in the second interval, and the value of probability of reinforcement in the first interval necessary to produce a shift in the modal IRT; (d) there was a response bias in favor of short IRTs; (e) the value of the lower bounds of the two intervals had little if any effect that could not be observed in simple non-concurrent schedules; (f) no simple relation existed between either rate or probability of reinforcement and location of the modal IRT.

REFERENCES

- Cumming, W. W. and Schoenfeld, W. N. Behavior stability under extended exposure to a time-correlated reinforcement contingency. *J. exp. Anal. Behav.*, 1960, 3, 71-82.
- Malott, R. W. The differential reinforcement of temporally spaced responses. Doctoral dissertation, Columbia University, 1963.
- Malott, R. W. and Cumming, W. W. Schedules of interresponse time reinforcement. *Psych. Rec.*, 1964, 14, 211-252.
- Malott, R. W. and Cumming, W. W. Interresponse time reinforcement: Transition and steady states. *Psychonom. Sci.*, 1965, 2, 257-258.
- Millenson, J. R. A further investigation of the correlation between IRT frequency and IRT reinforcement distributions. *Psychonom. Sci.*, 1964, 1, 373-374.
- Newman, E. B. and Anger, D. The effect upon simple behavior of different frequencies of reinforcement. Technical Report. Contract DA-49-007-MD-408. Psychological Laboratories, Harvard University, 1954.

Received January 4, 1965