TIME ALLOCATION AND RESPONSE RATE¹

CHARLES P. SHIMP

UNIVERSITY OF UTAH

Food reinforcement for key pecking by three pigeons was arranged by a variable-interval schedule and a device that assigned each reinforcement to one of 10 component response rates corresponding to 10 classes of equally reinforced interresponse times ranging from 1.0 to 6.0 sec in 0.5-sec classes. The overall number of reinforcements per hour was varied from one to more than 60. Overall response rate was a monotonically increasing, negatively accelerated function of the overall number of reinforcements per hour. This function was decomposed into two time-allocation functions: (1) the time allocated to all of the reinforced component response rates as a function of the total reinforcement rate, and, (2) the time allocated to a particular reinforced component response rate as a function of the reinforcement rate for that component. Asymptotic response rate was predicted by combining the asymptotes of the two separate time-allocation functions: virtually all of the time was spent responding, and the percentage of the time spent responding that was allocated to a particular reinforced component response rate roughly equalled the relative reinforcements per hour for that component.

Current analyses of free-operant schedules of reinforcement frequently concentrate on the way in which various schedule parameters determine mean rate of responding maintained by variable-interval (VI) schedules and by compounds of VI schedules (Baum, 1973; Catania, 1973; Herrnstein, 1970; Rachlin, 1973). Mean rate of responding is the reciprocal of the mean of a distribution of interresponse times (IRTs). Therefore, an analysis of mean response rate must remain incomplete, and risk the possibility of being seriously in error, if it disregards variables on which the distribution of IRTs importantly depends. Current analyses of schedules of reinforcement assume, usually implicitly, that the distribution of reinforced IRTs does not importantly determine behavior maintained by VI schedules (Baum, 1973; Catania, 1973; Herrnstein, 1970; Rachlin, 1973).

An alternative viewpoint maintains that, when a reinforcement parameter depends on

the time since the last response, *i.e.*, on the IRT, then that parameter is likely to affect the distribution of IRTs. In particular, this viewpoint maintains that the distribution of reinforced IRTs importantly determines the distribution of IRTs in a VI schedule (Anger, 1956, 1973; Shimp, 1969*a*, 1973*a*). The simplest way to test this hypothesis would seem at first glance only to require one to see if orderly relationships exist between IRT and reinforced-IRT distributions in VI schedules. However, this route is blocked by the interaction between IRT and reinforced-IRT distributions in VI schedules: the subject partly determines which IRTs are reinforced (Anger, 1956). Thus, the hypothesis must be tested indirectly. Evidence bearing on this hypothesis comes from experiments in which reinforcements are arranged by a VI schedule so modified that the experimenter retains control of the distribution of reinforced IRTs. A VI schedule modified in this way is called a synthetic VI (Anger, 1973; Shimp, 1973*a*). To date, it has been demonstrated that in synthetic VI schedules, the distribution of IRTs depends on: the absolute and relative durations of reinforced IRTs; the total density of reinforcement; the relative reinforcements per hour for the various reinforced IRTs, and; the duration of the IRT preced-

¹This work was supported in part by NIMH grant 16928 and by a fellowship from the Center for Advanced Study in the Behavioral Sciences. The author would like to thank Mr. Steven Smith and Mrs. Candace Lindsey for their assistance in conducting the experiment. Reprints may be obtained from the author, Dept. of Psychology, The University of Utah, Salt Lake City, Utah 84112.

ing reinforced IRTs (Hawkes and Shimp, 1974; Shimp, 1967, 1968, 1969*b*, 1970, 1973*a*, *b*; Staddon, 1968). These variables have been investigated with three different kinds of modified VI schedules. First, paced VI schedules have been employed in which reinforced IRTs were restricted to a single narrow class. In such a schedule, there is but one reinforced rate of responding. Second, concurrent variable-interval (*conc* VI VI) schedules for two IRTs have been employed in which there are two narrow, non-contiguous classes of reinforced IRTs. In such a schedule, there are two reinforced component response rates. Third, synthetic VI schedules have been employed in which an interval of IRTs consisted of 10 reinforced component response rates. With these schedules, one finds that mean response rate is determined in large measure by the properties listed above of the distribution of reinforced IRTs.

One may argue for the assumption that these variables are irrelevant to an understanding of behavior maintained by VI schedules, because the reinforcement contingencies of these modified VI schedules and of VI schedules are not precisely identical. However, this assumption of irrelevance would seem to become progressively less tenable as evidence of controlling relations continues to accumulate from a wide variety of modified VI schedules. Also, there is no empirical evidence to support the assumption of irrelevance; *i.e.*, there are no data directly suggesting that controlling relations in these modified VI are fundamentally different from those in VI schedules. Thus, it may reasonably be tentatively assumed that mean response rate in a VI schedule depends on several properties of the distribution of reinforced IRTs.

The present experiment was designed to investigate the manner in which the distribution of reinforced IRTs and the total density of reinforcement combine to determine mean response rate in a synthetic VI rather closely resembling a VI. In VI schedules, mean response rate is a negatively accelerated, monotonically increasing function of the total number of reinforcements per hour (Catania and Reynolds, 1968). This function confounds the effects of total reinforcement density and of the distribution of reinforced IRTs, if indeed one is correct in assuming that the latter distribution affects the distri-

bution of IRTs. In an attempt to estimate the effects of total reinforcement density unconfounded with effects of a changing and uncontrolled distribution of reinforced IRTs, Shimp (1970) employed a *conc* VI VI schedule for two IRTs with which a pigeon was trained to peck a single key at two different reinforced component response rates corresponding to the two classes of reinforced IRTs. The relative reinforcements per hour for the two reinforced component response rates were equal throughout the experiment. It was found that mean response rate was a negatively accelerated, monotonically increasing function of the total number of reinforcements per hour, as in the experiment by Catania and Reynolds (1968). But at the same time, it was found that preference between the reinforced component response rates, as measured by the relative frequency of IRTs conforming to a particular component response rate, was not constant: preference for a component increased from a level of approximate indifference (corresponding to equal frequencies of IRTs conforming to either reinforced rate) at low total reinforcements per hour, to an asymptote at high total reinforcements per hour. The asymptotic preference for the higher component reinforced response rate approximately equalled the relative reciprocal (see below) of the duration of the reinforced class of IRTs defining that component. Thus, the temporal patterning of behavior, as indexed by the IRT distribution, depended on the total reinforcements delivered per hour even when the reinforced-IRT distribution was held constant.

The present experiment resembled that described by Shimp (1970) except that here, the schedule was a synthetic-VI schedule with 10 component reinforced response rates corresponding to 10 classes of reinforced IRTs ranging in 0.5-sec classes from 1.0 to 6.0 sec.

METHOD

Subjects

Three White Carneaux pigeons were maintained at approximately 80% of their free-feeding weights. All three birds previously had served in various two- or three-key experiments.

Apparatus

Three, three-key Lehigh Valley Electronics pigeon chambers were interfaced to a Digital Equipment Corporation PDP-12 laboratory computer that arranged all stimuli and reinforcements, and recorded the data. Only the center keys were used in the experimental chambers.

Procedure

The computer program employed here was the same as that used by Shimp (1973a). In the previous experiment, the overall reinforcements per hour were held constant while the relative reinforcements per hour for the reinforced component response rates were varied. Here, the overall reinforcements per hour were varied while the relative reinforcements per hour were held constant.

Variable-interval schedule. A distribution of interreinforcement intervals was arranged by a single VI schedule in which there was a given probability, p , of assigning a reinforcement every t sec. Table 1 shows how p and t were varied over experimental conditions, and also shows the equivalent mean interreinforcement interval.

When a reinforcement was arranged by the VI schedule, it remained available until a response terminated the appropriate IRT (see below) and produced a reinforcer (1.6-sec access to mixed grain). The VI programmer stopped when a reinforcement was arranged. No additional reinforcement could be ar-

ranged until a previously arranged one was delivered.

The distribution of reinforced IRTs. Throughout the experiment, the reinforced IRTs ranged from 1.0 to 6.0 sec in 10, 0.5-sec classes. The interval from 1.0 to 6.0 sec was selected because it includes nearly all IRTs that typically occur in most experiments with interval schedules, except that it excludes those short IRTs that appear to depend to a unique extent on the natural topography of a pigeon's pecking for food (Blough, 1963; Shimp, 1973a). Each reinforcement arranged by the VI schedule was assigned to one of these 10 classes of reinforced IRTs. These assignments were random in that each class was equally likely to be reinforced and one assignment did not depend on the previous assignment.

Other procedural details. Visual response feedback was provided by darkening the white keylight and the houselight for 0.5 sec after each key peck. Each session lasted 1 hr and sessions were conducted six days a week. Experimental conditions were terminated when the number of responses per session appeared stable for several days for each of the three subjects.

Pretraining. Subjects' initial exposure to the schedule was with five, 1.0-sec classes of reinforced IRTs with parameters in the VI schedule that reduced it to a continuous reinforcement schedule. The classes of reinforced IRTs always extended from 1.0 to 6.0 sec. Immediately preceding the conditions for

Table 1
Experimental Conditions

Condition Number	Parameters in the Random-Interval Schedule		Equivalent Mean Interreinforcement Interval (sec)	Number of Days of Training
	p	t (sec)		
1	0.002	1	500	40
2	0.001	1	1000	44
3	0.001	2	2000	21
4	0.001	3	3000	20
5	0.100	1	10	41
6 ^a	0.017	1	60	33
7	0.008	1	120	30
8	0.100	1	10	37

^aBetween Conditions 6 and 7 was a condition intended to replicate Condition 5 but with uninterpretable results owing to a mistake in the computer program.

Table 2
Response Rates and Reinforcement Rates

Condition Number	Mean Number of Reinforcements Per Hour ¹			Mean Number of Responses Per Minute ²			Percentage of IRTs Within the Reinforced Interval (1.0 to 6.0 sec)		
	Bird 1	Bird 2	Bird 3	Bird 1	Bird 2	Bird 3	Bird 1	Bird 2	Bird 3
1	5	7	6	20.7	23.2	19.4	96	88	89
2	3	3	3	19.0	21.6	15.6	80	87	86
3	1	1	1	16.7	20.9	11.4	84	85	84
4	1	3	1	11.6	19.4	7.5	81	83	75
5	62	54	60	23.2	26.7	23.0	95	81	99
6	30	32	26	21.0	23.8	21.4	93	93	98
7	18	14	20	20.0	22.3	19.6	91	90	94
8	69	58	64	20.5	25.2	23.2	96	86	98

¹Each entry is a mean over the last six sessions.

²A table entry is the number of responses in a session divided by a time equal to the session duration minus time during which the food hopper was presented. Each entry is a mean over the last six sessions.

which data are presented here, the subjects received 36 days of training with 10, 0.5-sec reinforced classes of IRTs and with a 10-sec average programmed interreinforcement interval.

RESULTS

The results are presented here in a sequence beginning at the molecular level and ending with the conventional response-rate function. A preliminary to the molecular analysis of IRTs in the reinforced classes is a check to ensure that a respectable percentage of all responses in fact terminated IRTs in those classes. The last three columns of Table 2 show that this was the case.

The essential points regarding behavioral events at the molecular level are summarized in Figures 1 and 2. The left panel of Figure 1 shows the relative-frequency distributions of IRTs from the two conditions with the most extreme reinforcement densities; Condition 3 (the lowest reinforcements per hour) and Condition 5 (the highest reinforcements per hour). With regard to the results described below, the intermediate conditions produced intermediate results. The relative frequency of a class of IRTs was computed by dividing the total frequency of IRTs in that class by the total frequency of IRTs in all 10 classes. These computations included only those IRTs in the reinforced IRT classes. The mean relative frequency of IRTs in each class over the

last six days of each condition was computed for each subject. Figure 1 shows that the relative frequency of the shortest reinforced IRTs (those from 1.0 to 1.5 sec in duration) tended to be greater for the highest than for the lowest reinforcements per hour. This ranking was reversed for intermediate IRTs a little longer than the shortest IRTs. The ranking for most of the longest IRTs was the same as for the shortest IRTs. The left panels of Figure 1 show that the relative frequency of a class of IRTs approximately equalled the relative reciprocal of the duration of that class, and that the approximation was somewhat better for the highest reinforcements per hour. The relative reciprocal of a class simply equals the reciprocal of the lower bound of that class divided by the sum of all 10 reciprocals. Stated symbolically, the relative reciprocal of the i^{th} reinforced IRT class is

$$\frac{1}{T_i} \bigg/ \sum_{j=1}^{10} \frac{1}{T_j}$$

where T_j is the lower bound of the j^{th} class. The lower bounds, rather than the midpoints, were used in the computation because IRT distributions tend to be skewed, with a greater frequency of IRTs nearer to the lower bound than to the midpoint of a reinforced class. The matching-to-the-relative-reciprocal values

Fig. 1. The relative-frequency distributions of IRTs from conditions with the highest and lowest reinforcement rates (left panels). The per cent of the time spent responding allocated to each reinforced component response rate (right panels). Dashed lines correspond to matching-to-relative-reciprocal values. Each point for an individual subject represents a mean over the last six days of a condition.

are shown in the left panels of Figure 1 as dashed lines. The data shown in Figure 1 for Condition 5 also closely represent the data for Condition 8, the replication of Condition 5.

If the relative frequency of an IRT approximately equals the relative reciprocal of its duration, then the percentage of the time consumed by all 10 classes should be partitioned into 10 approximately equal groups. The right panels of Figure 1 show the extent to which this was so. The vertical axis is simply the time allocated to a class divided by the time allocated to all 10 classes. That is, it is the per-cent-of-the-time-spent-responding that was allocated to a particular class. Specifically, the vertical axis is the relative frequency of a class times the lower bound of that class, divided by the sum of all 10 similar products.

For the condition with the highest reinforcements per hour, the average percentage of the time spent responding allocated to a particular component response rate, or class of reinforced IRTs, was within 0.03 of the relative reinforcements per hour for that component, *i.e.*, 0.10, except for the two long-

est reinforced classes. It has been noted before that the frequencies of these longest classes are lower than predicted from matching-to-the-relative-reciprocal (Shimp, 1973a), and this same discrepancy arises here. For the condition with the lowest reinforcements per hour, the percentages of the time spent responding allocated to the various components did not so closely approximate the value of 0.10, and the discrepancy just noted was obtained over a wider range of reinforced classes of IRTs.

Figure 2 shows another way in which the temporal patterning of behavior depended on the total reinforcements per hour while the reinforced-IRT distribution was held constant. Figure 2 shows the IRTs/Op from the same two conditions as in Figure 1. The IRTs/Op for a class of IRTs was computed here as it was in Shimp (1973a). The number of IRTs in a given 0.5-sec class longer than 6.0 sec was too small to be reliable, but the total number of IRTs longer than 6.0 sec was sufficiently large so that it could serve as the denominator in the computation of IRTs/Op for the longest class of reinforced IRTs. In

Fig. 2. The interresponse times per opportunity from conditions with the highest and lowest reinforcement rates. The dashed line corresponds to matching-to-relative-reciprocal values. Each point for an individual subject represents a mean over the last six days of a condition.

short, the IRTs/Op for a class of IRTs was computed by dividing the number of IRTs in that class by the total number of IRTs in that class or any longer class, including those too long to be reinforced. The mean IRTs/Op was calculated for each class of reinforced IRTs for each subject over the last six days of a condition. As in Figure 1, the group mean is representative of each subject with regard to each result described here. Figure 2 shows that the IRTs/Op curve predicted by the matching-to-the-relative-reciprocal phenomenon is an asymmetric U-shaped function.² When the number of reinforcements per hour was large, the obtained IRTs/Op curve had such a shape, but when the number of reinforcements per hour was small, it did not. Instead, it was an inverted U with its mode near the middle of the interval of reinforced IRTs. Perhaps more im-

portantly, the range of variation in the IRTs/Op was reduced to less than 0.15. The IRTs/Op curve for the lowest reinforcements per hour simply may be a noisy flat line, and in any case is flatter for low than for high reinforcements per hour. Conditions intermediate between Conditions 3 and 5 produced results intermediate between the two curves shown in Figure 2. Recall that such was the case also with Figure 1.

To summarize the results at the molecular level, one may say that the temporal patterning of behavior, as indexed by the IRT distribution, depended on the overall reinforcements per hour even though the relative reinforcements per hour were constant. Roughly speaking, the temporal patterning of behavior at the two extremes on the input-output function may be described as follows. For the lowest reinforcements per hour, the IRTs/Op curve varied within a range of only 0.15, and in that sense was essentially flat. A flat IRTs/Op curve in turn may define a situation in which response probability is independent of time since the last response. Thus, a subject was more nearly indifferent to time since the last response when the total number of reinforcements per hour was small. For the highest reinforcements per hour, the relative frequency of an IRT, for most of the reinforced classes of IRTs, approximately equalled the relative reciprocal of the duration of that IRT. That is, the percentage of the time spent responding that was allocated to one class approximately equalled the relative reinforcements per hour for that class, *i.e.*, 0.10.

Now let us examine the data from a molar viewpoint. Figure 3 shows the percentage of the time allocated to IRTs in any of the 10 classes of reinforced IRTs plotted against the total reinforcements per hour. The dependent variable in Figure 3 was computed by first summing 10 products, where each product equalled the lower bound of a reinforced class of IRTs times its frequency. This sum estimates the time spent responding, *i.e.*, the time allocated to the various reinforced response rates. Next, the sum was divided by the entire session duration minus the time during which the food hopper was up. The resulting ratio estimates per cent-time-spent-responding. Figure 3 shows the mean of this ratio over the last six days of

²The predicted IRTs/Op for the longest class is unity and is omitted from the figure.

Fig. 3. Per cent time spent responding, *i.e.*, the per cent of the session duration allocated to any of the 10 reinforced components, as a function of the obtained reinforcements per hour.

each condition plotted against obtained total reinforcements per hour. This function is an input-output function with per-cent-time-spent-responding, rather than number of responses, as the output. Figure 3 shows that the per-cent-time-spent-responding tended to increase sharply over the interval from one to 10 reinforcements per hour, and reached approximately asymptotic values roughly between 20 and 30 reinforcements per hour. Greater reinforcement densities had no appreciable effect on the percentage of time a subject allocated to the reinforced behaviors. The asymptotic percentages were approximately 80% (Bird 2), 86% (Bird 1), and 95% (Bird 3).

Finally, let us examine the conventional input-output function relating responses per minute to reinforcements per hour. The number of responses per minute was computed by dividing the total number of responses in a session by the entire session duration (minus the time the food hopper was up). Figure 4 shows the mean of this ratio over the last six days of each condition. As in Figure 3, the curves in Figure 4 rise sharply as the reinforcements per hour increase from one to about 10, but then rise very slowly. The response rates for the highest reinforcements per hour ranged from 21 to 26 responses per minute.

DISCUSSION

The present experiment examined behavioral changes as a function of reinforcement

Fig. 4. The number of responses per minute as a function of the obtained reinforcements per hour.

density. Behavioral changes were attributable to two time-allocation functions. One time-allocation function shows how a subject partitions time between reinforced and unreinforced behaviors, as a function of the total density of reinforcement. Thus, the per cent of the time that a subject allocated to all of the reinforced component response rates was a monotonically increasing, negatively accelerated function of the total reinforcements per hour. The shape of this function qualitatively resembles those reported by Catania and Reynolds (1968) for VI schedules. A second time-allocation function shows how a subject partitions time among the various reinforced classes of behavior. Here, the reinforced classes consisted of different reinforced component response rates corresponding to 10 classes of reinforced IRTs. This second function was shown to depend on the value of the first function. For high reinforcement densities, the time allocated to all the reinforced component response rates was partitioned among them in such a way that the time allocated to a particular component approximately equalled the relative reinforcements per hour for that component.

These two time-allocation functions combine to determine mean response rate. The special case of asymptotic response rate has both theoretical and practical significance, and therefore merits detailed attention. This case is of practical importance because it corresponds to the common use in experiments on operant behavior of total reinforcement

densities of 20 or more reinforcements per hour. Its theoretical importance will be described below. Let us now derive asymptotic mean response rate from the two time-allocation functions. The session duration in the present experiment was 60 min, and by assumption, a subject allocated approximately 10% of the 60 min to each of the 10 reinforced component response rates, since each received 10% of the reinforcements. Thus, a subject allocated about 6 min to each of the 10 components. Each 6-min allotment generated a certain number of key pecks, depending on the duration of the corresponding class of IRTs. Summed over the 10 components, these numbers of key pecks give the predicted total number of key pecks. This discussion can be stated symbolically as follows. Let L_i represent the lower bound (in seconds) of component i , let π_i represent the relative reinforcements per hour for component i , and let D represent the session duration (in seconds). Then, the time allocated to the i^{th} component is $\pi_i D$ during which there would be approximately $\pi_i \cdot D/L_i$ key pecks. In the entire session, there would be

$$D \sum_{i=1}^{10} \frac{\pi_i}{L_i}$$

key pecks. This number, divided by the session duration (in minutes) is the mean number of key pecks per minute. This rate equals 24.2 responses per minute in the present experiment, and is represented in Figure 4 by a dashed horizontal line. As Figure 4 shows, the derived response rate rather accurately describes the steady state performance at the molar level for all three organisms. This description is noteworthy for two reasons: (1) it may be computed in advance of the experiment, because it requires the estimation of no free parameter, and, (2) it emphasizes the derivative rather than fundamental nature of the traditional, molar, input-output function reported by Catania and Reynolds (1968).

The preceding derivation of mean response rate in a modified VI schedule suggests that recent analyses of mean response rate in VI schedules may be in error in neglecting the distribution of reinforced IRTs. Let us examine in detail some tentative implications of the present results for theories proposed by Herrnstein (1970) and Rachlin (1973).

Herrnstein hypothesized a free parameter, k , which is to represent asymptotic behavioral output in terms of a particular response, such as a key peck. Herrnstein estimated k from the data of Catania and Reynolds (1968) to range in value from 68 to 115 for different subjects. That is, asymptotic pecks per minute ranged from 68 to 115. According to Herrnstein's theory, this asymptote is a constant, once estimated for a particular subject, response topography, and type of reinforcement. But we have seen above that this asymptote depends on the distribution of reinforced IRTs. Thus, there may be no such thing as an asymptotic mean rate of key pecking that is, as assumed in Herrnstein's theory, independent of reinforced IRTs. The implications of the present results for Rachlin's theory are related to those for Herrnstein's. Rachlin, confronted with apparently contradictory data, stated that "to solve the problem we have to consider, on a more molecular level, how changes in rate of response come about. There is some evidence that even with simple reinforcement schedules, rats and pigeons may spend more or less time responding, but, while responding, the rate remains fairly constant" (Rachlin, 1973). Rachlin presented a figure from which it is clear that by a constant response rate, he referred to a constant IRT. However, the references that he cited to support the assumption of an underlying constant response rate employed the cruder criterion of a straight-line cumulative record, and in none of the cited experiments did any reinforcement parameter depend on IRT, unlike the case in VI schedules. The absence of supporting evidence from contexts with reinforcement contingencies resembling those in a VI schedule, together with results such as the present ones, suggest that there may be no such thing as a response "tempo" that is independent of the distribution of reinforced IRTs.

Behavioral control exerted by a distribution of reinforced IRTs is an instance of a more general effect, the control of sequences, *i.e.*, of temporal patterns, of different response topographies: when a reinforcer is delivered for a free-operant, a preceding temporal pattern of response topographies may be strengthened. In the case of IRTs, the delivery of food after a sequence of a key peck, an x-sec non-key peck, and a key peck, tends

to strengthen that particular sequence of response topographies. The identification of the mechanism or mechanisms by means of which a pattern is strengthened, *i.e.*, whether by delay-of-reinforcement gradients, chaining, or direct reinforcement of the pattern itself, remains an important issue for further research.

REFERENCES

- Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 1956, 52, 145-161.
- Anger, D. The effect upon simple animal behavior of different frequencies of reinforcement. Part II: Separate control of the reinforcement of different IRTs. *Journal of the Experimental Analysis of Behavior*, 1973, 20, 301-312.
- Baum, W. M. The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, 1973, 20, 137-153.
- Blough, D. S. Interresponse times as a function of continuous variables: a new method and some data. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 237-246.
- Catania, A. C. Self-inhibiting effects of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 517-526.
- 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.
- Hawkes, L. and Shimp, C. P. Choice between response rates. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 109-115.
- Herrnstein, R. J. On the law of effect. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 243-266.
- Rachlin, H. C. Contrast and matching. *Psychological Review*, 1973, 80, 217-234.
- Shimp, C. P. The reinforcement of short interresponse times. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 425-434.
- Shimp, C. P. Magnitude and frequency of reinforcement and frequency of interresponse times. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 525-535.
- Shimp, C. P. Optimal behavior in free-operant experiments. *Psychological Review*, 1969, 76, 97-112. (a)
- Shimp, C. P. The concurrent reinforcement of two interresponse times: the relative frequency of an interresponse time equals its relative harmonic length. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 403-411. (b)
- Shimp, C. P. The concurrent reinforcement of two interresponse times: absolute rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 1-8.
- Shimp, C. P. Synthetic variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 311-330. (a)
- Shimp, C. P. Sequential dependencies in free-responding. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 491-497. (b)
- Staddon, J. E. R. Spaced responding and choice: a preliminary analysis. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 669-682.

Received 1 June 1973.

(Final Acceptance 7 December 1973.)