

SECONDARY REINFORCEMENT AND RATE OF PRIMARY REINFORCEMENT¹

R. J. HERRNSTEIN

HARVARD UNIVERSITY

Four pigeons were trained to peck at either of two response-keys. Pecking at either key occasionally produced a secondary reinforcer. Then, in the presence of the secondary reinforcer, further pecking occasionally produced the primary reinforcer, food. The relative rate at which each pigeon pecked to obtain a secondary reinforcer equalled the relative rate of primary reinforcement in its presence.

The learned behavior of animals differs from the learned behavior of human beings not only in degree of complexity, but also in its heavy dependence on biologically important events. While the learning of animals seems to depend on food, water, sex, and the other basic needs, the learning of human beings seems to range, in addition, over a host of biologically arbitrary goals. Some psychologists have attempted to bridge this gap with the concept of secondary, or conditioned, reinforcement. All learning, they say, is based on reinforcement, but reinforcers are either primary—arising in the intrinsic needs or drives of the organism—or secondary—acquiring their potency in the individual experience of a particular organism. It is said that any neutral stimulus may become a secondary reinforcer if it is present at the same time as a primary reinforcer. No one knows the exact specifications for this transfer of effectiveness from a primary reinforcer to an originally neutral stimulus, but the phenomenon itself has been demonstrated frequently. With this sanction to infer reinforcers where needed, the psychologist may hope to account for the diversity of human behavior in terms of a process that is common to animals and human beings.

Considering its key position in an explanation of human behavior, we might expect to

find secondary reinforcement the subject of many experiments. The nature of the phenomenon, however, makes it difficult to study quantitatively. A secondary reinforcer is created by the repeated pairings of a stimulus with a primary reinforcer, but, when the pairings are discontinued, the stimulus gradually loses its power to reinforce. One experimental approach is to establish a secondary reinforcer and then to discontinue the pairings. Because the effectiveness of the stimulus as a reinforcer gradually disappears as it is repeatedly presented, any extended study will be confronted with the knotty problem of evaluating a continuously changing parameter. The alternative approach—that of not discontinuing the pairings—usually permits no ready way of separating the effects of the primary from those of the secondary reinforcer.

Autor (1960) used a procedure that seemed to bypass the difficulties just mentioned. Pigeons, pecking at a pair of keys, gained access to either of two secondary reinforcers, in the presence of which the pairings with the primary reinforcer were sustained throughout the experiment. In a sense, the pigeon chose one or the other of the two secondary reinforcers by responding to one or the other of the two keys. Internal controls in the experiment made it possible to separate the effects of primary and secondary reinforcement. Autor found that the pigeon's preference for a secondary reinforcer was governed by the frequency of primary reinforcement in its presence. The present experiment is a further examination of this relationship.

¹This work was supported by grants from the National Science Foundation to Harvard University. The assistance of Mrs. Antoinette Papp and Mr. Wallace Brown is greatly appreciated. Reprints may be obtained from the author, Psychological Labs, Memorial Hall, Harvard University, Cambridge 38, Mass.

METHOD

Subjects

Four adult, male, White Carneaux pigeons were maintained at about 85% of normal weight. Two (Nos. 50 and 67) had no prior experimental training, and two (Nos. 055 and 231) had been in a lengthy series of experiments.

Apparatus

An experimental chamber contained two response-keys—translucent plastic discs—which required forces of at least 20 g to be operated, a solenoid-operated hopper that occasionally presented grain to the pigeon, and a speaker that sounded a continuous white noise. The chamber was illuminated by two 7-w, white bulbs during experimental sessions. In addition, the response-keys were transilluminated at different times with lights of various colors, correlated with various phases of the procedure. All subjects (*Ss*) were experimented upon daily over a period of 19 months. A switching circuit was used for programming and recording.

Procedure

The procedure is represented schematically in Fig. 1. Each box contains one of the three possible states of the two keys (L for left key




Fig. 1. Diagram of one cycle of the experimental procedure. Each box represents one of the three possible states of the two response-keys. At the start of a cycle (left box), both keys are illuminated and operative. Pecking at either key leads, on a VI 1', to the state represented by the box either at the upper right or lower right, during which pecking at one key leads intermittently to food (S^R), while pecking at the other, which is now darkened, has no consequences.

and R for right key). At the start of a session (left box) both keys were transilluminated with white light. Correlated with each key was a variable-interval schedule with an average interval of 1 min (VI 1'). The schedules for the two keys were equal, but mutually out of phase and independent. These schedules did not regulate the delivery of food, but, rather, a change in the stimuli. When, for example, a peck at the left key occurred after one of the scheduled intervals, the stimuli were altered as shown in the upper box at the right. The left key was then transilluminated with a different color—namely, red—and the right key completely darkened and made inoperative. For pecking at the left key, *S* was given food (4-sec presentation of the hopper and labelled S^R in Fig. 1) on a variable-ratio schedule (VR). The value of the variable ratio was systematically changed during the course of the experiment. Pecks at the right key during this phase of the procedure were totally ineffective and were virtually non-existent. After *S* twice earned access to food, the procedure reverted to that shown in the left box. A peck at the right key after one of the scheduled intervals caused the change of stimuli shown in the lower box at right. The left key was darkened and made inoperative, while the right key was transilluminated with yellow light. The schedule of food reinforcement was at times a variable ratio and at others a variable interval, and the values of these schedules were systematically changed. After two presentations of food, the original procedure was restored.

Experimental sessions ended when *S* obtained 60 presentations of food, so that the durations of the sessions depended upon the over-all rate of primary reinforcement. Since food was presented twice in each cycle of the procedure, there were 30 cycles per session. The procedure did not regulate directly the number of times *S* received food for responses to the one key as opposed to those received for responses to the other. Nevertheless, all four *Ss* produced virtually equal numbers of primary reinforcements by responses to the two keys, so that the factor of *number*—as distinguished from *rate* or *probability*—of food presentations can be, and is, disregarded in the present experiment.

This procedure comprises two chained schedules, one for each of two keys. For each

key, pecking in the presence of one stimulus intermittently produces a second stimulus, in the presence of which pecking intermittently produces food. Given two such chains, the experimenter must then select one of the several possible time relations between them. In the present experiment, the two first links are always concurrent, but the two second links are never. This procedure was first described by Autor (1960), who, in one experiment, used variable-interval schedules in all links and, in another, used variable-interval schedules in the first links and variable-ratio schedules in the second. The present experiment adds the case in which the two second links are dissimilar schedules. The value of the schedule in the first link is always VI 1'; the schedule values for the second links, the

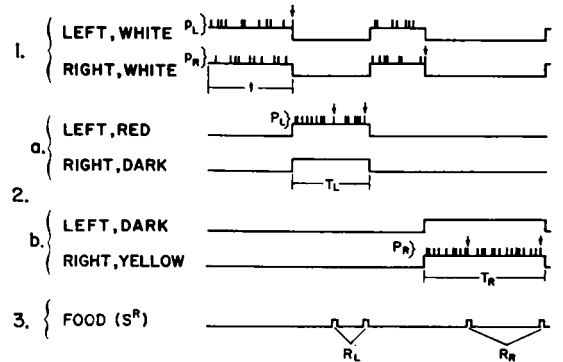


Fig. 2. Two cycles of the experimental procedure, showing schematically the relations among responses and the various stimuli. The pair of lines labelled 1 depict the two keys during the first links of the two chains. Responses are shown as discrete vertical deflections. A peck to the left key initiates the second link of the chain on this key (lines 2a). Further responding leads to the primary reinforcer (R_L , line 3) twice and then the first links recur (lines 1). Now, a peck to the right key initiates the second link of the chain on the right key (lines 2b). Once again, further responding leads to the primary reinforcer (R_R , line 3) twice.

Table 1

Order and Duration of Experimental Conditions

Left Key		Right Key		No. of Sessions
First link	Second link	First link	Second link	
VI 1'	VR 40	VI 1'	VR 40	77
	VR 16		VR 40	63
	VR 40		VR 16	91
	VR 40		VI 30"	55
	VR 40		VI 36"	59
	VR 40		VI 15"	71
	VR 60		VI 15"	42
	VR 80		VI 15"	44
	VR 80		VI 45"	71

number of sessions apportioned to each value, and the sequence of values, are shown in Table 1.

The pertinent quantitative measures are diagrammed in Fig. 2, which shows the sequence of events during two typical cycles. Each line, except the last, represents one of the two keys. The three pairs of lines represent the three phases of the procedure. The first pair of lines shows the concurrent first links of the chains. The pair labelled 2a shows the second link for the left key; the pair labelled 2b shows this for the right key. The last line, labelled 3, shows the presentations of food. Responses are shown as discrete vertical deflections on the upper six lines. The small arrows mark those responses which actually result in some change. Thus (see pair 1), it is a response to the left key that produces

the second link during the first cycle, and a response to the right key that does so during the second cycle. During the first cycle, responses to the left key produce food (pair 2a), and, during the second cycle, responses to the right key (pair 2b). The three basic dependent variables are number of responses, duration of conditions, and number of food presentations, symbolized by the letters P, T, and R, respectively. Lower-case letters refer to the first link of a chain, upper-case letters to the second link, and the subscripts L and R, to left and right key, respectively. Thus:

- p_L = number of responses to the left key during the first link;
- p_R = number of responses to the right key during the first link;
- p_{L+R} = sum of the responses to both keys during the first link;
- p = number of responses to either key during the first link;
- P_L = number of responses to the left key during the second link;
- P_R = number of responses to the right key during the second link;
- P = number of responses to either key during the second link;
- t = duration of the first link;
- T_L = duration of the left, second link, omitting the duration of food presentations;

T_R = duration of the right, second link, omitting the duration of food presentations;
 T_{L+R} = sum of the durations of both second links, omitting the duration of food presentations;
 T = duration of either second link, omitting the duration of food presentations;
 R_L = number of food presentations produced by responses to the left key;
 R_R = number of food presentations produced by responses to the right key;
 R_{L+R} = total number of food presentations;
 R = number of food presentations produced by responses to either key.

RESULTS

Inasmuch as this experiment deals with secondary reinforcement, the dependent variable is the rate of responding during the first link of each chain. The independent variable is the frequency of primary reinforcement during the second link, but herein is found a problem of definition. In what sense, for example, does the frequency decrease when the value of a variable-ratio schedule is increased? There is, first of all, a decrease in the number of reinforcements per unit response. But there is also likely to be a decrease in the number of reinforcements per unit time. Thus, the frequency of reinforcement has decreased with respect to both number of responses and amount of time, and the two measures may not be proportional. The same ambiguity is present with changes in the value of variable-interval schedules. Two questions explored by the present experiment are, therefore, whether there is a preferable measure of frequency of reinforcement and, if so, whether it is number of responses or time between reinforcements. As a convenience, the number parameter will be referred to as "reinforcements per response" (or as "probability of reinforcement"), symbolized by the ratio R/P , and the time parameter as "rate of reinforcement", symbolized by the ratio R/T .

The possibility of choosing a preferable measure of reinforcement frequency—either R/T or R/P —arises in a particular feature of the experimental design. It is well known that at equal frequencies of reinforcement, variable-ratio schedules and variable-interval schedules produce different rates of responding. Therefore, if two such schedules are

equated for rates of reinforcement, R/T , then their reinforcements per response, R/P , will be different, and *vice versa*. In the present experiment, the second link of one chain was sometimes a variable-interval schedule while the second link of the other was a variable-ratio schedule. At these times, a choice between the two measures of reinforcement frequency could be made on the basis of the performance during the first links of the chain. Stated simply: is responding equal during the two first links when the second links have equal R/T or equal R/P ?

The logic of this experiment is predicated on faster responding on variable-ratio schedules than on variable-interval schedules at a given R/T , or at a given R/P . Figure 3 shows the extent to which this condition is fulfilled. Data for the four *Ss* have been individually plotted. The ordinates show rates of responding during the second link, and the abscissas, either the rate of reinforcement, R/T , or the reinforcements per response, R/P . Each point is the median rate over the last 15 sessions of each experimental condition. The filled points show the rates of responding on variable-ratio schedules; the open points, on variable-interval schedules. Eighteen points are plotted for each *S*, since there were nine pairs of second links during the experiment (see Table 1). For all four *Ss*, variable-ratio schedules produced faster responding than variable-interval schedules at either a given R/T or a given R/P . The extent of the rate difference varies among the pigeons, being minimal for pigeon No. 67. Rate of responding does not, however, appear to be greatly influenced by changes in the value of either R/T or R/P , as suggested by the low slope of any plausible curve that might be drawn through either the filled or open points.

With the basic logical requirement of the experiment satisfied, it now remains to be shown which of the measures of reinforcement frequency is preferable. Figure 4 gives the relative frequency of responding to a key during the first link as a function of R/T and R/P . The nine data points, corresponding to the nine experimental conditions, are medians from 15 sessions. The coordinates are each plotted as relative measures, going from 0 to 1.0. The lighter 45-degree line traces the path of perfect matching between relative frequency of responding during the first link and

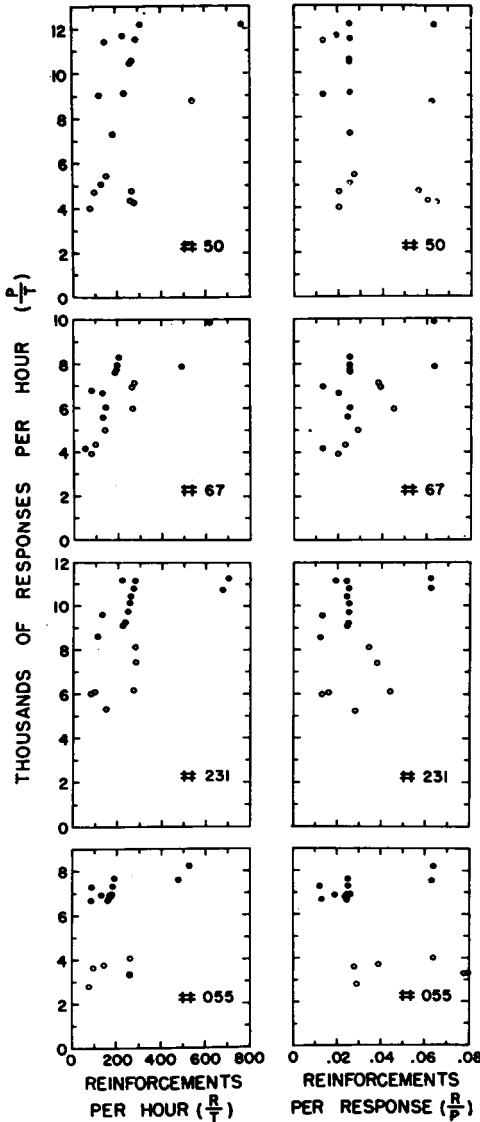


Fig. 3. Rate of responding during the second link as a function of the rate of primary reinforcement (left column of graphs) and the probability of primary reinforcement (right column of graphs). Open points denote responding on variable-interval schedules of reinforcement; filled points that on variable-ratio schedules.

relative frequency of primary reinforcement (by each of the two definitions of this concept) during the second. The thicker line is the calculated linear regression line through the data points. The linear equations are on each plot, as well as the standard deviations around regression, $\sigma_{y.x}$. For all Ss, there is less scatter around regression with R/T as the independent variable than with R/P. In addition,

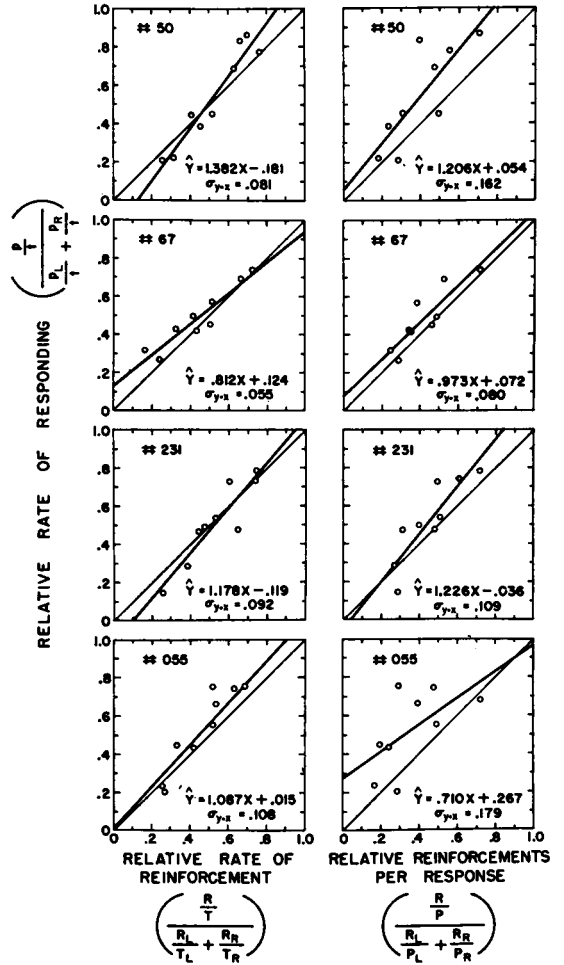


Fig. 4. Relative rate of responding during the first link as a function of the relative rate of primary reinforcement (left column of graphs) and the relative probability of primary reinforcement (right column of graphs). The lighter diagonal lines show the locus of perfect matching between ordinate and abscissa. The heavier lines are the calculated linear regression lines through the data points. The calculated linear equations and the standard deviations around the regression lines are shown on each graph.

R/T gives regression lines that appear, on the whole, to approximate the perfect-matching function better than does R/P. This last assertion is more clearly supported in Fig. 5, which gives the average regression lines for the four Ss. Again, the perfect-matching line is shown; the solid, heavy line is the average regression line for R/T, while the dashed line is that for R/P. The R/T function is, at all values of the abscissa, closer to the perfect-matching line than the R/P function.

Figures 4 and 5 show the relative frequency of responding during the first link as a function of frequency of reinforcement; Fig. 6 presents the absolute frequencies (i.e., rates of responding). Fifteen-session median rates, and the best-fit linear regression lines, are again plotted. Once again, R/T gives linear regression lines with the smaller standard deviations. In Fig. 7 are plotted the average linear regression lines for the four Ss, the solid line for R/T, the dashed line for R/P. Reinforcements per time, R/T, is clearly the preferable independent variable: it accounts for more of the variance in the data (see Fig. 6), and it yields an average curve that closely approximates a simple proportionality between the rate of responding in the first link of the chain and the relative rate of reinforcement in the second (see Fig. 7).

Is there some way of measuring the independent variable that is different from, but as good as, the relative rate of reinforcement? Figures 8 and 9 suggest a negative answer. Figure 8 shows rate of responding during the first link as a function of rate of responding during the second. Perhaps, it might be said, the Ss do not distinguish between the two

links of each chain, and whatever causes the rate to change in one link will therefore cause it to change in the other. Or, perhaps, as Premack (1959) has suggested, the effectiveness of a reinforcer is based on the rate of responding in its presence. Either view implies a function of positive slope in Fig. 8 and is thus refuted by the scatter of points.

Another view is refuted by Fig. 9. In a chain, it could be argued, the responding in all links is governed directly by the primary




Fig. 5. The average linear regression lines for relative rate of responding during the first link as a function of the relative rate of primary reinforcement (heavy solid line) and the relative probability of primary reinforcement (dashed line). The locus of perfect matching (light solid line) is shown. The subscripts of X in the equations identify the independent variables as relative rate (X_T) and relative probability (X_P).




Fig. 6. Rate of responding during the first link as a function of the relative rate of primary reinforcement (left column of graphs) and the relative probability of primary reinforcement (right column of graphs). The calculated linear regression lines, their equations, and the standard deviations around them, are all shown.




Fig. 7. The average linear regression lines for rate of responding during the first link as a function of the relative rate of primary reinforcement (solid line) and the relative probability of primary reinforcement (dashed line). The subscripts of X in the equations identify the independent variables as relative rate (X_T) and relative probability (X_p).

reinforcer and not by the successive conditioned reinforcers. This reasoning leads to the conventional experimental paradigm for the study of secondary reinforcement, in which the conditioned reinforcer is tested after pairings with the primary reinforcer have been discontinued. Figure 9 shows the 15-session-median rates during the first link as a function of both the rate of reinforcement, R/T , and the reinforcements per response, R/P , during the second link.




Fig. 8. Rate of responding during the first link as a function of rate of responding during the second link.




Fig. 9. Rate of responding during the first link as a function of rate of primary reinforcement (left column of graphs) and probability of primary reinforcement (right column of graphs).

R/P , during the second link. If the primary reinforcement were in direct control of first-link responding, orderly positive functions should have been observed. Although some sort of increasing functional relations can be seen in Fig. 9, they are considerably more erratic than those in Figs. 4 and 6.

DISCUSSION

The present experiment indicates that the strength of a secondary reinforcer is influenced by the frequency of primary reinforcement in its presence. Moreover, frequency is to be understood as a rate of reinforcement, rather than as a probability of reinforcement (the more usual term for reinforcements per re-

sponse). A quantitative statement of the relation between effectiveness of secondary reinforcement and the frequency of primary reinforcement is suggested by the data. On the basis of Fig. 5, we may conclude that the relative rate of responding during the first link of a chain is equal to the relative rate of primary reinforcement during the second link:

$$\frac{\frac{p}{t}}{\frac{P_L}{t} + \frac{P_R}{t}} = \frac{\frac{R}{T}}{\frac{R_L}{T_L} + \frac{R_R}{T_R}} \quad (1)$$

Furthermore, on the basis of Fig. 7, we may conclude that the absolute rate of responding during the first link is directly proportional to the relative rate of primary reinforcement during the second link:

$$\frac{p}{t} = k \frac{\frac{R}{T}}{\frac{R_L}{T_L} + \frac{R_R}{T_R}} \quad (2)$$

Equations 1 and 2 are equivalent if it can be shown that the total number of responses during the first link, p_{L+R} , is a constant, since the duration of the first link, t , was deliberately held constant throughout the experiment. Figure 10 shows that, in fact, the over-all




Fig. 10. Rate of the sum of the responding during the two first links as a function of the over-all rate of primary reinforcement.

rate of responding during the first link was independent of the over-all rate of reinforcement during the second link. The results in Fig. 10 may be written as

$$\frac{P_L}{t} + \frac{P_R}{t} = k \quad (3)$$

which makes equations 1 and 2 equivalent.

Equation 2 implies that the rate of responding during either one of the first links depends on the rate of primary reinforcement during both of the second. If the rate of primary reinforcement in one chain is held constant, while that in the other varies, then the rates of responding during the two first links should change in opposite directions. Of the nine pairs of reinforcement rates, five shared an approximately common value for each S . Since on variable-ratio schedules, the reinforcement rate is a function of the response rate, neither the size of this common value, nor even its presence, could have been foretold. By chance, then, it was possible to plot Fig. 11, which shows the rate of responding during the first link as a function of the rate of reinforcement during one of the second links when the rate of reinforcement during the other second link was held constant at the value shown on the figure for each S . The smooth curves were drawn according to equation 2, after ascertaining the value of k (see Equation 3) for each S . The rising curves show how the rate of responding during one first link varies with the rate of reinforcement during the second link of the same chain. The falling curves show the relation




Fig. 11. Rate of responding during the first link as a function of rate of primary reinforcement. The curves are the predictions of Equation 2 in the text. The solid lines are supposed to predict the positions of the filled points; the dashed lines, the positions of the open points.

between the rate of responding during the first link of one chain and the rate of reinforcement during the second link of the other. The points are 15-session medians and they conform reasonably well to the predicted curves.

The present experiment may be viewed as a two-key concurrent schedule of secondary reinforcement. As such, it may be compared with the more customary procedure, in which concurrent responding leads to primary reinforcement. Such a comparison, based on Herrnstein's (1961) data for the simple concurrent schedule, shows that variations in either the frequency of primary reinforcement or in the frequency with which a secondary reinforcer is paired with a primary, have essentially identical effects. In either case, there is equality between the relative rate of responding and the relative rate of reinforcement, and direct proportionality between the absolute rate of responding and the relative rate of reinforcement. In this sense, the present experiment merely equates some properties of primary reinforcement with some of secondary reinforcement. It also, however, demonstrates that frequency of reinforcement is properly measured as a rate, rather than as a probability—a conclusion implied by the earlier experiment, but not demonstrated experimentally.

The measurement of reinforcement frequency as a rate, rather than as a probability, may prove to be relevant even to experiments not concerned with secondary reinforcement. For example, many experiments purport to investigate the influence of probability of reinforcement on probability of response. In such "probability learning" experiments, the subject usually must choose between one of two response-alternatives during each of a series of discretely presented trials. To separate the trials, there is an inter-trial interval when neither the response nor the reinforcer is permitted to occur. The independent variable is always assumed to be the probability of reinforcement associated with each of the response-alternatives—in the present terminology, R_L/P_L and R_R/P_R . It may be, however, that in these experiments, as in the present one, it is R_L/T_L and R_R/T_R that are critical. When the experimenter alters R/P , he is inadvertently altering R/T as well. The resulting changes in the subject's behavior may be

better explained by considering R/T as the independent variable.

Using a modification of the present procedure, Autor (1960) obtained results favoring the measurement of reinforcement frequency as a rate. In one phase of his experiment, the second links of the two chains did not require any overt responding by the pigeons; primary reinforcement was simply delivered at specified rates, independently of the pigeon's behavior. Autor found that, even under these conditions, the responding during the first links was governed by the relative rates of primary reinforcement. If one assumes that this procedure was free of superstitious responding during the second links, then the results imply, in agreement with the present experiment, that time, and not probability, determines the effectiveness of the secondary reinforcer. Anger (1956) came to a comparable conclusion on the basis of experiments using a single-response situation and rats as subjects. He studied the distribution in time of responses reinforced on a simple variable-interval schedule and concluded that, ". . . the relative Reinforcements/Hour, not the relative Reinforcements/Interresponse Time, determine[s] the response probability" (p. 161, Anger, 1956).

Although the present formulation is encouragingly general, it is by no means a complete account of the effects of frequency of reinforcement on response-output. When the rate of primary reinforcement is zero for one of the two chains, then Equation 2 should predict responding on a single, two-link chain. As can be readily seen, Equation 2 under these circumstances becomes:

$$\frac{P}{t} = k. \quad (4)$$

Findley (1954) has studied single, two-link chains and has found that the rate of responding during the first link does not remain constant, but varies systematically with the rate of reinforcement during the second link. At the present time, no explanation can be given of this discrepancy.

REFERENCES

- Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *J. exp. Psychol.*, 1956, 53, 145-161.

- Autor, S. M. The strength of conditioned reinforcers as a function of frequency and probability of reinforcement. Unpublished doctoral dissertation, Harvard University, 1960.
- Findley, J. D. Rates of response in a two-member chain as a function of mean variable-interval schedule of reinforcement on the second member. Unpublished doctoral dissertation, Columbia University, 1954.
- Herrnstein, R. J. Relative and absolute strength of response as a function of frequency of reinforcement. *J. exp. Anal. Behav.*, 1961, 4, 267-272.
- Premack, D. Toward empirical behavior laws: I. Positive reinforcement. *Psychol. Rev.*, 1959, 66, 219-233.

Received March 5, 1963