JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1976, 26, 441-449 NUMBER 3 (NOVEMBER)

THE EFFECTS OF UNSIGNALLED DELAYED REINFORCEMENT[1]

BEN A. WILLIAMS

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Pigeons' pecks were reinforced according to a variable-interval schedule. A delay-of-reinforcement procedure was then added to the schedule, or a yoked-control procedure was arranged where the reinforcers occurred independently of responding according to the same variable-interval schedule. During the delay-of-reinforcement procedure, the first peck after a reinforcer was scheduled began a delay timer and the reinforcer was delivered at the end of the interval. No stimulus change signalled the delay interval and responses could occur during it, so that the obtained delays were often shorter than those scheduled. Responding under this procedure was highly variable but, in general, behavior was substantially reduced even with the shortest delay used, 3 sec. In addition, the rates maintained by delayed reinforcement were only slightly greater than those maintained by the yoked-control procedure, suggesting that adventitious pairings of response and reinforcer were responsible for some of the maintenance of behavior that did occur. The results challenge recent conceptions of reinforcement as involving response-reinforcer correlations and re-emphasize the role of temporal proximity between response and reinforcer.

Key words: delay of reinforcement, correlational learning, temporal contiguity, response-independent reinforcement, VI schedules, key peck, pigeons

Of all parameters of reinforcement, temporal proximity between response and reinforcer is most central to an understanding of the fundamental principles of conditioning. The premier axiom of operant conditioning has been that behavior is strengthened automatically simply by virtue of its temporal contiguity with the reinforcer. Nevertheless, the function relating response strength to the degree of temporal proximity remains open to question. For years, most investigators have agreed that delayed reinforcement effects follow a negatively accelerated exponential decay function, although they often have disagreed over the temporal limits of the function (*cf.* Mowrer, 1960). Recent work with free-operant procedures has questioned the exponential decay function, however, on the basis of evidence that reinforcement delay obeys the matching law (Chung and Herrnstein, 1967). As noted by these investigators, the matching of relative response rate to relative reinforcement immediacy is incompatible with an exponential decay function (but see Herbart, 1970). The result has been that reinforcement delay is regarded by some as simply one more determinant of reinforcement "value" (Baum and Rachlin, 1969), and thus no more central to the response-reinforcer contingency than other parameters of reinforcement.

The challenge to the exponential decay notion of delayed reinforcement is one of several factors leading to a questioning of the supposedly automatic effects of response-reinforcer temporal contiguity. Several writers now dispute temporal contiguity as the fundamental determinant of conditioning and have substituted in its place an emphasis on correlations between molar response rates and molar reinforcement rates (Baum, 1973; Rachlin, 1970). According to this view, a reinforcer does not strengthen behavior simply because it is temporally contiguous, but rather only insofar as greater temporal contiguity is related to a stronger correlation between behavior and reinforcement. Thus, contiguity between response and reinforcer is not a necessary condition for reinforcement to occur; it is instead a derived condition subsumed under the more general notion that molar correlations are the fundamental units of behavior.

---

[1] This research was supported by National Science Foundation research grant GB-42887 to the University of California, San Diego. Reprints may be obtained from the author, Department of Psychology C-009, University of California, San Diego, La Jolla, California 92093.

The major evidence in favor of the correlation notion is the successful demonstration of regular functions between response rate and reinforcement rate in a variety of free-operant situations (*cf.* Herrnstein, 1970). While such demonstrations in themselves do not preclude a contiguity-based reinforcement principle, their continued resistance to explanation by more molecular conditioning principles argues against the adequacy of a contiguity conception. Both historically and conceptually, temporal contiguity as the fundamental aspect of conditioning has been tied to the view that molar behavior can be reduced to some set of elementary events (*e.g.,* each response-reinforcer pairing increases response strength, whereas each response–no-reinforcer pairing decreases it). To the extent that no such reduction is successful, the molecular principles supposedly governing the elementary events, of which temporal contiguity is one, are themselves weakened. As of now, the reduction of free-operant behavior to molecular principles has been notably unsuccessful (*cf.* Jenkins, 1970).

Quite apart from the general success of a molar *versus* molecular approach, the empirical status of delayed reinforcement effects remains to be resolved. While several studies have examined delay of reinforcement, almost all have been confounded by changes in the response-reinforcer correlation. The most typical delay-of-reinforcement procedure, for example, involves removing the response opportunity during the delay interval (*cf.* Ferster, 1953). This interjects a cued temporal period where no reinforcement is delivered, so that it is impossible to determine if effects are due to the temporal delay itself or to the properties of the stimulus change. The other frequently used delay procedure has similar difficulties. Here, no stimulus change cues the delay interval, but instead each response during the delay interval resets the delay timer (a DRO contingency). Again, delayed reinforcement effects are confounded by changes in the response contingency.

Perhaps the only delay-of-reinforcement procedure that unconfounds correlation and contiguity effects is that in which no stimulus change or additional response requirement is interposed during the delay interval. Thus, whenever a reinforcer is scheduled, the next response begins the delay timer and the reinforcer is delivered at the end of the interval, regardless of the animal's behavior during the interval. The disadvantage of the procedure, of course, is that the obtained delay values differ from those scheduled because responses are likely to occur during the delay interval. The general effects of the procedure still should be of interest, however, because it leaves the dependency of reinforcement on responding unchanged. To the extent response rate is reduced by the procedure, therefore, the effect must depend on some facet of changes in the temporal separation of response and reinforcer.

The present study examined the effects of the just-described delay-of-reinforcement procedure on VI performance. For Experiment I, the delay values investigated were 3, 8, and 15 sec; for Experiment II, the value was 5 sec. In addition, for both experiments, a yoked procedure was used to control for the effects of the reinforcers when they were delivered independently of responding.

METHOD

Subjects

Sixteen White Carneaux pigeons, maintained at 80% of their free-feeding weights, had experience in a variety of experimental situations.

Apparatus

Two standard conditioning chambers were constructed from plastic picnic chests. The bird's chamber for each box was approximately 30.5 cm in all dimensions. On the front panel was mounted a single transparent pigeon key, 1.7 cm in diameter, which required a force of at least 0.10 to 0.12 N for operation. The stimulus, a diffuse yellow light, was projected onto the rear of the key by a standard 28-V, 12-stimulus projector. Ten cm below the key was a 5- by 5-cm aperture through which the birds were fed when the grain magazine was operated. Located on the rear wall of the inner chamber was a 28-V houselight that was illuminated at all times during a session.

Procedure: Experiment I

In the first session, all responses were reinforced, to a total of 50. The schedule was changed gradually during the next four sessions through increasing VI values until a VI 2-min schedule was in effect. An additional six sessions of training were given with the VI 2-min schedule before beginning the first delayed-reinforcement training. During this baseline training, the first peck after a reinforcer was scheduled was followed immediately by 4-sec access to the grain hopper. The schedule consisted of 18 intervals constructed from the distribution of Fleshler and Hoffman (1962).

During the delayed-reinforcement procedure, the same VI 2-min schedule remained in effect. Now, however, the first peck after a reinforcer was scheduled started a delay timer, and the reinforcer was delivered when the interval had ended. The VI tape restarted once the delay timer began and thus continued to run during the delay and during delivery of the reinforcer. If a second reinforcer set up during the delay interval, it was cancelled. The delay values used were 3, 8, and 15 sec. The order of the values was counterbalanced across subjects, and each value was used for 10 consecutive sessions.

Pigeons were paired so that the schedule of reinforcement for a second pigeon, in a second chamber, was yoked to each pigeon receiving the delayed reinforcement condition. The yoking procedure consisted of presenting a response-independent reinforcer whenever the delayed-reinforcement subject received its response-dependent reinforcer. Each pigeon received both the delayed reinforcement and yoked procedure in alternating periods of training. Separating each experimental condition were five sessions of baseline retraining, where all pigeons received response-dependent immediate reinforcement on the VI schedule. Each pigeon was trained in the same chamber throughout the experiment, so that the two chambers were alternately used for the delayed reinforcement and yoked procedure. Table 1 summarizes the training schedule for the 12 subjects used in Experiment I.

Sessions ended after 55 min. Throughout training, the first peck of each session produced a reinforcer to reduce warm-up effects.

Procedure: Experiment II

Four subjects from Experiment I and four new subjects, all with extensive experience, were used. Intervening between the training of Experiment I was a six-month period of training on a multiple schedule involving immediate reinforcement.

Because all pigeons had extensive experience with intermittent reinforcement, they were placed immediately on a VI 2-min schedule with immediate reinforcement. A total of 21 sessions of baseline training was given where each session terminated after 90 min.

At the end of baseline training, four pigeons were assigned to the same delay-of-reinforcement condition, with a 5-sec delay, as in Experiment I. The schedules for each of the remaining pigeons were yoked to one of the delay subjects, as in Experiment I. After 30 sessions

Table 1

Schedule of training for each subject. Subjects worked in pairs with each member alternately serving as a delay-of-reinforcement subject and then as a yoked subject. Each experimental condition continued for 10 sessions, followed by five sessions of retraining on the baseline VI 2-min schedule.

	Order					
Subject	*1*	*2*	*3*	*4*	*5*	*6*
R-12	yoked	3 sec	yoked	8 sec	yoked	15 sec
R-20	3 sec	yoked	8 sec	yoked	15 sec	yoked
B-1	yoked	15 sec	yoked	3 sec	yoked	8 sec
R-2	15 sec	yoked	3 sec	yoked	8 sec	yoked
B-16	3 sec	yoked	15 sec	yoked	8 sec	yoked
R-5	yoked	3 sec	yoked	15 sec	yoked	8 sec
R-13	yoked	8 sec	yoked	15 sec	yoked	3 sec
B-5	8 sec	yoked	15 sec	yoked	3 sec	yoked
R-11	yoked	15 sec	yoked	8 sec	yoked	3 sec
Y-22	15 sec	yoked	8 sec	yoked	3 sec	yoked
Y-91	yoked	8 sec	yoked	3 sec	yoked	15 sec
B-3	8 sec	yoked	3 sec	yoked	15 sec	yoked

of training, the procedures for the pairs of subjects were reversed. Thus, the subject that had received the yoked procedure now received the delay-of-reinforcement procedure, and *vice versa*. Thirty additional sessions were presented.

RESULTS

Experiment I

Table 2 shows the rate of responding during the last two sessions of each condition, as well as the mean rate averaged across the last sessions of each baseline retraining period. The data were highly variable across both subjects and conditions. The most likely source of the variability was the delay contingency, which specified only the *maximum* delay between response and reinforcer. Thus, subjects could fail to come into contact with the scheduled delays by responding consistently throughout the delay intervals (*e.g.*, B-5, for the 15-sec condition). Such an interpretation is supported by the similar degree of variability produced by the yoked procedure involving response-independent reinforcement, where presumably any "strengthening" effects of reinforcement were due to occasional contiguous occurrences of the response and reinforcer.

In spite of the variability, the main effect of delayed reinforcement is evident. Namely, regardless of the delay value, the median response rate was reduced by 70 to 80%. Even the low response rate that was maintained cannot be ascribed unequivocally to the response-reinforcer contingency, because similar low rates of responding also were maintained by the yoked procedure. Apparently, response-independent reinforcement was almost as effective in maintaining the behavior, at least within the limits of the number of training sessions employed.

Two aspects of the data warrant more detailed consideration. The first concerns whether the differences in delay values caused differences in behavior. Some slight evidence of a greater decrease with longer delays is suggested in Table 2, but this difference was not statistically reliable ($F < 1$). In addition, the control conditions also produced a gradient, although it too did not approach statistical reliability ($F < 1$).

A second issue is whether behavior was maintained at a higher level with the delayed reinforcement condition than with the yoked condition. Table 2 indicates a slight effect immersed in great variability. Two comparisons are available to test the difference statistically. For a between-subject comparison, the absolute response rates of each subject were normalized with respect to its VI baseline rate and then the normalized rate for each delayed-reinforcement subject was compared to the normalized rate of its yoked partner. The difference scores were then submitted to a *t* test and found not to be statistically reliable ($t = 1.42$, $p > 0.05$). For a within-subject comparison, the absolute rates of responding for all three delayed-reinforcement values were averaged and compared to the average rate for

Table 2

Mean response rate (responses per minute) for individual subjects during the last two sessions of each experimental condition. Also shown is the mean baseline rate, taken as the average of the last session of each baseline condition.

		Response-contingent			*Yoked*		
Subject	*Baseline*	*3 sec*	*8 sec*	*15 sec*	*3 sec*	*8 sec*	*15 sec*
R-12	77.2	41.2	32.1	16.2	55.2	34.9	12.9
R-20	79.7	20.8	19.8	30.5	18.9	10.4	25.4
B-1	54.5	10.8	15.1	16.6	9.9	10.2	2.7
R-2	39.1	17.1	13.6	14.5	22.9	15.5	9.4
B16	35.0	7.5	1.9	5.7	4.7	2.2	3.6
R-5	41.7	5.3	11.9	5.0	4.4	9.9	3.1
R-13	52.3	30.5	8.7	20.9	2.5	2.2	4.4
B-5	53.5	15.9	32.1	64.9	5.2	41.3	22.5
R-11	52.8	17.2	8.1	10.6	8.5	21.8	12.1
Y-22	62.4	15.1	14.9	9.7	19.2	7.5	4.3
Y-91	77.1	24.0	8.2	11.8	19.3	8.1	6.8
B-3	73.6	12.3	7.2	10.0	0.9	2.0	7.8
Median	54.0	16.5	12.8	13.2	12.1	10.1	7.3

the three yoked conditions. Here, the difference was statistically significant ($t = 2.37$, 11 df, $p < 0.05$). A cautionary note about the second test is in order, however. As is evident from Table 2, it was possible to have large differences in the degree of responding between yoked and delayed reinforcement subjects. A subject that responded at a substantial rate during the delay condition might thus have its rate reduced during the yoked condition by the failure of its partner to respond consistently during the delay condition, since such a failure would change both the number and distribution of reinforcers. A difference for the same subject between the delayed reinforcement and yoked condition might not reflect the contingency of reinforcement, therefore, but instead differences in the obtained distribution of reinforcement.

A final question concerning Experiment I is whether the actually obtained rates of reinforcement were affected by the delay-of-reinforcement contingency. Since all 12 subjects were exposed to three delay values, there were 36 experimental conditions available for comparison. During only four of these were obtained rates of reinforcement less than 25 per session, compared to the median rate during the VI baseline of 28.4. The median rates of reinforcement were 27.5, 26.5, and 26.5 for the 3, 8, and 15-sec conditions, respectively.

Experiment II

Experiment II differed from Experiment I mainly in the longer training given for each experimental condition. Figure 1 shows the individual subject data normalized with respect to the rates of responding during the baseline condition; Table 3 summarizes the absolute rates of responding for the last five sessions of each condition.

In general, the overall rate of responding was higher under the delayed reinforcement condition than under the yoked condition. The higher rate of responding occurred for the delayed-reinforcement subject for six of the eight subject pairings. The exceptions were R-1/B-19 and R-2/R-18. For the within-subject comparison, seven of the eight subjects had higher rates under the delayed condition than under the yoked condition (R-18 was the exception).

The absolute level of responding maintained by the two conditions varied considerably across subjects. Under the delayed reinforcement condition, four of the eight subjects (R-12, R-1, B-19, Y-91) had stable rates of responding that were less than 10% of their baseline values. The remaining subjects had higher rates (20 to 40% of the baseline) but this varied considerably from session to session. For the yoked condition, only three subjects had any significant responding at the end of the 30 sessions of training (R-18, B-19, R-2). One of these, R-18, continued to respond at a rate that was 30 to 50% of its VI baseline.

A major feature of the data was the variability. The rate of responding often varied by a factor of 2 to 3 over consecutive sessions. This variability generally was greater under the delayed-reinforcement condition, but was not restricted to it. For example, R-18 displayed a pattern of variability under the yoked condition highly similar to that exhibited under the delay condition.

The variability that occurred was not random from session to session but often cyclic. That is, increases in rate typically were sustained for a few sessions followed by a sustained decline. This pattern occurred not only for delay subjects but for yoked subjects as well. Most notable was R-2, where the first five sessions of the yoked condition produced a regular decline in responding, followed by a regular increase in responding over the next 10 sessions to approximately 50% of the baseline rate.

DISCUSSION

The major finding was that interjecting an unsignalled delay between response and reinforcer substantially reduces rate of responding

Table 3

Mean response rate (responses per minute) over the last five sessions of each experimental condition for individual subjects.

Subject	*VI Baseline*	*5-sec Delay*	*Yoked*
R-1	60.1	5.1	0
R-12	78.4	6.8	0.1
R-9	40.4	1.3	0.6
B-19	38.1	9.9	4.4
R-18	36.0	15.6	15.8
R-2	43.5	16.8	3.3
R-5	26.2	8.6	0.6
Y-91	66.0	5.3	0.3
Median	42.0	7.7	0.6

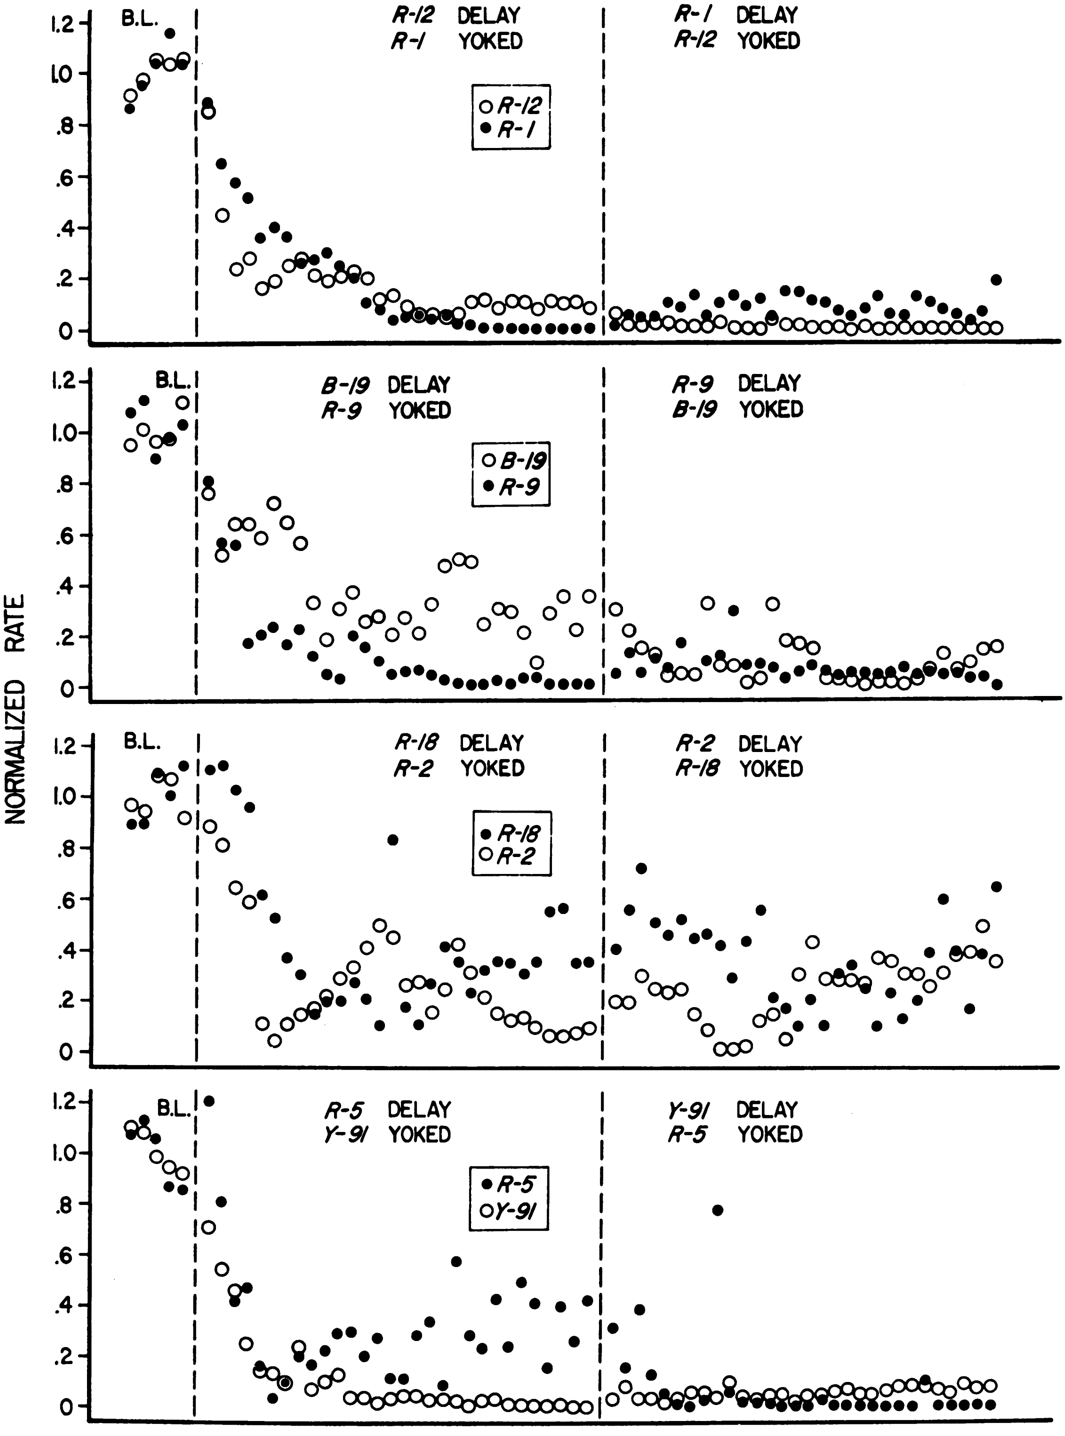

Fig. 1. Rate of responding for the VI baseline, the delayed reinforcement, and yoked conditions. In each panel are pairs of subjects yoked together. Response rates were normalized by calculating the average rate during the last five sessions of the baseline training and dividing that value into the response rates for individual sessions.

in a free-operant situation. The 70 to 80% reduction in rate with the shortest delay (3 sec in Experiment I) is particularly remarkable because 3 sec was only the nominal value. The pigeons could respond during the delay interval, so that the delays actually obtained could be much shorter.

The delayed-reinforcement effects found in the present study stand in striking contrast to the effects found previously with free-operant procedures. Delays shorter than 10 sec have produced little reduction in response rates in a variety of situations, including simple schedules (Ferster, 1953; Pierce, Hanford, and Zimmerman, 1972), multiple schedules (Nevin, 1974; Richards, 1972; Wilkie, 1971), and concurrent schedules (Chung, 1965; Neuringer, 1969). Such studies typically used procedures where the delay was signalled by some stimulus change (*e.g.*, a blackout during the delay interval). Although no direct comparison was made between the signalled-delay procedure and the present unsignalled-delay procedure, the magnitude of the reduction in responding found in the present study suggests strongly that unsignalled delays maintain much less responding. In the only previous experiment to use an unsignalled-delay procedure (Dews, 1960), behavior was maintained with the procedure up to delays of 100 sec. A comparison with the present study is difficult, however, because Dews used a continuous reinforcement schedule in contrast to the present VI, and because he reported no baseline data in which reinforcement was presented at zero delay.

The major question raised by the results is their implication for the underlying basis of the effects of delayed reinforcement. In particular, are the effects best understood by the concept of correlation, such as that proposed by Baum (1973), or by the traditional theory based solely on the temporal proximity between response and reinforcer?

At one level of analysis, the data provide strong evidence against the correlation view. That is, the *overall* correlation between response rate and reinforcement rate was the same for the delayed-reinforcement condition as for the VI baseline condition with immediate reinforcement, and yet responding under the delayed reinforcement was severely reduced. This analysis cannot be taken literally as evidence against the correlation concept, however, because it need not be free of temporal parameters. Baum (1973) argued, for example, that the goodness of the correlation will be affected by a delay between response and reinforcer if the size of the temporal sample used to construct the correlation coefficient is relatively short. Thus, with delayed reinforcement, the response causing the reinforcer may occur in one temporal sample and the reinforcer itself may occur in the succeeding sample. The correlation coefficient would thus be reduced, in spite of the fact that the actual dependency of reinforcement on responding remained unchanged. According to Baum, therefore, *both* the correlation and temporal contiguity notions would predict a decline in responding under delayed reinforcement.

At a different level of analysis, the data seem to support the correlation concept. That is, with sufficient training, the rate of responding under response-independent reinforcement did generally decline to near-zero levels and was maintained under delayed reinforcement. However, the two conditions were very similar in the early sessions of training (as with Experiment I) and responding under response-independent reinforcement was sustained for some subjects. In addition, a difference between the two conditions need not be taken as support for the correlation view, because it is predicted equally well by the effects of temporal contiguity. This can be seen by considering the feedback cycle begun for the two conditions on the change from the VI baseline. In both cases, at the beginning of training, the average delay of reinforcement was equal to one-half of the average interresponse time. Since substantial responding occurred under the VI baseline (40 to 50 responses per minute) the average delay was slightly less than 1 sec. This delay in turn reduced the average rate of responding still further, which in turn decreased the average delay of reinforcement. With the response-independent reinforcement, this feedback cycle was unlimited, so that the average delay of reinforcement should gradually increase and all responding should eventually cease. With the delayed reinforcement condition, however, an upper limit on the maximum delay was imposed by the delay-of-reinforcement interval, and responding should thus be maintained indefinitely.

The results of Experiments I and II are consistent with such an analysis. In Experiment I, little difference occurred between the delay and yoked conditions, presumably because sufficient responding was still maintained under the yoked condition to guarantee some degree of temporal proximity between response and reinforcer whenever the reinforcers were freely presented. With Experiment II, however, a larger difference did occur, since the behavior under the yoked condition was allowed to continue further into the feedback cycle just described.

If the obtained overall difference between the delay and yoked conditions is of no importance, what then remains to distinguish between the correlation and temporal contiguity theories? As noted by Baum, it may be impossible to distinguish empirically between the two conceptions. It nevertheless may be instructive to attempt such a distinction to make explicit the assumptions underlying the correlation concept.

Consider the overall level of responding maintained by each of the three experimental conditions. To the extent that the correlation concept is meaningful, differences in the rate of responding should be related to the degree of scatter (the correlation coefficient) around the regression of rate of reinforcement on rate of responding. As noted by Baum, the degree of scatter in turn will be a function of the size of the time sample used to construct the correlation. Although there is no way of specifying in advance what the size of the time sample actually is, limitations are imposed on the range of possible values by the actually obtained results. If the sample size were smaller than the scheduled delay, the correlation coefficient would always be zero, and thus like that of response-independent reinforcement. This is true because no one sample would include both the reinforcer and the response causing the reinforcer. On the other hand, as the size of the time sample becomes larger than the scheduled delay, the correlation coefficient approaches that of the VI schedule of immediate reinforcement. The limiting case is when the sample size is sufficiently large *always* to include both the reinforcer and the response causing the reinforcer. If it is assumed that responses are independent of each other, it can be shown that the response-reinforcer correlation can be specified as a limit function with the size of the temporal sample and the scheduled delay of reinforcement as the two parameters. If $s =$ the temporal sample, and $d =$ the scheduled delay, the function is $s - d/s + d$. The function represents the probability that the entire response-reinforcer interval will be included in the sample, given that at least part of the interval is included. Its best interpretation is the degree to which the correlation under delayed reinforcement approximates that with the VI schedule of immediate reinforcement.

The advantage of specifying the above function is that it determines what the size of the temporal sample must have been, given the behavior produced by a particular delay of reinforcement, assuming that rate of responding and the correlation coefficient represent fundamental scales of measurement. Consider, for example, the 3-sec delay condition of Experiment I. There, the median rate of responding maintained by the schedule was 0.27 of the VI baseline. As this should correspond to the degree of correlation relative to the VI schedule, the value can be inserted into the above equation, which then can be solved for the temporal sample. In this case, the sample size would be 6.0 sec. Such a value seems plausible. The problem, however, is its generality. Consider the results for the 15-sec delay condition of Experiment I. There, the median rate was 0.23 of the VI baseline. When that value is inserted into the equation, the necessary sample size would be 24 sec. It is apparent, therefore, that to account for the data produced by the different delay conditions of Experiment I, the size of the temporal sample must be free to vary across conditions. Such an assumption seems highly implausible, because the temporal sample presumably should be a fixed property of the pigeon that should be constant across experimental conditions.

The preceding analysis of course need not be accepted literally. It is likely that rate of responding would be related to the index of correlation only ordinally, in which case the algebraic formulation has little empirical utility. Moreover, it is possible that the organism uses not just a single size of temporal sample, but a distribution of samples, which in turn might be differentially weighted in the computation of the correlation coefficient. Once such concessions are made, however, the concept of correlation ceases to be empirically testable. At such a point, arguments in its favor must rest solely on its theoretical and conceptual merits (*e.g.*, parsimony).

There are other empirical reasons for rejecting correlations as the underlying basis of the principle of reinforcement. One concerns the maintenance of responding of at least some of the pigeons under the response-independent reinforcement. It can be argued, of course, that the 30 sessions were not sufficient to produce stable behavior. While it is indeed likely that all behavior under the yoked condition eventually would have ceased (but see Lattal, 1972), such an appeal restricts the domain of the principle of reinforcement solely to "steady-state" behavior. The same general problem arises for the session-to-session variability. Such variability would be expected according to the temporal contiguity conception, because the actually obtained delays between responses and reinforcers were themselves variable. The heart of the correlation concept, on the other hand, is that moment-to-moment variations are of little significance because correlations represent an average of the relations occurring over multiple observations.

An example of the difficulty with the correlation concept is the pattern of variability that occurred for Subject R-2 (third panel of Figure 1) when trained under response-independent reinforcement. When transferred to the yoked condition, its rate decreased regularly over the first five sessions, then regularly increased over the next 10 sessions. The cycle then repeated once again, although with a shorter period. Because by definition there was no correlation between the response rate and reinforcement rate, this regularity in the changes in no way can be explained by the correlation concept. It is easily explained, however, by the temporal relations between responses and reinforcers. One need only assume that at some point, a chance pairing of the response and reinforcer occurred. This would then cause the average rate of responding to increase, which in turn would decrease the average delay of reinforcement. Such a feedback cycle would continue until broken by the chance occurrence of a sustained period without immediate pairings of the response and reinforcer. A second feedback cycle would then begin working in the opposite direction. While the postulation of such a mechanism is admittedly *post hoc*, without it the principle of reinforcement is reduced to an extremely weak principle that accounts for little of the variance in the actual behavior.

REFERENCES

Baum, W. M. The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, 1973, **20**, 137-153.

Baum, W. M. and Rachlin, H. Choice as time allocation. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 861-874.

Chung, S. Effects of delayed reinforcement in a concurrent situation. *Journal of the Experimental Analysis of Behavior*, 1965, **8**, 439-444.

Chung, S. and Herrnstein, R. J. Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 67-74.

Dews, P. B. Free-operant behavior under conditions of delayed reinforcement: I. CRF-type schedules. *Journal of the Experimental Analysis of Behavior*, 1960, **3**, 221-234.

Ferster, C. B. Sustained behavior under delayed reinforcement. *Journal of Experimental Psychology*, 1953, **45**, 218-224.

Fleshler, M. and Hoffman, H. S. A progression for generation of variable interval schedules. *Journal of the Experimental Analysis of Behavior*, 1962, **5**, 529-530.

Herbart, E. W. Two-key concurrent responding: response-reinforcement dependencies and blackouts. *Journal of the Experimental Analysis of Behavior*, 1970, **14**, 61-70.

Herrnstein, R. J. On the law of effect. *Journal of the Experimental Analysis of Behavior*, 1970, **13**, 243-266.

Jenkins, H. M. Sequential organization in schedules of reinforcement. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules*. New York: Appleton-Century-Crofts, 1970. Pp. 63-109.

Lattal, K. A. Response-reinforcer independence and conventional extinction after fixed-interval and variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 1972, **18**, 133-140.

Mowrer, O. H. *Learning theory and behavior*. New York: Wiley, 1960.

Neuringer, A. J. Delayed reinforcement *versus* reinforcement after a fixed interval. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 375-383.

Nevin, J. A. Response strength in multiple schedules. *Journal of the Experimental Analysis of Behavior*, 1974, **21**, 389-408.

Pierce, C. H., Hanford, P. V., and Zimmerman, J. Effects of different delay-of-reinforcement procedures on variable-interval responding. *Journal of the Experimental Analysis of Behavior*, 1972, **18**, 141-146.

Rachlin, H. *Introduction to modern behaviorism*. San Francisco: W. H. Freeman, 1970.

Rachlin, H. and Baum, W. M. Effects of alternative reinforcement: does the source matter? *Journal of the Experimental Analysis of Behavior*, 1972, **18**, 231-241.

Richards, R. W. Reinforcement delay: some effects on behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1972, **17**, 381-394.

Wilkie, D. M. Delayed reinforcement in a multiple schedule. *Journal of the Experimental Analysis of Behavior*, 1971, **16**, 233-239.

Received 15 January 1976.
(Final Acceptance 14 May 1976.)