REINFORCEMENT OF BEHAVIORAL PATTERNS: SHAPING A SCALLOP¹

LARRY HAWKES AND CHARLES P. SHIMP

UNIVERSITY OF UTAH

Temporal patterns of key pecking by pigeons were shaped by a schedule in which the delivery of food was contingent upon a measure of the overall extent to which the temporal pattern of behavior within a 5-sec trial conformed to a required pattern. This pattern approximated a constant rate of change in the rate of key pecking throughout the 5-sec trial. In comparison with behavior maintained by a classical fixed-interval 5-sec schedule, the new schedule controlled a better approximation to a "scallop" within individual trials and greatly reduced intersubject variability. These results are consistent with the view that the delivery of a reinforcer after a behavioral pattern a few seconds in duration may strengthen the entire pattern as a unit, or operant. The response topography contiguous with reinforcement may be a negligible fraction of the strengthened operant. One implication of this view is that mean response rate for such brief responses as key pecks and lever presses is a byproduct of whatever patterns are strengthened, and generally will not reveal fundamental controlling relationships, whenever a reinforcer is not contiguous with all the behavior on which it is contingent.

A fundamental issue in a science of behavior must be the nature of the fundamental behavioral unit (Skinner, 1938). This unit, or operant as it is called here, may be defined to be the unit in terms of which functions relating behavior to reinforcement parameters assume their most pleasing, simple, forms. The practical working definition of an operant has been different. It is traditionally assumed that the operant as defined above is the same as the behavior upon which the delivery of a reinforcer is contingent. Thus, the traditional view maintains that a key peck or lever press is the fundamental behavioral unit in schedules such as fixed-ratio, variable-ratio, fixed-interval, and variable-interval schedules, because these are the behaviors in terms of which the reinforcement contingencies are defined (Skinner, 1938). However, all of these contingencies share one property, and since operants traditionally have been defined in terms of reinforcement contingencies, this property also has been attributed to operants in general. This property is that of temporal contiguity

between response and reinforcer. It is traditionally assumed that this temporal contiguity is essential for learning to take place, and that therefore, when larger behavioral patterns are observed, they must be maintained by additional processes such as delay-of-reinforcement gradients (Catania, 1971; Dews, 1966, 1969) or "mediating" chains of stimuli and responses (Ferster and Skinner, 1957).

The traditional assumption that an operant is temporally contiguous with a reinforcer can be incompatible with the traditional definition of an operant, because not all reinforcement contingencies are such that a reinforcer is temporally contiguous with all the behavior that produces it. A familiar example of such a contingency is an interresponse-time contingency. With any contingency such as this, one confronts a dilemma. On the one hand, the traditional definition of an operant holds that it is the behavior upon which the delivery of a reinforcer is contingent that is the fundamental behavioral unit, or operant. One would therefore conclude that interresponse-time classes would form the fundamental units in the case of interresponse-time contingencies. But on the other hand, the assumption of response-reinforcer temporal contiguity would imply that only a key peck or lever press terminating an interresponse time could be an oper-

¹This research was supported in part by NIMH grants 16928 and 19901 and by a fellowship from the Center for Advanced Study in the Behavioral Sciences to the second author. Reprints may be obtained from Charles P. Shimp, Department of Psychology, The University of Utah, Salt Lake City, Utah 84112.

ant. In order to resolve the dilemma, we must return to the basic definition of the fundamental behavioral unit as that which provides us with the most parsimonious set of functions relating behavior and reinforcement parameters. That is, to resolve the dilemma, one must empirically ascertain which behavioral unit leads to the most harmonious view of behavioral controlling relations. Observe that classical schedules are of little or no value in resolving the dilemma because they confound the behavior on which reinforcement is contingent with the behavior temporally contiguous with reinforcement.

One is therefore required to broaden reinforcement contingencies beyond those of classical schedules. In particular, one needs to unconfound the behavior on which reinforcement depends and the behavior temporally contiguous with reinforcement. We have noted above that interresponse-time contingencies satisfy this need. With such a contingency, a reinforcer is delivered for a member of response class A, a key peck, only if a mutually exclusive response class B, any response topography but a key peck, consumes a specific amount of time immediately preceding the instance of class A. In this sense, such a contingency specifies the response topography at each moment in time for some interval preceding reinforcement. Thus, an interresponse-time contingency requires a specified sequence, or temporal pattern, of two response classes, and a reinforcer is contiguous with only one of the two classes. Indeed, a reinforcer is temporally contiguous with only a very small fraction of the behavior on which it is contingent. Does the literature on interresponse-time contingencies suggest that we abandon the assumption that temporal contiguity is essential to the definition of an operant? Stated differently, does this literature suggest that interresponse times can be fundamental behavioral units? The data make a rather cogent case for an answer of "Yes". We are not aware in the literature on interresponse-time contingencies of any elegantly simple functional relationship with a key peck or lever press as the behavioral unit. However, when units of interresponse times are used, simple functional relations, including matching, may be obtained (Shimp, 1969, 1974). These results support the utility of the traditional definition of an operant as that behavior on which rein-

forcement is contingent, and also suggest that the assumption of temporal contiguity is not essential to the definition of an operant.

The preceding discussion implies a need to extend our understanding of the effects of reinforcement contingencies in which reinforcement is not contiguous with all the behavior on which it depends. Our understanding of such contingencies presently leans excessively on interresponse-time contingencies. The purpose of the present experiment was to generalize these contingencies to more complex temporal patterns of behavior. Unlike an interresponse-time contingency, the contingency developed here did not specify which of two response topographies had to occur at each moment over an interval of time preceding the delivery of a reinforcer. Instead, the present method somewhat relaxed this requirement. The method required only that rates of key pecking within narrow subintervals of 5-sec periods conformed sufficiently closely to required rates of key pecking in those subintervals. Thus, a pigeon was rewarded at the end of a 5-sec trial only if the pattern of key pecking throughout the 5-sec trial sufficiently conformed to a required pattern. The required patterns consisted of key pecking at a rate that changed at a constant rate throughout the trial.

METHOD

Subjects

Three experimentally naive male White Carneaux pigeons were maintained at 80 to 90% of their free-feeding weights (see Table 1).

Apparatus

Only the left keys were used in three standard two- or three-key Lehigh Valley Electronics pigeon chambers. These chambers were interfaced to a Digital Equipment Corporation PDP-8E computer, which arranged all experimental contingencies and recorded the data on magnetic tape for subsequent analysis. White noise helped to mask noises from outside the chambers.

Procedure

The general objective of the procedure was to shape desired temporal patterns of key pecking. That is, the procedure made the de-

livery of food contingent upon whether or not a *pattern* of key pecking within a 5-sec "trial" met an overall criterion, *i.e.*, upon whether or not an entire pattern produced by a subject closely approximated a particular required pattern. Reinforcement was *not* contingent upon the occurrence or nonoccurrence of any single response at any particular moment: it depended instead on how well an entire temporal pattern of responses approximated, or matched, a required pattern. An index of overall "goodness-of-fit" estimated the extent to which, on any particular trial, the observed pattern matched the required pattern. Food was delivered at the end of a trial, provided that this match was sufficiently close. This procedure is described below in detail.

Required Response Patterns

The required temporal pattern consisted, in different conditions, of an increasing or a decreasing response rate. That is, a subject was required to produce a temporal pattern of key pecks which, when plotted on a cumulative record, showed a "scallop"—either a concave-upwards or a concave-downwards curve. In a baseline condition, the reinforced temporal pattern was not controlled, and reinforcement

was delivered simply for the first response after 5 sec elapsed since the beginning of a trial. That is, in the baseline condition, reinforcement was arranged as in a fixed-interval (FI) 5-sec schedule.

In the conditions in which reinforcement was contingent upon temporal patterns of responding, the required response-rate function increased (or decreased) at a specified constant rate. Let $f(t)$ represent the required response pattern. That is, let $f(t)$ represent the response rate that we wish to establish at t sec after the beginning of a trial. Then, the schedule can be summarized by the first derivative of $f(t)$, which by definition is the rate of change of $f(t)$ at t . The derivative of $f(t)$, $f'(t)$, was either -1 or $+1$ key pecks per second per second. Reinforcement was contingent not only upon the rate of change of response rate, but also upon the absolute response rates at various times after the beginning of a trial. Table 1 shows the required response-rate functions, $f(t)$, for each condition, and gives the corresponding number of key pecks within 1-sec subintervals of a 5-sec trial. That is, a continuous function, $f(t)$, was approximated by a discrete function defined over five 1-sec subintervals.

Table 1
Experimental Conditions

							Required Responses per Second, f(t), as a Function of Time, t(sec), since the Beginning of a 5-Sec Trial							
Condition Number	Number of Days	Criterion			Trials per Day	Depri- vation Level (%)	Response- Rate Function, f(t)	f'(t) ¹	Required Responses per Second in Five, 1-Sec Sub- intervals with Upper Bound Equal to t					
		Bird 1	Bird 2	Bird 3					1	2	3	4	5	
1	165 ^a	1	1	1	300	80	t-1	1	0	1	2	3	4	
2	103	—	—	—	100	90 ^a	uncon- trolled	uncon- trolled	uncontrolled					
3	22	2	2	2	300	90	t-1	1	0	1	2	3	4	
4	29	7	2	7	300	85	5-t	-1	4	3	2	1	0	

¹ $f'(t)$ is the first derivative of $f(t)$.

^aThe criterion was self-adjusting during the early part of this condition. The tabled parameters were in effect for more than 100 days for all three birds.

^bIt proved impossible to maintain subjects at an 80% level of deprivation in this condition. The variations in body weight in the present experiment did not have any measurable effect in the sense that Condition 3, a replication of Condition 1 except for differences in body weight, maintained patterns of behavior that were indistinguishable from those maintained in Condition 1.

Goodness-of-Fit

Reinforcement was contingent upon a measure of the overall extent to which a response-rate function produced by a subject in a 5-sec trial conformed to the corresponding required response-rate function shown in Table 1. A subject was not required to generate a response-rate function that perfectly matched the desired function. The subject was required to approximate the desired function only within a tolerance determined by an adjustable "criterion".

The sum of the squared deviations. The overall extent to which the obtained response-rate function on a trial deviated from the required response-rate function was measured by a sum of squared deviations. After each 1-sec subinterval, the computer calculated the difference between the number of key pecks calculated from $f(t)$ in Table 1 and the number of pecks actually observed in that subinterval. The computer then squared this difference, or deviation. At the end of a 5-sec trial, the computer obtained the sum of the five squared deviations, one for each subinterval, to obtain a measure indicating how closely a subject's overall performance on that trial conformed to the required response-rate function. This measure can be expressed symbolically as

$$D = \sum_{i=1}^5 (f_i - o_i)^2,$$

where o_i is the observed frequency of key pecks in the i^{th} 1-sec subinterval and f_i is the corresponding required frequency from Table 1. The larger the value of D on a trial, the poorer the match between required and obtained temporal patterns of responding on that trial. A value of D equal to zero on a trial corresponds to a perfect match between required and obtained response-rate functions on that trial.

The criterion. The criterion was an integer, $C > 0$, such that a reinforcer was presented at the end of a 5-sec trial if, and only if, $D \leq C$ on that trial: if $C < D$, no reinforcer was presented on that trial. The criterion enabled the reinforcement of behavioral patterns not perfectly conforming to the required pattern. The value of C at the end of each condition is shown in Table 1.

During the initial phase of Condition 1, the first 20 trials of each session were warm-up trials, in the sense that on these trials, C was set to a larger value than on subsequent trials. A higher criterion enabled a greater variety of behavioral patterns to be reinforced, and tended to ensure that a subject would be rewarded near the beginning of a session. It was found that a subject might stop responding altogether if it did not experience a reinforcing event early in a session. The warm-up trials were unnecessary later in Condition 1 and therefore were discarded. The larger criteria used in Conditions 3 and 4 made the use of warm-up trials unnecessary in those conditions. No data reported here are from sessions during which warm-up trials were presented.

The criterion was set to a large value at the beginning of each condition, and was reduced over days only as it was judged that a subject's patterns were more closely approximating the desired pattern. This judgement was based exclusively on the number of trials in a session on which $D \leq C$, i.e., on the number of trials on which a reinforcer was delivered. The value of C typically was lowered following a day on which more than 30 reinforcements were obtained and raised following a day on which fewer than 20 reinforcements were obtained.

Intertrial Interval

If a trial ended with the presentation of a reinforcer, a 2.5-sec blackout preceded the illumination of the key at the beginning of the next trial. If a trial did not end with reinforcement, a 5.0-sec blackout preceded the beginning of the next trial. Reinforcement consisted of 2.5-sec access to mixed grain. Thus, a constant 5.0-sec interval separated the end of each trial from the beginning of the next.

Other Arrangements

The left key was illuminated by white light and the houselight was on except during reinforcement and between trials. Also, both the keylight and the houselight were turned off for 0.1 sec after every key peck. Pecks during this blackout were not counted, as they terminated interresponse times so short that they conform to different laws than do longer ones (Shimp, 1973). The 50-min sessions were conducted six or seven days per

week. Pretraining before Condition 1 consisted of exposure for six days to an FI 5-sec schedule.

RESULTS

Cumulative-Response Curves

Figures 1, 2, and 3 show average cumulative-response curves for steady state behavior maintained in each condition. That is, the total number of responses within each 0.25-sec subinterval over the last N days of a condition was divided by the total number of trials over those same days. The value of N was 2 for Conditions 1, 3, and 4, and 6 for Condition 2. Thus, Figures 1, 2, and 3 (and subsequent figures) show data averaged over the last 600 trials for each condition (see Table 1). The average number of responses per 0.25-sec subinterval was then plotted against the duration of the subinterval. Figures 1, 2, and 3 show that all three subjects in both Condition 1 and its replication, Condition 3, produced very similar mean cumulative-response curves.

Furthermore, the mean cumulative-response curves shared quantitative properties in common with the corresponding required curves. For example, consider the total output of key pecks. The total number of responses produced by a subject perfectly conforming to the required function in Conditions 1 and 3 was 10 (see Table 1). This value is represented in Figures 1, 2, and 3 by a horizontal dashed line. The mean total number of responses actually ranged from about 8.5 to about 10.5. Thus, the total output of key pecks varied within the narrow interval 9.5 ± 1.0 , where the desired total output was 10.0. This narrow interval within which the total output varied contrasts sharply with the wide limits produced by the FI schedule in Condition 2. Figures 1, 2, and 3 show that the FI schedule produced approximately the same, significantly more, and significantly fewer, key pecks than did Conditions 1 and 3, for Birds 1, 2, and 3, respectively. Figures 1, 2, and 3 show that the FI 5-sec schedule in Condition 2 produced

Fig. 1. Mean cumulative number of responses for Bird 1 per 0.25-sec subinterval, averaged over the last 600 trials of each condition. The horizontal dashed line represents the cumulative number of responses produced by the required response-rate functions in Conditions 1 and 3.

Fig. 2. Mean cumulative number of responses for Bird 2 per 0.25-sec subinterval, averaged over the last 600 trials of each condition. The horizontal dashed line represents the cumulative number of responses produced by the required response-rate functions in Conditions 1 and 3.

mean cumulative-response curves that had the same positively accelerated shape as those produced by much longer FI schedules. These same figures show that the mean cumulative-response curve was negatively accelerated in

Condition 4, where $f'(t)$ was -1 : in this condition, the mean response rate was greatest at times most remote from reinforcement and was least at times closest to reinforcement. In summary, the mean cumulative-response

Table 2
Mean Time (seconds) to the n^{th} Response

Condition Number	Bird Number	First	Second	Third	Fourth	Fifth	Sixth	Seventh	Eighth	Ninth	Tenth
		REQUIRED ¹ :									
		1.5	2.3	2.7	3.3	3.5	3.7	4.2	4.4	4.6	4.8
		CONTINUOUS THEORETICAL ² :									
		1.4	2.0	2.5	2.8	3.2	3.5	3.7	4.0	4.2	4.5
1	1	1.6	2.1	2.5	3.0	3.3	3.6	3.8	4.0	4.2	4.4
	2	1.5	2.0	2.5	3.0	3.4	3.7	3.9	4.1	4.3	4.4
	3	1.7	2.1	2.5	2.9	3.3	3.6	3.9	4.1	4.4	4.5
2	1	2.4	2.7	3.0	3.2	3.4	3.6	3.8	4.0	4.1	4.3
	2	0.5	1.0	1.4	1.8	2.1	2.5	2.9	3.3	3.6	3.9
	3	3.1	3.5	4.1	4.6	5.0	5.2	(- - - - - ³)			

¹These entries are those for a hypothetical subject distributing the required number of responses in a subinterval equally throughout that subinterval.

²These entries are those for a hypothetical subject distributing its responses according to the continuous function of which the "required function" in Table 1 is a discrete approximation.

³This subject failed to emit seven or more responses sufficiently often to make these entries reliable.

Fig. 3. Mean cumulative number of responses for Bird 3 per 0.25-sec subinterval, averaged over the last 600 trials of each condition. The horizontal dashed line represents the cumulative number of responses produced by the required response-rate functions in Conditions 1 and 3.

curves suggest that differential reinforcement of temporal patterns maintained behavioral patterns that in some ways approximated the required patterns and also reduced intersubject variability, in comparison with an FI schedule. Other differences, such as those in terms of reinforcement density and response-reinforcer contiguity, also could have contributed to the differences in data between Conditions 1 and 2. However, the differential reinforcement of behavioral patterns presumably accounted for most of the results.

Table 2 provides a second way in which one may compare the extent to which the delivery of a reinforcer after a temporal pattern appears to have maintained temporal patterns similar to those preceding reinforcement. Table 2 shows the mean time to the first, second, third, . . . , ninth, and tenth responses in a trial. Table 2 shows that the mean times for all three subjects in Condition 1 closely approximated those that would be produced by a hypothetical subject that distributed its re-

sponses according either to the continuous function $f(t) = t - 1$ or to the discrete approximation of that function shown in Table 1. Table 2 also shows the much greater intersubject variability that occurred in Condition 2 with the FI contingency than in Condition 1.

Interresponse-time Distributions

The same cumulative-response curve can represent very different kinds of behaviors. On the one hand, a mean cumulative-response curve may accurately represent the behavior within individual intervals. On the other hand, Schneider (1969) suggested that a positively accelerated mean curve is obtained by averaging over many intervals, during each one of which behavior is described by a two-state process: an interval begins with a near-zero response rate, which is followed by a "break point" after which responses occur at what appears in a cumulative record to be a steady rate. Interresponse-time distributions

provide a means by which it can be seen whether the scallops in Figures 1, 2, and 3 are attributable to averaging over many intervals each having a two-state, break-run pattern, but having different "break points".

The distributions presented below are averages over trials following reinforcement and trials following nonreinforcement. This averaging is justified by there having been no important differences between the distributions based on the two classes of trials. Most trials in fact followed a nonreinforced trial. The percentages of reinforced trials in Condition 1 were 16, 3, and 7, for Birds 1, 2, and 3, respectively. The corresponding percentages in Condition 3 were 23, 20, and 13. In Condition 4, they were 21, 32, and 5. In Condition 2, the FI condition, 100% of the trials were of course reinforced. The 3% of reinforced trials earned by Bird 2 in Condition 1 was the lowest in the experiment. Three per cent corresponded approximately to 11 reinforcements per hour, since there were 300 trials per session and each trial lasted 10 sec. Eleven reinforcements per hour may be slightly less than that required to ensure that a subject allocates all of its time to reinforced patterns of behavior (Shimp, 1974). This low reinforcement density therefore may partly explain why the behavior of

Bird 2 in Condition 1 conformed somewhat less precisely to that required by the present reinforcement contingency.

Figures 4, 5, and 6 show the relative-frequency distributions of interresponse times terminated by the second, third, fourth, . . . , tenth response. The distribution of latencies shorter than 2 sec in duration terminated by the first response also is shown. As is the case with the curves in Figures 1, 2, and 3, the relative-frequency distributions of interresponse times were very similar across Condition 1 and its replication, Condition 3. In order to conserve space, Figures 4, 5, and 6 only show data from Condition 1. If behavior approximated $f(t)$ in Condition 1, Figures 4, 5, and 6 should reveal a progressive increase in short interresponse times throughout the interval. If, on the other hand, the scallops in Figures 1, 2, and 3 for Conditions 1 and 3 are artifacts of averaging over many different break-run patterns, Figures 4, 5, and 6 should reveal an unsystematic change, or no change at all, in the interresponse-time distributions as a function of the ordinal response number. Figure 4 shows that the interresponse-time distributions for Bird 1 were unimodal and that the mode progressively shifted toward shorter interresponse times as the response

Fig. 4. The relative-frequency distribution of interresponse times for Bird 1 for the first, second, third, . . . , tenth response. The numbers to the side of the ordinal-response-number axis are the frequencies of responses in the corresponding distributions. Points are averages over the last 600 trials of Condition 1.

number increased. The interresponse-time distributions for Birds 2 and 3 were bimodal, or perhaps trimodal. These distributions show a progressively greater preponderance of shorter interresponse times as a function of the response number. In summary, Figures 4, 5, and 6 show that response rate did in fact increase within individual 5-sec trials: the scallops in Figures 1, 2, and 3 for Conditions 1 and 3 are not artifacts of averaging over intervals. At the same time, the schedule of reinforcement did not specify *how* response rate had to increase, and different subjects responded faster in different ways: in the case of Bird 1, the mode of a unimodal interresponse-time distribution shifted toward shorter interresponse times, and in the case of Birds 2 and 3, the percentage of interresponse times belonging to the shortest of two or three subdistributions increased throughout the interval.

Figures 7, 8, and 9 are the equivalents of Figures 4, 5, and 6 for Condition 2, in which reinforcements were arranged according to an FI 5-sec schedule. Figure 7 shows for Bird 1 that there was no clear change in the interresponse-time distribution after the first or second response. That is, Bird 1 responded at an approximately constant rate throughout most of the interval, and the scallop for the

FI schedule in Figure 1 is therefore an artifact of averaging over different break points. Figure 8 shows for Bird 2 that the interresponse-time distributions were very irregular and changed unsystematically for the first four or five pecks. Thereafter, response rate was essentially stable, with no systematic increase. Bird 3 began to respond much later within the 5-sec trials than did Birds 1 and 2. Bird 3 produced interresponse-time distributions with modes in about the same places as in Figure 6, but Figure 9 fails to show, after the second response, any systematic growth in the shortest subdistribution at the expense of the two longer ones. In short, Figures 7, 8, and 9 indicate that after the first few responses, response rate either was approximately constant or varied unsystematically in the FI 5-sec schedule. In summary, the interresponse-time distributions in Figures 4 through 9 show that response rate tended to increase systematically in those conditions in which reinforced temporal patterns closely approximated a required, constantly increasing, response-rate function. On the other hand, the interresponse-time distributions for the FI condition, in which reinforced temporal patterns were not experimentally controlled, failed to reveal any systematic increase within individual intervals.

Fig. 5. The relative-frequency distribution of interresponse times for Bird 2 for the first, second, third, . . . , tenth response. The numbers to the side of the ordinal-response-number axis are the frequencies of responses in the corresponding distributions. Points are averages over the last 600 trials of Condition 1.

Fig. 6. The relative-frequency distribution of interresponse times for Bird 3 for the first, second, third, . . . , tenth response. The numbers to the side of the ordinal-response-number axis are the frequencies of responses in the corresponding distributions. Points are averages over the last 600 trials of Condition 1.

Interresponse-time distributions for Condition 4 are not presented here, as they were not steady state: Condition 4 was conducted only long enough to see if the present schedule with $f'(t)$ equal to -1 could establish a negatively accelerated, mean cumulative-response curve.

Fig. 7. The relative-frequency distribution of interresponse times for Bird 1 for the first, second, third, . . . , tenth response. The numbers to the side of the ordinal-response-number axis are the frequencies of responses in the corresponding distributions. Points are averages over the last 600 trials of Condition 2.

Fig. 8. The relative-frequency distribution of interresponse times for Bird 2 for the first, second, third, . . . , tenth response. The numbers to the side of the ordinal-response-number axis are the frequencies of responses in the corresponding distributions. Points are averages over the last 600 trials of Condition 2.

It was not conducted for a time sufficient to determine whether or not one can establish a clearly decreasing response rate within individual intervals.

Sequential Statistics

The conclusions drawn from the interresponse-time distributions are generally sup-

Fig. 9. The relative-frequency distribution of interresponse times for Bird 3 for the first, second, third, . . . , tenth response. The numbers to the side of the ordinal-response-number axis are the frequencies of responses in the corresponding distributions. Points are averages over the last 600 trials of Condition 2.

ported by Table 3, which shows a second way of determining whether response rate increased within individual trials when reinforcement was contingent upon such an increase. Table 3 shows the probability that, within a given trial, an interresponse time was shorter than the preceding interresponse time.

Table 3

The probability that the n^{th} interresponse time was shorter than the $(n-1)^{\text{st}}$ interresponse time.

<i>n</i>	<i>Bird 1</i> <i>Condition</i>		<i>Bird 2</i> <i>Condition</i>		<i>Bird 3</i> <i>Condition</i>		<i>Average</i> <i>Condition</i>	
	1	2	1	2	1	2	1	2
1	0.98	0.97	0.85	0.48	0.99	1.00	0.94	0.82
2	0.64	0.62	0.50	0.49	0.77	0.35	0.64	0.49
3	0.75	0.42	0.56	0.58	0.66	0.45	0.66	0.48
4	0.70	0.50	0.60	0.56	0.52	0.38	0.61	0.48
5	0.70	0.51	0.59	0.55	0.53	0.45	0.61	0.50
Average	0.75	0.60	0.62	0.53	0.69	0.53	0.69	0.55

The first row in Table 3 shows the probability that the first interresponse time was shorter than the latency between the onset of a trial and the first response. The second row shows the probability that the second interresponse time was shorter than the first, and so on. The probability that the n^{th} interresponse time was shorter than the $(n-1)^{\text{st}}$ interresponse time simply equals the percentage of all pairs of the n^{th} and $(n-1)^{\text{st}}$ interresponse times in which the n^{th} interresponse time was shorter than the $(n-1)^{\text{st}}$. Probabilities for successive interresponse times in the last 1-sec subinterval are omitted from Table 3. Response rate was so high in the last 1-sec subinterval that no reliable increase in response rate in that subinterval could be anticipated. From Table 3, one may compare the probabilities from Conditions 1 and 2. If the reinforcement of temporal patterns was effective, the probabilities should be higher in Condition 1 than in Condition 2. Indeed, according to Schneider's (1969) analysis of FI behavior, subsequent to the latency in Condition 2, the probabilities should not importantly deviate from 0.50. Table 3 shows that indeed the probabilities tended to be higher in Condition 1 than in Condition 2, and that after the latency, there was no reliable tendency over the three subjects for the average response rate to change in Condition 2, the FI condition.

DISCUSSION

The delivery of a reinforcer in classical reinforcement schedules is contingent upon the occurrence or nonoccurrence of very brief behavioral episodes, *e.g.*, key pecks and lever presses. A reinforcer traditionally is temporally contiguous with the response producing it. In the present experiment, the delivery of a reinforcer was contingent upon the occurrence or nonoccurrence of specific patterns of key pecking. That is, reinforcement depended on particular temporal distributions of key pecks. This reinforcement schedule successfully established and maintained temporal patterns of behavior that approximated the required patterns. The general extent to which explicit reinforcement of temporal patterns can successfully strengthen those patterns is of course a question only further research can answer. The generality of the present results presumably depends on experimental parameters such as the trial duration, subinterval duration, the value of the criterion, the formula for the index of goodness-of-fit, and the required function that behavior is to match. An exploration of the effects of these schedule parameters could elucidate the extent to which behavioral patterns, *i.e.*, sequences of response topographies, can be established and maintained by reinforcement contingencies in which reinforcement is not temporally contiguous with all of the behavior producing it.

Any account of the present results must allow for the rather precise control obtained here over behavior that was not temporally contiguous with reinforcement. One traditional account appeals to a hypothetical delay-of-reinforcement gradient. According to this account, responses having the same topography as that which is temporally contiguous with reinforcement, let us say key pecks or lever presses, are differentially strengthened at different times via different delays of reinforcement. This notion has been proposed by Catania (1971) as "an alternative to the analysis of reinforcement schedules in terms of the differential reinforcement of interresponse times". Let us examine this proposal in the light of the present experiment. Catania assumed that a reinforcer delivered after a temporal pattern of key pecks independently strengthens each peck in the sequence to an extent depending inversely on the time sepa-

rating that peck from the reinforcer. The present writers have no quarrel with this first assumption. But Catania assumed secondly that the various increments in strengths contribute independently to a subsequent strength of a single operant, key pecking. He assumed thirdly that this strength is measured by the subsequent mean rate of key pecking. These two assumptions are incompatible with the present results, from all experiments of which we are aware that have used interresponse-time reinforcement contingencies, and indeed, in general from experiments in which reinforcement is not temporally contiguous with all the behavior producing it. How is this so? The difficulty is that for these schedules, assumptions two and three attribute to mean response rate a property it does not have. As we have seen, schedules in which reinforcement is not temporally contiguous with all the behavior that produces it tend to establish and maintain temporal patterns of behavior that resemble those required for reinforcement. Mean response rate then is a byproduct of these patterns. The patterns required by the reinforcement contingency are the behavioral units in such schedules, and mean response rate is derived from the patterns: mean response rate does not reveal the controlling relationships and does not measure the strength of a key-pecking or lever-pressing operant (see Shimp, 1974; Shimp and Hawkes, 1974).

As noted in the Introduction, an operant traditionally is defined as that behavior on which reinforcement is contingent and which is temporally contiguous with reinforcement. In the present experiment, an operant cannot have both of these properties. The present results suggest that the unessential property is that of response-reinforcer temporal contiguity: behavior resembling that required by the contingency was maintained, yet only a small fraction of that behavior was contiguous with reinforcement. This conclusion agrees with results of other experiments, such as experiments with interresponse-time reinforcement contingencies, in which reinforcement is not contiguous with all of the behavior producing it. Response-reinforcer temporal contiguity does not seem to be essential to the definition of the behavioral unit, or operant. At a general level, this conclusion is compatible with other recent suggestions indicating a superfluity of response-reinforcer contiguity (Baum, 1973;

Bloomfield, 1972). However, the details of the present interpretation are to be sharply contrasted with those of Baum's (1973) position. He advocated the abandonment of response-reinforcer contiguity in favor of an analysis in terms of the "molar correlation" between response and reinforcer. His analysis, like that of Catania (1971) described above, presupposes that mean response rate is a valid measure of reinforcement value. As we have seen, if the unit of behavior is a temporal pattern of pecking or lever pressing, as in the present experiment and in experiments with interresponse-time reinforcement contingencies, then mean response rate tells one very little of value: total output of key pecks in a fixed amount of time determines mean response rate, but the pattern of responding determines the total output in a fixed amount of time, and the reinforcement contingency determines the pattern of responding. Baum's position, like Catania's fails to accommodate these facts. In short, mean response rate obscures the controlling relations when the behavioral unit is a behavioral pattern, as it is very likely to be, whenever reinforcement is not contiguous with all the behavior that produces it.

We shall briefly consider one final topic; the ontogeny of behavioral patterns that function as integrated units, or operants. We wish simply to suggest the following possibility. It has been shown that a pigeon to some extent can remember stimulus-response associations for a few seconds (Shimp and Moffitt, 1974). In particular, a pigeon can remember for a few seconds which of two responses was associated with which of three stimuli. Results such as these encourage one to consider the possibility that a pigeon's short-term memory for its recent behavior may constitute a portion of the functional stimulus when a reinforcer is delivered. Repeated presentations of a reinforcer after the same remembered pattern of free-operant behavior might establish that pattern as an organized unit in long-term memory (Shimp and Moffitt, 1974). Such an interpretation predicts a correlation between short-term memory capacity and limitations on the duration and complexity of behavioral patterns that can be established as operants.

REFERENCES

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

- Bloomfield, T. M. Reinforcement schedules: contingency or contiguity? In R. M. Gilbert and J. R. Millenson (Eds.), *Reinforcement*. New York: Academic Press, 1972. Pp. 165-208.
- Catania, A. C. Reinforcement schedules: the role of responses preceding the one that produces the reinforcer. *Journal of the Experimental Analysis of Behavior*, 1971, 15, 271-287.
- Dews, P. B. The effect of multiple S^A periods on responding on a fixed-interval schedule: V. Effect of periods of complete darkness and of occasional omissions of food presentations. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 573-578.
- Dews, P. B. Studies on responding under fixed-interval schedules of reinforcement: the effects on the pattern of responding of change in requirements at reinforcement. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 191-199.
- Ferster, C. B. and Skinner, B. F. *Schedules of reinforcement*. New York: Appleton-Century-Crofts, 1957.
- Schneider, B. A. A two-state analysis of fixed-interval responding in the pigeon. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 677-687.
- Shimp, C. P. The concurrent reinforcement of two interresponse times: the relative frequency of an interresponse time equals its relative harmonic mean. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 403-411.
- Shimp, C. P. Synthetic variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 311-330.
- Shimp, C. P. Time allocation and response rate. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 491-499.
- Shimp, C. P. and Hawkes, L. Time allocation, matching, and contrast. *Journal of the Experimental Analysis of Behavior*, 1974, 22, 1-10.
- Shimp, C. P. and Moffitt, M. Short-term memory in the pigeon: stimulus-response associations. *Journal of the Experimental Analysis of Behavior*, 1974, 22, 507-512.
- Skinner, B. F. *The behavior of organisms*. New York: Appleton-Century-Crofts, 1938.

Received 8 November 1973.

(Final Acceptance 22 July 1974.)