JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1981, 36, 303-315 NUMBER 3 (NOVEMBER)

THE LOCAL ORGANIZATION OF BEHAVIOR: DISCRIMINATION OF AND MEMORY FOR SIMPLE BEHAVIORAL PATTERNS

CHARLES P. SHIMP

UNIVERSITY OF UTAH

A procedure was developed to enable nonverbal organisms to report what they remember of the temporal organization of their recent behavior. A baseline behavior with known temporal structure was established by a concurrent variable-interval variable-interval schedule for two temporal patterns of behavior (two different classes of reinforced interresponse times). The five pigeon subjects emitted these two temporal patterns on a center key and were occasionally given a short-term memory probe for their most-recently-emitted pattern. The probes consisted of symbolic delayed matching-to-sample tests, in which a response on a green side key was reinforced if the most recent pattern belonged to the shorter reinforced class, and a response to a red side key was reinforced if the most recent pattern belonged to the longer reinforced class. All subjects could report with over ninety percent accuracy what their most recently emitted behavioral pattern was when a retention interval separating the pattern from the memory probe was only .1 seconds. The retention interval was then manipulated, and it was found that recall for a pattern was frequently above chance after a delay of as much as eight seconds. Thus, pigeons can remember their most recent interresponse time not only right after it is emitted, but for several seconds thereafter. In other conditions, the patterns themselves were manipulated. It was found that as the patterns became more similar, discrimination became poorer. These results agree with the view that reinforcement tends to organize and integrate the local structure of behavior to the extent to which that structure is remembered.

Key words: short-term memory, delayed matching to sample, concurrent schedules, variable-interval schedules, interresponse times, pigeons

Morse (1966) proposed that there is a shaping property of reinforcement which is reflected in the way reinforcement establishes local temporal patterns in behavior. More recently, it has been suggested that the behavior that is shaped is the behavior an organism remembers having recently emitted when a reinforcer is delivered (Shimp, 1975, 1976, 1978, 1979). A pattern might get shaped as an integrated unit if a subject remembered it when a reinforcer was delivered. This idea has been applied to the analysis of interresponse times, interchangeover times, and more elaborate sequences of various kinds of patterns (Shimp, 1975, 1978, 1979).

An experiment by Reynolds (1966) has appeared to be an anomaly from this point of view. In Reynolds's experiment, a pigeon pecked a red key twice. The second peck changed the color to blue for 30 sec. If the two pecks in red were spaced at least 18 sec apart, so that there was at least an 18-sec interresponse time, the reinforcement schedule in effect during blue was a variable-interval schedule for food. If the interresponse time in red was less than 18 sec the schedule in the subsequent period of blue was extinction. A differential-reinforcement-of-low-rate (DLR) 18-sec schedule was therefore in effect in red for just one interresponse time. Reynolds found that his pigeons had very poor DRL performance in red in the sense that there was little or no sign that the contingency controlled the local patterning of behavior in red. And, what is puzzling from the shaping point of view described above, a pigeon was able to report what interresponse times it was producing in red, in the sense that the subsequent response rate in blue was higher after longer interresponse times in red than after shorter interresponse times in red. Reynolds concluded

---

This research was supported in part by NIMH grant 16928. I would like to thank Stephen J. Hutton for his outstanding work in helping to conduct this research. A part of this research was described in a different form in a talk delivered at the Psychonomic Society Twentieth Annual Meeting, Phoenix, Arizona, November, 1979. Reprints may be obtained from Charles P. Shimp, Department of Psychology, University of Utah, Salt Lake City, Utah 84112.

that the subjects "discriminated the duration of the interresponse times" in red, yet failed to produce interresponse times appropriate to the DRL contingency. It was as though the subject knew what reinforcement depended on, i.e., a particular temporal pattern of behavior, yet did not use that information: the local reinforcement contingency did not shape the local patterning of behavior even though a subject could indicate by its appropriate response rate in blue, that it remembered what interresponse time it had just produced. But this should not happen if reinforcement shapes the local patterning of behavior to the extent to which a subject remembers it.

It seems possible to resolve this anomaly. A first step is to note that the temporal pattern of responding in the presence of red was perfectly confounded with the duration of the red stimulus: the interresponse time in red determined how long red was on. It is therefore entirely possible that a pigeon discriminated not its own interresponse time but the duration of red. Indeed, pigeons are sufficiently oriented toward visual stimuli so that it seems altogether likely that, when confronted with a multidimensional discrimination task that can be solved on the basis either of color or on the basis of the temporal organization of recent behavior, control would be exerted by the former, i.e., color. One could therefore reasonably dispute Reynolds's contention that pigeons were shown to have "discriminated the duration of the interresponse times." They may have discriminated only the duration of red. The shaping hypothesis described above therefore may remain intact; one may still maintain that reinforcement shapes the behavior a subject remembers having recently emitted. There was no shaping of the temporal patterning of behavior in Reynolds's experiment because conditions were arranged so that when a reinforcer was delivered, a subject did not remember its recent temporal patterning.

If this argument potentially resolves the anomaly, it still leaves open the question Reynolds asked: What, if anything, can a pigeon remember about the temporal patterning of its recent behavior, and how is that memory related to the behavior it produces? The experiment described below is designed to provide a method to answer this question and to provide some preliminary results.

The method required a baseline behavior. This was established by a shaping procedure with known properties (Hawkes & Shimp, 1974; Shimp, 1968, 1970). Two simple behavioral patterns of key pecks on a center key were intermittently reinforced according to a variable-interval schedule. Once a reinforcer was arranged, it was assigned randomly to one of the two patterns, and then that pattern had to occur and the reinforcer had to be delivered before the variable-interval schedule began to run again. The patterns in this case were just interresponse times—a shorter one and a longer one. These arrangements produced a clearly bimodal interresponse-time distribution. A pigeon in this situation may be thought of as producing a stream of behavior consisting of a succession of two integrated patterns; some are the shorter interresponse time and the others are the longer interresponse time.

It has been known for a long time that such a reinforcement contingency produces this kind of behavior. What is not known is why the contingency does this. The shaping hypothesis suggests that it is because when a reinforcer is delivered, a pigeon remembers the most recent interresponse time, so that the reinforcer operates on that pattern as a whole. Can it be shown that a pigeon in this context remembers its most recent interresponse time? To answer the question, occasional symbolic matching-to-sample probes were delivered in which a choice was reinforced provided it was to a key corresponding to the preceding interresponse time. The details of this procedure are described next.

METHOD

Subjects

Five White Carneaux pigeons were maintained at 80% of their free-feeding weights, plus or minus 10 g. The subjects had served previously in a variety of experiments on short-term memory for visual stimuli.

Apparatus

Five, three-key Lehigh Valley Electronics pigeon chambers were interfaced to a Digital Equipment Corporation PDP 12/30 that arranged all experimental contingencies and recorded the data. A minimal force of roughly .15N was required to operate the keys. White noise and ventilator fans helped to mask extraneous sounds.

Procedure

The procedure consisted of two parts, a shaping contingency to establish and maintain two behavioral patterns, and a short-term memory contingency to probe memory for the most recently emitted of these two patterns. These contingencies are described separately.

Shaping contingency. This contingency was modeled after earlier experiments in which two temporal patterns, a shorter and a longer, were concurrently reinforced (e.g., Hawkes & Shimp, 1974; Shimp, 1968). During this contingency the houselight was on and the center key appeared white. This contingency had two parts: a variable-interval schedule to arrange a distribution of minimum interreinforcement intervals and a pattern-selection procedure to assign each of the reinforcements arranged by the variable-interval schedule to a particular one of the reinforced patterns. These two features of the shaping contingency are described separately.

Variable-interval schedule: A single variable-interval schedule arranged reinforcement with probability .05 every 1.0 sec so that the mean interreinforcer interval, as arranged by the schedule, was 20 sec. When reinforcement was arranged, the variable-interval timer stopped and reinforcement remained available until a key peck terminated the selected pattern. No additional reinforcement could be arranged until a previously arranged one was delivered.

Pattern-selection procedure: The durations of the reinforced pairs of patterns are displayed in Table 1 for each condition. Each reinforcement arranged by the variable-interval schedule was assigned to one of the two reinforced patterns. The assignments were random in the sense that each pattern was equally likely to have a particular reinforcement assigned to it, and one assignment did not depend on the previous one. After initial pretraining established the desired baseline (see below), all reinforcers arranged in this manner were replaced by short-term memory probes, so that this contingency became the means of arranging probes rather than of arranging reinforcers.

Short-term memory procedure. Symbolic delayed matching-to-sample probes were designed to measure a pigeon's memory for its most recent pattern, i.e., the shorter or longer reinforced class of interresponse times. In general, this procedure worked as follows. A key peck initiated a retention interval provided that the variable-interval schedule had timed out and that the key peck terminated an interresponse time falling in the class chosen by the pattern selection procedure (see above). The red and green side keys appeared after the termination of the retention interval. A bird was required to peck the green key if the previous interresponse time belonged to the shorter class and to peck the red key if the previous interresponse time belonged to the longer class. The details of this procedure were as follows.

Retention interval: A short-term memory probe began with a retention interval during which the experimental chamber was dark and during which a peck on any key had no programmed consequence. Table 1 shows for every condition the duration of the retention interval. It ranged from .1 to 8 sec.

Test phase: When the retention interval timed out, red and green appeared on the side keys, with the position of a particular color varying randomly from one test to another. A peck to the green key turned off the side keys and initiated reinforcement provided that the most recent pattern had been a shorter one, and a peck to the red key initiated reinforcement provided that the pattern had been a longer one. After reinforcement, the center key again appeared white and the houselight was turned back on. A peck to the incorrect color initiated a correction procedure.

Correction procedure: A choice of the incorrect key was followed by a 5-sec correction interval, during which the keylights were off but the houselight was on. After the 5 sec elapsed, the side keys appeared again, with the same assignment of colors to positions. This correction procedure continued to recycle until a correct response produced reinforcement and ended the test. This procedure, coupled with tests that were delivered equally often after shorter and longer patterns, ensured that reinforcement was delivered equally often after pecks to a particular color and position.

Experimental conditions. Two classes of experimental parameters were varied: the retention interval and the durations of the two

Table 1

Experimental Conditions

Condition Number	*Reinforced Classes of Interresponse Times (Lower and Upper Bounds in sec)*		*Retention Interval (sec)*	*Number of Days*	*Mean Relative Frequency of the Shorter Class and Standard Error*	
	Shorter	*Longer*				
1	1.50, 2.00	4.50, 7.00	.1	20	.68	.01
2	1.50, 2.00	4.50, 7.00	.5	20	.70	.01
3	1.50, 2.00	4.50, 7.00	1.5	15	.66	.02
4	1.50, 2.00	4.50, 7.00	3.0	15	.64	.01
5	1.50, 2.00	4.50, 7.00	.1	10	.67	.01
6	1.50, 2.00	10.00, 12.50	.1	15	.87	.01
7	1.50, 2.00	10.00, 12.50	2.0	10	.87	.01
8	1.50, 2.00	10.00, 12.50	4.0	15	.90	.00
9	1.50, 2.00	10.00, 12.50	8.0	15	.90	.01
10	1.50, 2.00	3.50, 5.00	.1	15	.70	.01
11	1.50, 2.00	2.50, 3.50	.1	15	.64	.01
12	1.50, 1.75	2.00, 2.50	.1	15	.64	.01
13	1.50, 1.60	2.25, 2.50	.1	15	.62	.02
14[1]	1.50, 2.00	4.50, 7.00	.1	15	.76	.01
15	1.50, 2.00	4.50, 7.00	6.0	15	.77	.00

[1] Two conditions between 13 and 14 were deleted due to a programming error.

temporal patterns. Table 1 summarizes how these values of the experimental parameters were varied over conditions.

Miscellaneous details. Reinforcement consisted of 1.5-sec access to the same fixed grain that provided the subjects' daily diet. During reinforcement, all lights in the chamber were off except for one directly over the food hopper. The duration of each session was one hour, and sessions were conducted at the same time six days a week. Experimental conditions typically lasted a standard 15 days, except in a few cases described in Table 1. There was no programmed visual or auditory response feedback at any time.

Pretraining. When the experiment began, it was unclear that a suitable method could be developed for the study of memory for the temporal organization of behavior, and it was by no means clear what kind of pretraining should be administered. While the pretraining that was given did indeed work, there is no way to determine at present whether the procedure was particularly efficient, and it certainly was not executed according to any specific predetermined plan. Note also that there is no guarantee this procedure will work for subjects with different histories. A flexible approach, with a sharp eye out for the development of position biases, should help. Not every parameter change is described in the following summary of pretraining.

The first attempt at pretraining consisted simply of placing the subjects on the schedule described above, with a .1-sec retention interval, a higher density of reinforcement arranged by a variable-interval 10-sec schedule, and with half of the reinforcements allocated for interresponse times and half for probes. Not surprisingly, this was an overly optimistic attempt and after about 10 days there was a beginning of the two-interresponse-time baseline, but severe position biases were emerging in the short-term memory probes. It was then decided to establish first the baseline performance of behavioral patterns and only subsequently to introduce the short-term memory probes. Therefore, the percentage of reinforcers replaced by short-term memory probes was reduced from 50 to 5, so that 95% of the reinforcers were for behavioral patterns. Over several days, the percentage of reinforcements allocated to the shorter pattern was varied from 20 to 80 and an appropriate baseline performance was quickly established. The percentage of reinforcers replaced by probes was then changed back to 50. Attention was then focused for several days on establishing control by the correction procedure in the probes. At this stage, the correction interval was varied over a range of 5 to 20 sec. Next, an attempt was made to establish the interresponse time preceding a probe as the discriminative stimulus for behavior during the probe. To do this, the probe contingency was, from the subject's point of view, simplified: the relative frequency of reinforcement for longer patterns was changed from .50 to .90 and the probability that red would appear on the right was changed from .50 to .90. This implied that it was appropriate for a pigeon to emit many long patterns and usually to peck the right key in the probes. Altogether, 81% (9 times 9) of the probes were after long patterns and required a pigeon to peck the right key. Then, the percentage of reinforcers replaced by probes was changed to 100%; all reinforcers were those delivered during probes—none was delivered directly for the termination of a pattern. This arrangement soon resulted in an appropriate baseline behavior consisting largely of the longer pattern, and of appropriate behavior in the probes, consisting of choosing the right, red key. Then, the location of red was allowed to vary randomly between left and right, and after several days, performance appropriately tracked this change, so that after a longer pattern, a pigeon chose the red key, whatever its location. The next step was to have the subject choose the green key after a shorter pattern. This was achieved by arranging a situation complementary to that just described: most of the probes were arranged after shorter patterns. After about a week of such training, the contingency was reversed back so that most probes followed longer patterns. Two more such reversals followed, after progressively fewer days in any one condition. Lastly, the relative frequency of probes following shorter patterns was fixed at .50, after which Condition 1 began. This entire period of pretraining took 62 sessions.

Summary of procedure. A pigeon responded on a center key in such a way that its behavior stream consisted largely of shorter and longer patterns. Once in a while, following a randomly selected shorter or longer pattern, the computer that controlled the experiment delivered a symbolic matching-to-sample test. There was a retention interval consisting of a blackout after the pattern and before the red and green side keys were turned on.

A peck to a green key was reinforced if the preceding pattern belonged to the shorter category and a peck to a red key was reinforced if the preceding pattern belonged to the longer category.

RESULTS

Independence of Baseline Behavior from Short-Term Memory Probes

The relevance of the present results for a general understanding of local behavioral patterning would seem to depend on the extent to which the local patterning obtained here was representative of that which would be expected in the absence of the probes; to the extent that the occasional short-term memory probes changed the local structure of behavior, any conclusions based on these results might not generalize to situations where an organism is not confronted with probes. Thus, a preliminary to the presentation of the major results is the determination of whether the baseline established here was representative of behavior maintained by concurrent variable-interval variable-interval schedules for two behavioral patterns. We ask, that is, if the occasional short-term-memory probes changed the baseline behavior. The kind of behavior one would expect in the absence of the probes is known in some detail (Shimp, 1968, 1969b, 1970, 1973). This literature suggests that we use two criteria to answer our question. First, we can determine whether the shapes of the interresponse-time distributions appeared importantly changed by the presence of the probes and second, we can ask if preference between the two patterns appeared changed.

First, the top row in Figure 1 shows relative-frequency distributions of interresponse times for two birds. These two examples, from Condition 10, reveal the qualitative features of the corresponding distributions in previous experiments where two classes of interresponse times were concurrently reinforced. In these earlier experiments, the two reinforced classes generally have been sufficiently different so that the interresponse-time distributions were bimodal. In particular, the difference between the lower bounds of the two reinforced classes has been less than two sec in only very rare cases. In Condition 10, this difference was 1.5 sec and the top row of Figure 1 shows in this case the present curves were appropriately bimodal. In other respects as well, the shapes of these two curves match up with expectations built up from the corresponding previous experiments (Gibbon, 1977; Shimp, 1978), so that by the first criterion, the baseline behavior

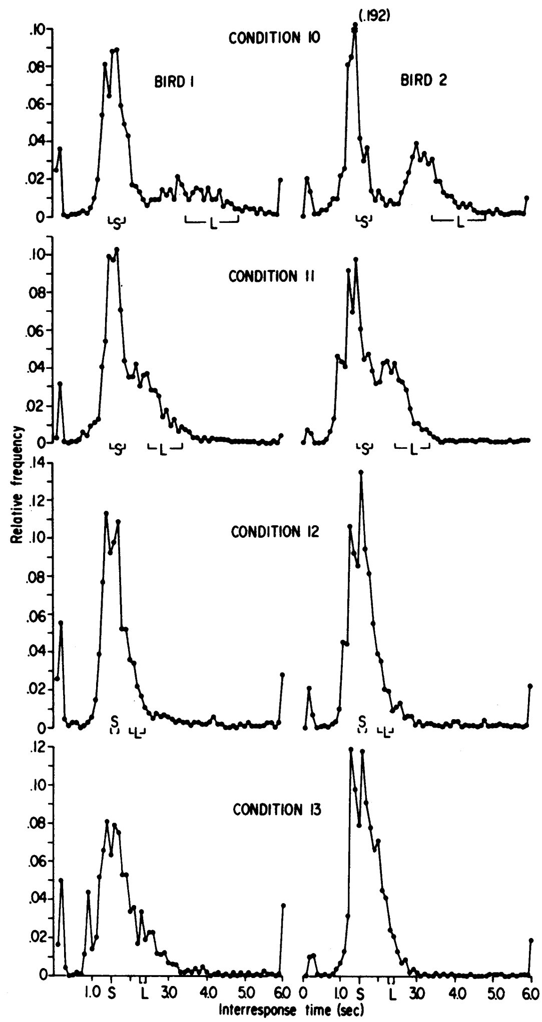

Fig. 1. Relative frequencies of occurrence of interresponse times in 0.1-sec bins, on the last day of each of four conditions. The two reinforced classes in each condition are marked by brackets labeled "S" and "L" for "shorter" and "longer." The left and right columns show distributions for Birds 1 and 2, respectively. In this series of four conditions, the time between the two reinforced classes was varied.

was unaffected by the short-term memory probes.

Table 1 shows a second way in which the baseline behavior resembled that which would have been obtained without the short-term memory probes. The relative frequency of the shorter pattern is the total number of interresponse times in the shorter reinforced class divided by the total number of interresponse times in either reinforced class. Table 1 shows these relative frequencies averaged over the five last days of a condition and the five subjects. The principal effect is that preference for the shorter class increased as that class became relatively shorter. This effect is the outcome one would expect if the short-term memory probes had not changed the baseline behavior (Hawkes & Shimp, 1974; Shimp, 1969b). Notice, incidentally, that the duration of the retention interval in the probes seems not to have had any systematic effect on preference between the two reinforced behavioral patterns. On the basis of the two criteria, then, one may tentatively conclude that the baseline behavior in the present experiment either was not affected at all by the introduction of the probes or, if it was, it was affected in some relatively subtle way unrelated to present purposes. Therefore, we tentatively may conclude that the results of the probes are applicable to cases where there are no probes; one would say in ordinary English that the probes tell us something about what a pigeon knows about the temporal patterning of its recent behavior (Shimp, 1976).

The two features of the baseline behavior described so far indicate ways in which reinforcement shaped the local organization of behavior. Therefore, by the hypothesis motivating the present experiment, one should be able to show that a subject could remember features of the organization of its recent behavior. The next step is, therefore, to ask if the results support this hypothesis. Specifically, can a pigeon remember its most recent interresponse time?

Effects on Recall of the Retention Interval

Figure 2 shows two functions representing the effects on the probability of correct recall in the short-term memory probes of changes in the retention interval. These results are based exclusively on the first choice in the probes; correction responses are not included. One function shows recall for the previous pattern for conditions when the shorter and longer patterns were relatively closer together, i.e., more similar: the shorter and longer patterns were interresponse times between 1.5 and 2.0 sec and between 4.5 and 7.0 sec, respectively. The other function is for conditions when the two patterns were farther apart, 1.5 to 2.0 sec and 10.0 to 12.5 sec. Both curves are generally decreasing functions of the retention intervals.

When the two patterns were less similar, the

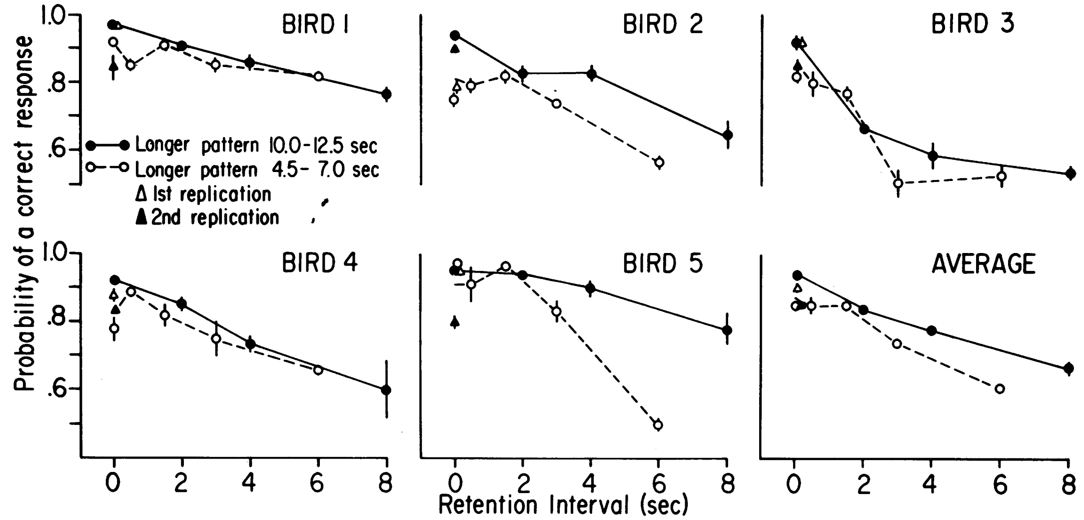

Fig. 2. The probability of a correct response in the symbolic-delayed-matching-to-sample probes as a function of the retention interval, i.e., the blackout between the termination of a shorter or longer interresponse time and the beginning of a probe. Each data point is an average over the last five days of a condition. The vertical lines on the panels for individual subjects indicate plus and minus one standard error where the variability is over the last five days of a condition. The vertical lines on the panel for averaged data indicate plus and minus one standard error where the variability is over the five-day averages for the five subjects. Where no line is visible, its length is less than the diameter of the symbol. The dotted line has been drawn through the average of the three conditions with the retention interval equal to 0.1 sec.

chance level of .50 was at least two standard errors below the obtained mean probability of recall for all birds except Bird 4, even for a retention interval of 8 sec; even 8 sec after the most recent pattern, several of the subjects could remember whether that pattern was a shorter or a longer one. When the two patterns were more similar, all but one subject's performance was still two standard errors above chance at a retention interval of 3 sec. The two replications of Condition 1 produced performances that adequately approximated that in Condition 1. Figure 2 suggests that the probability of recall of the most recent pattern is higher when the two alternative patterns are less similar. (This possibility is later examined in greater detail in Figure 4.)

Figure 3 describes the behavior over the same conditions as Figure 2, but in a different way. Figure 3 displays the probability of a hit plotted against the probability of a false alarm (Green & Swets, 1966). A hit may be defined as a peck at a green key after a short pattern, and a false alarm may be defined as a peck at a green key after a long pattern. The probability of a hit is the relative frequency of pecks at a green, as opposed to a red, key when the preceding pattern was short. The probability of a false alarm is the relative frequency of pecks at a green key after a long pattern. In a plot such as that in Figure 3, one can see two separate effects at work. A reduction in memory for the previous pattern is seen as movement of data points down an isobias line, represented as the minor diagonal in each panel of the figure. The isobias line represents the locus of points corresponding to all levels of accuracy with no bias for either of the two stimuli corresponding to the two to-be-remembered patterns. Figure 3 reveals little systematic bias for either pattern. No systematic bias emerges at all in the panel for the averaged data. Figure 3 indicates that the curves in Figure 2 are not attributable primarily, if at all, to changes in a bias for one stimulus over the other.

Effects on Performance of the Similarity Between Patterns

The similarity between the shorter and longer reinforced patterns, defined in terms of the time between the two corresponding lower bounds, was increased over Conditions 10 to 13; that is, the two patterns became closer together and more similar. The results from these conditions and from 1, 5, 6, and 14 appear in Figures 4 and 5. The retention interval was .1 sec in all eight of these conditions.

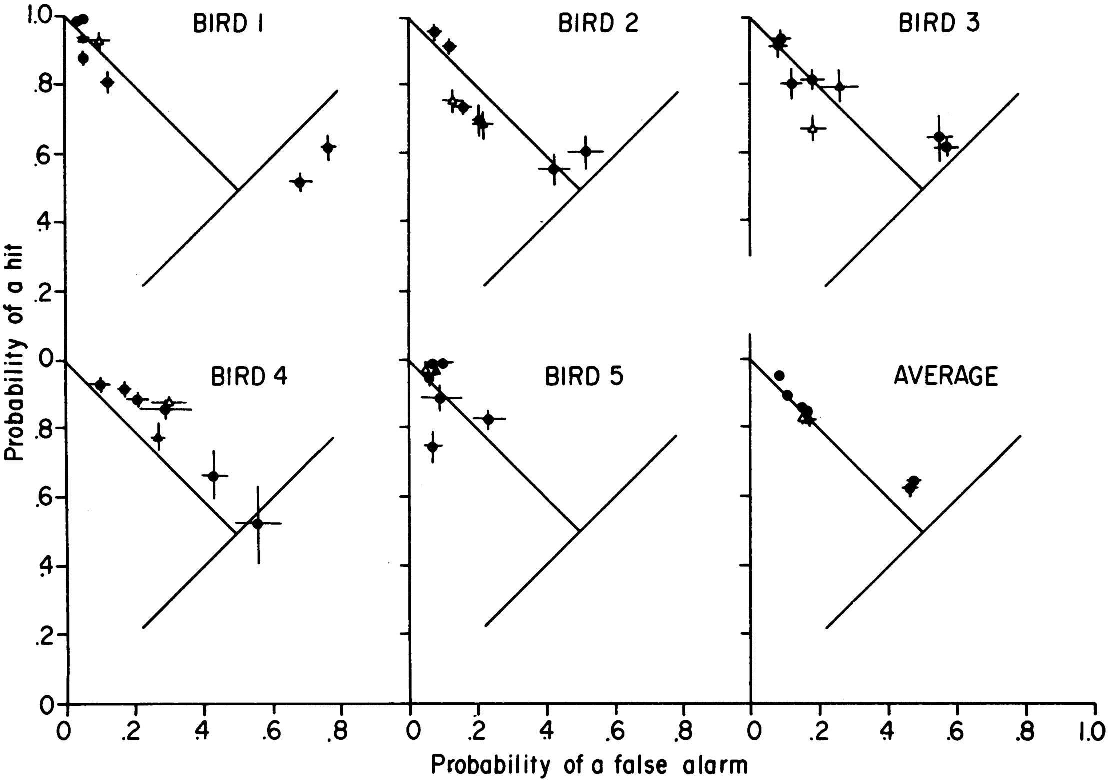

Fig. 3. The probability of a hit as a function of a false alarm, in the symbolic-delayed-to-sample probes. Each point is an average over the last five days of a condition. The vertical and horizontal lines indicate plus and minus one standard error, in the same manner as in Figure 2.

Figure 4 shows the probability of a correct response as a function of the similarity between the shorter and longer reinforced patterns. Figure 4 shows that performance decreased as the patterns became more similar, with most of the decrease taking place after the two lower bounds were 2 sec or less apart. The probability of a correct response in Figure 4 is correlated with the extent to which different patterns are discriminable in Figure 1; the bottom three rows of Figure 1 suggest that poorer performance on the test phase of a trial was correlated with a baseline behavior that itself revealed scarcely any discrimination between the two reinforced patterns.

Figure 5 shows the probability of a hit plotted against the probability of a false alarm, for the same conditions as shown in Figure 4. Figure 5 suggests that the changes in the probability of a correct response shown in Figure 4 are not in general attributable to changes in bias for a stimulus corresponding to one or the other behavioral pattern: with only a few notable exceptions in the panels for Birds 1 and 5, changes in the similarity between reinforced patterns affected a subject's discrimination between the two patterns, not its bias for a stimulus associated with a pattern.

DISCUSSION

Both the present results and those obtained by Reynolds (1966) are consistent with the idea that reinforcement shapes the local temporal organization of behavior to the extent to which a subject remembers that organization when a reinforcer is delivered. That is, to the extent to which temporal properties of recent behavior are represented in working memory, those properties are reinforceable. In Reynolds's experiment, it is reasonable to assume the temporal structure of recent behavior was not represented in working memory when a reinforcer was delivered, so reinforcement did not shape the local patterning of behavior. The present experiment can be interpreted as showing that the local organization was represented in working memory when a

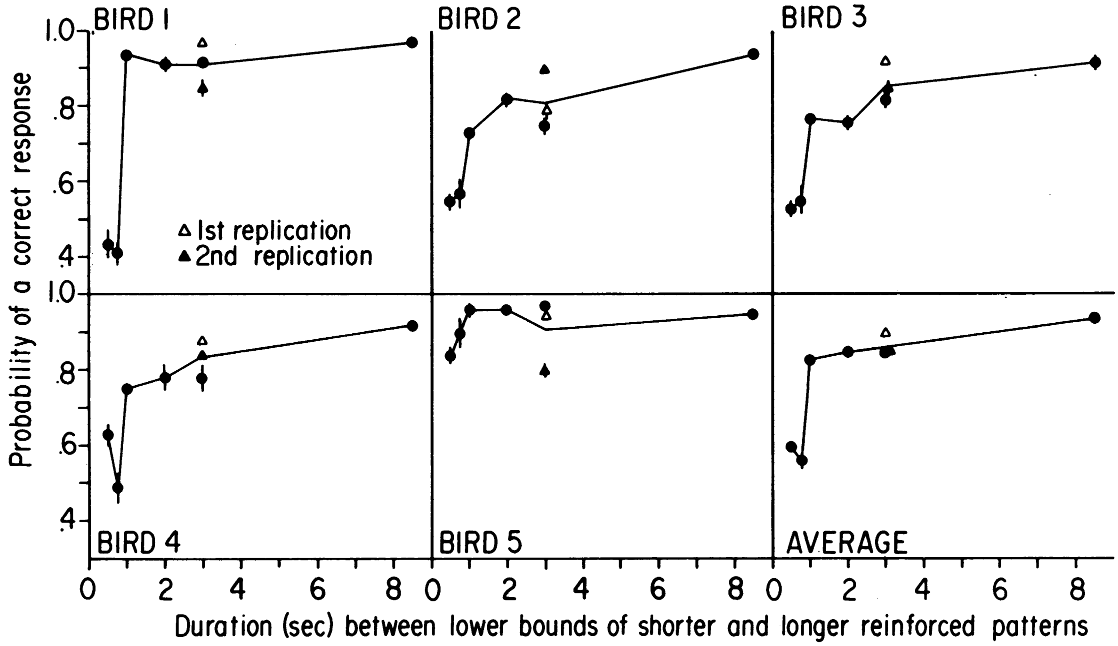

Fig. 4. The probability of a correct response in the symbolic-delayed-matching-to-sample probes as a function of the time between the lower bounds of the two classesof reinforced interresponse times. Each point is an average over the last five days of a condition and the vertical lines indicate plus and minus one standard error, as in Figure 2. The first and second replications are indicated by unfilled and filled triangles, respectively. The line has been drawn through the average of the three conditions at a value of 3 sec.

reinforcer was delivered and that reinforcement did shape the local structure of behavior.

A growing body of data from several different contexts is amenable to interpretation by this same idea; namely, that if a temporal pattern can be shaped by reinforcers that follow it, then that pattern was in working memory when the reinforcers were delivered. Pliskoff and Tierney (1979) conducted an experiment closely resembling Reynolds's. Unlike Reynolds, they obtained good evidence of temporal patterning and consistent with the present idea about the necessary conditions for such patterns, the subjects appeared able to report what pattern they had just emitted. Nelson (1974) conducted an experiment remarkably similar to the present one. He developed a procedure which, like the present one, required a pigeon to emit either shorter or longer interresponse times on a center key for access to symbolic-matching-to-sample probes on two side keys. The shorter interresponse time was varied over conditions from 1 to 2 sec to 5 to 7 sec, while the longer always was 5 to 7 sec. The color of the side key to which a response was reinforced depended on the duration of the preceding interresponse time on the center key. Nelson's data agree with the corresponding data in the present experiment: he found both that the temporal organization of behavior on the center key was shaped by the interresponse-time contingency and that the probe results indicated that a pigeon could report whether it had just emitted a shorter or longer pattern. The present results go beyond Nelson's in describing the time interval that can be interposed between a pattern and a pigeon's successful report of which pattern occurred.

Other interresponse-time data are compatible with the thesis we are considering. If one pattern is shorter than another, then it should be remembered better and a reinforcer delivered at the moment of its termination should more effectively shape it and integrate it. Other things being equal, such as biological preparedness for different kinds of patterns, shorter patterns should evidence greater accuracy of control by reinforcement than longer ones. A rather extensive body of data, including the top two rows of the present Figure 1, testify that this is indeed the case (Gibbon, 1977; Shimp, 1968, 1978, 1979).

There is yet another possible interpretation

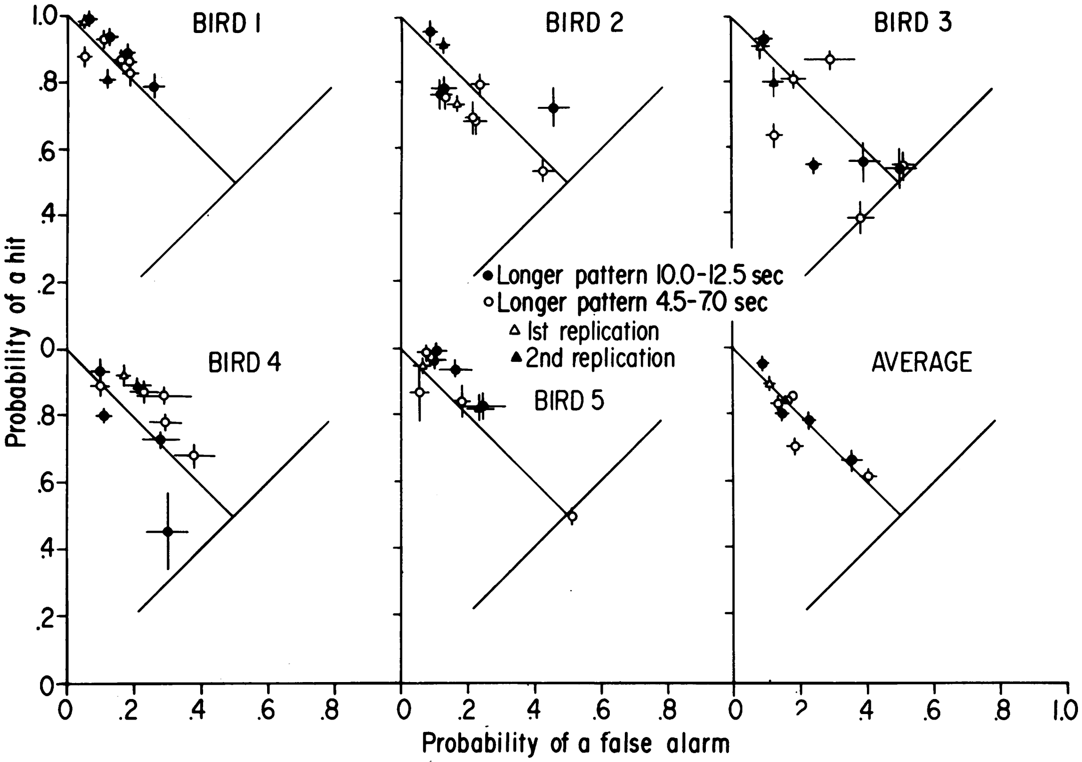

Fig. 5. The probability of a hit as a function of a false alarm, in the symbolic-delayed-matching-to-sample probes, for the same conditions as in Figure 4. Each point is an average over the last five days of a condition. The vertical and horizontal lines indicate plus and minus one standard error, as in Figure 2. The first and second replications are indicated by unfilled and filled triangles, respectively.

of the present thesis that patterns are shaped to the extent to which they are remembered when a reinforcer is delivered. The thesis would seem to imply for the present experiment that patterns would be shaped more effectively in conditions with shorter retention intervals than in those with longer ones, because in the latter conditions, a subject could not remember so well which pattern it had just emitted. However, the method of the experiment makes the evaluation of this interpretation difficult. First, retention intervals were introduced here only after a pattern was established, so that it is possible to examine the effects of retention intervals only on the maintenance of already established patterns. Second, and perhaps more important, the termination of a pattern was followed immediately by what presumably was the conditioned reinforcing stimulus of the retention interval. This stimulus could shape the pattern by an amount sufficient to maintain it even in cases where the retention interval was long enough so that a subject could no longer remember, at its termination, which pattern it had most recently emitted. In short, the procedure does not appear to equate the effects of the length of a pattern itself, which affected both recall of a pattern and the effectiveness with which it was shaped, and the length of a retention interval following a pattern. Table 2 shows some evidence on how the retention interval affected the shaping of patterns. The measure in Table 2 of the effectiveness of shaping is simply the percentage of all interresponse times that fell in the reinforced categories. If this number decreases for a fixed pair of reinforced categories, as the retention interval increases, the effectiveness of shaping can be said to decrease and to be correlated with poorer recall of the pattern. The top part of the Table shows no evidence to support the interpretation that shaping effectiveness should decrease as the temporal interval between a pattern and the delivery of a reinforcer increases. The bottom part of the Table, on the other hand, hints at such an effect for at least four of the five subjects and also for

Table 2

A measure of the effectiveness of the shaping procedure: percentage of all interresponse times that fell in the reinforced categories on the last day of conditions over which the retention interval was varied.[1]

Condition Number	*Reinforced Interresponse Times*		*Retention Interval (sec)*	*Bird 1*	*Bird 2*	*Bird 3*	*Bird 4*	*Bird 5*	*Average*
	Short	*Long*							
5	1.5-2.0	4.5- 7.0	.1	.53	.47	.57	.33	.53	.49
14	1.5-2.0	4.5- 7.0	.1	.33	.38	.29	.37	.40	.35
2	1.5-2.0	4.5- 7.0	.5	.50	.45	.28	.23	.29	.35
3	1.5-2.0	4.5- 7.0	1.5	.43	.36	.39	.21	.35	.35
4	1.5-2.0	4.5- 7.0	3.0[2]	.48	.55	.22	.27	.49	.40
15	1.5-2.0	4.5- 7.0	6.0	.38	.54	.46	.32	.49	.44
6	1.5-2.0	10.0-12.5	.1	.43	.39	.30	.32	.45	.38
7	1.5-2.0	10.0-12.5	2.0	.41	.38	.29	.31	.42	.36
8	1.5-2.0	10.0-12.5	4.0	.45	.36	.27	.29	.39	.35
9	1.5-2.0	10.0-12.5	8.0	.38	.36	.23	.22	.29	.30

[1] The data-recording format for Condition 1 was incorrect and prevented a suitable analysis.

[2] Data reported are for the next-to-last day of this condition; the data for the last day were accidentally not stored permanently on tape.

the group average. Thus, the superimposition of a retention interval between a behavioral pattern and a subsequent reinforcer does not have as great an effect on shaping effectiveness as does the length of the pattern itself.

The applicability of the idea we are examining is not limited to interresponse times. Ziriax and Silberberg (1978) employed a procedure much like Nelson's, but with center-key key-peck duration rather than interresponse-time duration serving as the cue for the correct side-key response. These experimenters obtained the same result as did Nelson: key-peck durations were shaped and subjects were able to report the durations of their responses. These results significantly extend the generality of the correlation between successful shaping of temporal patterns and an ability to recall those patterns, because the shorter and longer patterns in this experiment were only 0 to 20 msec and 60 to 180 msec, respectively. In the opposite direction from simple response durations, Wasserman, Deich, and Cox (in press) have advanced similar ideas for a context in which different sequences of choices are reinforced. However, memory limitations presumably restrict the level of complexity and duration of the temporal patterns that can be shaped. Further work will be required before it is clear how the present idea can be extended to the very large scale patterns studied by Marr (1979) and Zeiler (1979).

The short-term memory probes in the present experiment had no detectable effect on the local structure of the baseline behavior; the baseline behavior resembled that obtained in other experiments where reinforcement shaped different local temporal patterns of behavior (Shimp, 1968, 1969b, 1970). The present results are therefore in good agreement with the conjecture that in these earlier experiments reinforcement succeeded in establishing new behavioral units, classes of interresponse times, because a preceding interresponse time was represented in working memory when a reinforcer was delivered. Perhaps this is a model of what happens, in more complex ways, in operant conditioning in general: reinforcement may chunk remembered behavioral patterns into integrated units. If so, it might be these units, not single key pecks, that enter into associative relations with other events (Shimp, 1975, 1976, 1978, 1979).

Whereas the view presented here is intended to be reasonably general, it is not intended to be universally applicable. Several important qualifications need to be described. First, the present view of the relation between short-term memory, reinforcement, and the local structure of behavior is one that emphasizes ontogeny, not phylogeny. Other approaches to the temporal organization of behavior quite appropriately focus on the role of natural selection (Fentress, 1976; McFarland, 1976). It is to be hoped that future research will suggest how to integrate what at present are these two different approaches to the analysis of the structure of behavior. At present, neither approach can assimilate the phenomena handled by the other. Second, it bears emphasizing that the hypothesis examined here is designed to deal with the chunking or integration of temporal patterns, not simply with the emergence of any type of temporal pattern. The hypothesis asserts that if a behavioral pattern is integrated, then a subject can, or at one time could, remember it. It does not assert that all behavioral patterns can be remembered. Looked at this way, the hypothesis becomes a means for determining which patterns are, or can be, integrated. Third, the hypothesis deals with a phenomenon that might be labeled as "automatization" of behavior and should not be confused with other hypotheses about "automatic processes" in recent research on human perception (e.g., Shiffrin & Schneider, 1977). One may speak of an integrated class of interresponse times as "automatic" in the sense that one speaks of a decision made at the beginning of an interresponse time. It may be said that a subject chooses how long to wait before the next key peck; the duration of an interresponse time is apparently determined at its beginning rather than at its end (Gibbon, 1977; Shimp, 1969a, 1975, 1976, 1978). This implies that after commencing an interresponse time, a subject does not thereafter continue to decide whether or not to respond. This kind of automated behavioral pattern is sometimes called an integrated motor program (Stelmach, 1978): once started, the pattern runs off automatically until it is completed.

In one context with human subjects, it has already been noted that the present hypothesis does not seem to explain all the data. This is the context of work on serial pattern learning that interrelates remembered and emitted behavioral patterns (e.g., Restle & Brown, 1970). But the theoretical alternatives that have been proposed in the context of recent research with human subjects are more complex than the present hypothesis, frequently invoke linguistic variables, and do not yet seem required by the results obtained to date with infrahumans on short-term memory and the effects of reinforcement on the local organization of behavior.

REFERENCES

Fentress, J. C. Dynamic boundaries of patterned behavior: Interaction and self-organization. In P. P. G. Bateson & R. A. Hinde (Eds.), *Growing points in ethology*. Cambridge: Cambridge University Press, 1976.

Gibbon, J. Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, 1977, **84**, 279-325.

Green, D. M., & Swets, J. A. *Signal detection theory and psychophysics*. New York: Wiley, 1966.

Hawkes, L., & Shimp, C. P. Choice between response rates. *Journal of the Experimental Analysis of Behavior*, 1974, **21**, 109-115.

Marr, M. J. Second-order schedules and the generation of unitary response sequences. In M. D. Zeiler & P. Harzem (Eds.), *Advances in analysis of behavior* (Vol. 1), *Reinforcement and the organization of behaviour*. Chichester, Eng.: Wiley, 1979.

McFarland, D. J. Form and function in the temporal organization of behavior. In P. P. G. Bateson & R. A. Hinde (Eds.), *Growing points in ethology*. Cambridge: Cambridge University Press, 1976.

Morse, W. H. Intermittent reinforcement. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application*. New York: Appleton-Century-Crofts, 1966.

Nelson, T. D. *Interresponse time as a stimulus: Discrimination and emission of interresponse time by pigeons*. Unpublished doctoral dissertation, University of Maine, 1974.

Pliskoff, S. S., & Tierney, T. J. On Reynolds' "Discrimination and emission of temporal intervals by pigeons." *Bulletin of the Psychonomic Society*, 1979, **13**, 173-174.

Restle, F., & Brown, E. Organization of serial pattern learning. In G. H. Bower (Ed.), *Psychology of learning and motivation* (Vol. 4). New York: Academic Press, 1970.

Reynolds, G. S. Discrimination and emission of temporal intervals by pigeons. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 65-68.

Shiffrin, R. M., & Schneider, W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 1977, **84**, 127-190.

Shimp, C. P. Magnitude and frequency of reinforcement and frequencies of interresponse times. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 525-535.

Shimp, C. P. Optimal behavior in free-operant experiments. *Psychological Review*, 1969, **76**, 97-112. (a)

Shimp, C. P. The concurrent reinforcement of two interresponse times: The relative frequency of an interresponse time equals its relative harmonic length. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 403-411. (b)

Shimp, C. P. The concurrent reinforcement of two interresponse times: Absolute rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1970, **13**, 1-8.

Shimp, C. P. Sequential dependencies in free-responding. *Journal of the Experimental Analysis of Behavior,* 1973, **19**, 491-497.

Shimp, C. P. Perspectives on the behavioral unit: Choice behavior in animals. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 2). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1975.

Shimp, C. P. Organization in memory and behavior. *Journal of the Experimental Analysis of Behavior,* 1976, **26**, 113-130.

Shimp, C. P. Memory, temporal discrimination, and organization in behavior. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 13). New York: Academic Press, 1978.

Shimp, C. P. The local organization of behaviour: Method and theory. In M. D. Zeiler and P. Harzem (Eds.), *Advances in analysis of behaviour* (Vol. 1). *Reinforcement and the organization of behavior.* Chichester, Eng.: Wiley, 1979.

Stelmach, G. E. *Information processing in motor control and learning.* New York: Academic Press, 1978.

Wasserman, E. A., Deich, J. D., & Cox, K. E. The learning and memory of response sequences. In M. L. Commons, R. Herrnstein, & A. Wagner (Eds.), *Harvard symposium on quantitative analysis of operant behavior* (Vol. 3), in press.

Zeiler, M. D. Output dynamics. In M. D. Zeiler & P. Harzem (Eds.), *Advances in analysis of behavior* (Vol. 1). *Reinforcement and the organization of behaviour.* Chichester, Eng.: Wiley, 1979.

Ziriax, J. M., & Silberberg, A. Discrimination and emission of different key-peck durations in the pigeon. *Journal of Experimental Psychology: Animal Behavior Processes,* 1978, **4**, 1-21.

Received April 28, 1981
Final acceptance June 4, 1981