

An analysis of serial pattern learning by rats

ROLAND SELF and E. A. GAFFAN
University of Reading, Reading, England

When rats receive a sequence of rewards of different magnitudes for traversing a runway, they learn to "track" the sequence, showing anticipation of the forthcoming reward by appropriate running speed. There is disagreement as to whether this behavior depends on rats' encoding and recalling a complete sequence of foregoing hedonic events or just the immediately preceding one. The present experiments showed that rats can remember more hedonic events than the most recent one. In Experiment 1, when exposed concurrently to the sequences 10-1-0 (pellets) and 0-1-10, they were faster on Run 3 of the increasing than of the decreasing sequence, a discrimination which cannot be made on the basis of the preceding (1-pellet) reward alone. Experiment 2 showed that this behavior reflects genuine anticipation of the Run 3 reward, not simultaneous contrast or other simple aftereffects of Runs 1 and 2. It is argued, however, that these results, together with related findings by Capaldi and Verry (1981), show merely that rats can recall a hedonic event other than the most recent one, not that a sequence of such events is fully recalled in order.

Hulse (1978) has reviewed a number of experiments from his laboratory, where food-deprived rats received sequences of food rewards of different magnitudes for successive runs along a runway. For example, one series or trial might consist of five runs, closely spaced in time, for which rewards of 14, 7, 3, 1, and 0 food pellets were given in succession. After undergoing a number of such trials, rats show systematic changes in running speed across the runs of each trial. In the above example, they run more slowly towards the end of the series than at the beginning, with an especially marked drop in speed for the final nonrewarded run. Conversely, increasing reward magnitude within trials results in progressively faster running across a trial (Hulse & Campbell, 1975; Wike & King, 1973).

Evidently, rats are able to anticipate in some way the forthcoming rewards within a series and adjust their running speed appropriately. Hulse (1978) has raised the question of whether such anticipation is based on simple chaining, that is, using the outcome of the immediately preceding run as a cue for reward magnitude on the current run, or on more complex encoding of the magnitude sequence in its entirety. He has argued for the second interpretation, on the basis of the following evidence. First, rats show more accurate tracking (anticipatory adjustment of speed) with monotonic sequences, such as 14-7-3-1-0, than with nonmonotonic ones, such as 14-1-3-7-0 (Hulse & Dorsky, 1977, 1979). The simple chaining notion

provides no obvious rationale for this difference, but the sequence-encoding view suggests that a monotonic series, having a more simply describable structure, should be easier to encode and follow. Second, Hulse and Dorsky (1979) have found that rats previously exposed to monotonic sequences of various lengths show positive transfer (as measured by acquisition of sequence tracking) to a new monotonic series, compared with controls who had been pretrained with random sequences. Finally, Fountain and Hulse (1981) demonstrated a form of sequence extrapolation. Rats who had been trained with a four-element monotonic series, 14-7-3-1, showed better tracking of the extended five-element series 14-7-3-1-0 than did animals who had the 0 element added to a nonmonotonic training series.

All these data, though they are consistent with Hulse's (1978) sequence-encoding view, can also be explained in terms of more elementary associative and discriminative processes, according to Capaldi's sequential theory (Capaldi, 1967, 1971), which has been applied to a wide range of phenomena in instrumental learning. This account (see Capaldi, Verry, & Davidson, 1980) starts from the assumption that animals learn to use the memory of the reward outcome of a preceding trial as a discriminative cue which predicts the outcome, and hence controls the level of responding, on the current trial. Thus, in a 14-7-3-1-0 series, the memory of 14 pellets predicts that 7 will be received for the next run, while the memory of 1 pellet predicts nonreward. The subject therefore runs faster on the second than on the last run of the series, and so on. So far the theory is apparently of the "simple chaining" variety and offers no obvious explanation of why, for example, monotonic series are tracked more accurately than equally invariable but non-

The research was conducted by the first author and formed part of a thesis submitted in partial fulfillment of the requirements for the BA (psychology) at Reading University. Reprint requests may be addressed to E. A. Gaffan, Department of Psychology, Building 3, Earley Gate, University of Reading, Whiteknights, Reading RG6 2AL, England.

monotonic ones as described above. Such effects can, however, be encompassed if generalization between similar reward memories is considered. For example, we may assume that the memory of 14 pellets resembles the memory of 7 more closely than it does that of 1, and so forth. In a monotonic sequence, similar memories will predict similar outcomes and hence control similar response levels; in the current example, 14 and 7 signal 7 and 3, while 3 and 1 signal 1 and 0, so generalization between memories will not result in highly incompatible response tendencies and the discrimination between early and late runs in the sequence, with their appropriate behavior, will be easy to learn. In a nonmonotonic sequence such as 14-3-1-7-0, however, cross-generalization will cause problems because similar memories, such as 14 and 7, may signal consequences such as 3 and 0, which induce very different response tendencies. It is possible, moreover, to devise sequences in which considerations of generalization between memories lead to predictions different from those based on monotonicity and formal structure of the sequence, and the results from studies employing these favor Capaldi's rather than Hulse's predictions (Capaldi & Molina, 1979; Capaldi et al., 1980).

The data on transfer and extrapolation between series (Fountain & Hulse, 1981; Hulse & Dorsky, 1979) are also amenable to the generalization analysis. For example, Fountain and Hulse (1981) report that animals trained on a monotonic 14-7-3-1 series showed better tracking of an added 0 reward than did those trained with "weakly monotonic" 14-5-5-1 or nonmonotonic 14-3-7-1. This could be because extrapolation depends on the response tendency controlled by the memories most similar to that of 1 pellet during training. For the first group, this is 3 pellets, which was always followed by 1 and controlled slow running. For the second group, it is 5 pellets, which was sometimes followed by 5 and sometimes by 1; and for the third group, it is 3, which had been followed by 7. So, clearly, the generalized tendency evoked by the memory of 1 will be for faster running (less efficient extrapolation) in the latter two groups.

In attempting to decide between the alternative theories, the first point to be established is whether behavior on each run is controlled only by the outcome of the immediately preceding run, as assumed by Capaldi et al. (1980), or whether a sequence of at least two events is involved. Only if the latter turns out to be true need we consider whether the sequence undergoes abstract encoding, as proposed by Hulse (1978). The experiments reported here approach that simple question as follows. Rats were trained concurrently on two different series of reward magnitudes, for example, a decreasing series 10-1-0 was given on half the trials and an increasing series 0-1-10 on the remaining half. The critical test concerns be-

havior on the last run of each trial. All such runs are preceded by a 1-pellet reward, so the memory of the preceding run's outcome alone cannot determine whether 0 or 10 pellets will be received for the third run. If the subjects show some discrimination, running more slowly on the third run of a decreasing than of an increasing sequence, their behavior must be under the control of the sequence of rewards, 10-1 or 0-1, received for the first two runs. Exactly what feature of the two-reward sequence may be responsible is discussed later.

A similar technique has been independently devised by Capaldi and Verry (1981), using rather different types of sequence. Their results are discussed in relation to ours below.

EXPERIMENT 1

The main experimental group received the treatment just described. Each trial consisted of three closely spaced runs along a runway, the rewards for the three runs increasing in magnitude—0, 1, and 10 45-mg food pellets—on half the trials and decreasing (10-1-0) on the other half. There were two control groups, C(I) and C(D), which each received only one type of trial, increasing and decreasing, respectively. This was to verify that the simple sequence-tracking reported by Wike and King (1973) and Hulse and Campbell (1975) was shown by our subject population, and to provide a basis for comparison of the tracking, if any, shown by the experimental group.

Method

Subjects. These were 16 experimentally naive male hooded Lister rats obtained from Olac (Southern) Ltd., Bicester, Oxon. They were approximately 9 weeks old at the start of pretraining and were housed in pairs. They were food-deprived from 1 week prior to the experiment and were maintained at 80% of free-feeding weight (with a correction for growth estimated from the weight of nondeprived controls) by postexperimental feeding. Water was constantly available in the home cages and intertrial holding cages.

Apparatus. The runway, of unpainted plywood, had the following internal dimensions: total length, 214 cm; width, 10 cm; height, 23 cm. The first 30 cm and the final 41 cm of the runway were separated off by hand-operated guillotine doors and served, respectively, as startbox and goalbox. The entire apparatus was covered with transparent Plexiglas lids. A timing circuit, which recorded run times to the nearest .01 sec, was started by the raising of the startbox door and stopped by the interruption of an infrared photobeam 2.5 cm beyond the goalbox door. Testing was conducted under normal artificial ceiling lights with windows occluded. Forty-five-milligram Noyes diet pellet rewards were placed in an opaque bowl at the far end of the goalbox. Intertrial waiting periods were spent in individual holding cages with water bottles.

Procedure. On Days 1 and 2, the subjects were placed in the runway in pairs, with all doors open and 10 45-mg pellets distributed along the floor. They were left in the apparatus for 10 min and handled before being returned to the home cage. On Days 3 and 4, each rat underwent two preliminary trials as follows. After the rat was placed in the startbox, both doors were opened and it was allowed to traverse the runway and enter the goalbox, where the bowl contained four pellets; the doors were then closed. The rat

remained in the goalbox until it had eaten all the pellets; it was then placed in a holding cage until the next trial.

The main experiment commenced the day after pretraining and comprised Phase 1 (14 days) and Phase 2 (4 days). The subjects were nonsystematically assigned to an experimental group, E, of eight rats and a control group consisting of two subgroups C(I) and C(D) of four rats each. Both groups were trained during Phase 1, but only Group E proceeded to Phase 2.

Group E rats were given four trials per day, each consisting of three runs with different reward magnitudes, as described below. Groups C(I) and C(D) received two trials a day. For Group C(I), the sequence of rewards during every trial was increasing—0, 1, 10 pellets (I). Group C(D) always had a decreasing sequence, 10, 1, 0 (D). Group E, during Phase 1, had four I trials one day and four D trials the next, alternating by days. Half the subjects started with each type of sequence. Thus, during Phase 1, Group E experienced as many I trials in total as Group C(I) and as many D trials as Group C(D).

At the end of Phase 1, the C groups' behavior was stable and training was discontinued, but Group E's performance suggested the need for a procedural modification (see Results), and they received 4 further days of training (Phase 2) in which I and D sequences were alternated over the four trials of each day, half the subjects commencing with I and half with D.

Throughout, the rats were trained in squads of eight, either E or C. Members of each squad were placed in holding cages and received their trials in rotation, the order of rotation being varied daily. At the start of a trial, the rat was moved from its holding cage to the startbox. The doors were raised 5-10 sec later and lowered when the rat entered the goalbox. If it took more than 30 sec to complete the run, it was placed in the goalbox and a run time of 30 sec was recorded; this never occurred after Day 4. The rat was removed from the goalbox as soon as it had finished eating (or after 10 sec confinement on 0-pellet runs) and replaced in the holding cage while the goal was baited for the second run. Baiting was carried out in such a way as to avoid giving cues to reward magnitude; for example, on 0-pellet runs, the experimenter acted as if depositing pellets in the bowl. Runs 2 and 3 were given in the same fashion, and the next rat's trial then commenced. The interval between successive runs within a trial late in training was 15-20 sec and between successive trials for the same rat, 30-40 min. Supplementary feeding was given in the colony room 2 h after the end of training.

Results

Figure 1 shows the running speed of the control groups during the three runs of the I or D series, pooled over trials within days for the last 4 days of Phase 1. The performance resembles that reported by Wike and King (1973) and Hulse and Campbell (1975) with similar procedures. The C(I) subjects ran slowly on the first 0-pellet run and increased their speed over the next two runs; the C(D) group ran fast initially and then decreased. Thus, they were able to anticipate not only the changing reward magnitude within a trial, but also the absolute magnitude on the first run of the trial; the fact that each trial is preceded by a relatively long intertrial interval is presumably the cue for that discrimination (Haggblom, 1978). A mixed-design analysis of variance on speeds (reciprocal run times), I/D groups (2) \times runs (3) (pooling over days as that factor was nonsignificant in a preliminary analysis), showed significant effects of I/D [$F(1,6) = 11.55$, $p < .05$], runs [$F(2,12) = 5.38$, $p < .05$], and the I/D \times runs interaction [$F(2,12) = 16.00$, $p < .01$].




Figure 1. Experiment 1: Mean running speed of control groups on the three runs of each trial during the last 4 days of Phase 1. C(I) always received the reward magnitude sequence 0-1-10 pellets, whereas C(D) received a 10-1-0 sequence. (Speed is measured from start-door opening to passing goal entrance.)

Figure 2 shows the performance of the experimental group during the last 4 days of Phase 1, this time pooled over days but displaying the variation across trials within the day. The I vs. D comparison is here made within subjects rather than between, as with the controls. Evidently, at this stage, the rats were showing no discrimination between I and D sequences; indeed, on the first trial of each day they showed a tendency to behave in a way appropriate to the opposite sequence, increasing speed between the first two runs of a D sequence and vice versa on I sequences. Over the following trials, a pattern developed which was similar regardless of sequence, only slight and nonsignificant signs of I-D discrimination appearing by the fourth trial. Analysis of variance, I/D (2) \times trials (4) \times runs (3), confirmed significant effects of runs [$F(2,14) = 6.48$, $p < .05$] and the trials \times runs interaction [$F(6,42) = 3.57$, $p < .05$], but no effects or interactions involving the I/D factor were signif-




Figure 2. Experiment 1: Mean running speed of the experimental group on 0-1-10 (I) and 10-1-0 (D) sequences during the last 4 days of Phase 1. Each day's training comprised four I or four D trials, alternating by days within subjects. Speeds are shown separately for the four trials, pooled over 4 days.

icant (that is, no sequence discrimination was shown).

We concluded that the alternation of I and D sequences by days—four trials per day of each type—might be interfering with discrimination. There appeared to be signs that the current day's sequence was being acquired within the day, but the inappropriate behavior on Trial 1 suggests transfer from the previous day and consequently a need for reversal learning. This difficulty might be overcome if the sequences were alternated more frequently, as in Phase 2, in which I and D alternated every trial.

Figure 3 shows that this resulted in more successful discrimination, even though only 4 days' training was given under the new condition (suggesting that "latent" discrimination may have occurred in Phase 1, although masked by the need for reversal learning; cf. Capaldi, 1979). As in Figure 2, the data are pooled over 4 days but segregated by trials within a day. It should be recalled that Trial 1 was an I sequence for half the subjects and a D sequence for the rest, and so on for the other trials, but as a result of pooling over days the points for each trial represent data from every subject on both I and D sequences. It is evident that the rats generally ran more slowly on Runs 2 and 3 of D sequences than on the corresponding runs of I sequences. Analysis of variance indicated significant effects of I/D [$F(1,7)=8.09$, $p < .05$], runs [$F(2,14)=4.83$, $p < .05$], and the I/D \times runs interaction [$F(2,14)=9.10$, $p < .01$]. Simple main effects analyses confirmed that I and D did not differ on the first run of a trial ($F < 1$) but that D speeds were significantly slower on Runs 2 and 3 [$F(1,7)=6.70$ and 10.74 , $ps < .05$].

Comparison of Figures 1 and 3 shows that the I-D discrimination by the experimental group differed from that of the controls in two ways. First, the controls anticipated the reward magnitude on Run 1 but the experimentals did not. Second, the controls showed monotonic increases and decreases in speed over trials,




Figure 3. Experiment 1: Mean speed of the experimental group during Phase 2. I and D trials now alternated within each day. Format is as in Figure 2.

while the experimentals' pattern was more irregular, exhibiting in particular a V-shaped function on D sequences with the slowest speed on Run 2 and not Run 3. Incidentally, a similar V-shaped function was evident on both I and D sequences during most trials of Phase 1 (cf. Figure 2) when discrimination had not yet appeared.

Discussion

The results from Group E in Phase 2, particularly from the critical third run of each trial, suggest that the rats' behavior was being influenced by more than the immediately preceding trial. Similar conclusions have been drawn by Capaldi and Verry (1981) from a series of experiments in which rats were concurrently trained on pairs of sequences such as 0-0-20 and 20-0-0, 1-0-20 and 20-0-0, or 5-0-20 and 20-0-0. In all such concurrent pairs, the outcome (0 or 20 pellets) on Run 3 of each sequence cannot be predicted from the preceding run's outcome, which is always 0 pellets. Capaldi and Verry's subjects also discriminated on Run 3, running faster for 20 than for 0 pellets in all the above examples. Thus, our findings are in agreement with theirs, and extend the result to the case in which Run 2 has a 1- rather than 0-pellet reward.

Inspection of Capaldi and Verry's data reveals other points of similarity to ours. First, no discrimination was evident on Run 1, but the rats responded differentially on Run 2; they ran more slowly when the Run 2 nonreward followed a large Run 1 reward (20) than after a relatively small one (5, 1, or 0). Also, the V pattern of running speed in decreasing sequences was usually apparent, in that rats ran more slowly on the second than on the third run of a 20-0-0 sequence. (Similar phenomena appeared in other experiments reported by Capaldi & Verry, 1981, which employed longer sequences of nonreward, such as 20-0-0-0-0.) Capaldi and Verry noted these effects and described them as "nonrewarded-trial serial position effects"; the present data show that they are not peculiar to nonrewarded runs, as our rats ran most slowly on the 1-pellet run of a 10-1-0 sequence.

This pattern of results suggested to us some alternative accounts of the data which both we and Capaldi and Verry obtained from the concurrent-sequence paradigm. These should be evaluated before we can safely conclude that the rats' behavior on Run 3 reflects memory of and response to the sequence of rewards on Runs 1 and 2. The differential speeds on Run 2 suggest the operation of a simultaneous negative contrast effect (Mackintosh, 1974) whereby an anticipated 1-pellet reward (in our study) or nonreward (Capaldi & Verry, 1981) leads to slower running when it follows a larger Run 1 reward. Since the first of a series of small rewards following a large one would be most subject to negative contrast, this

process might also contribute to the V pattern observed on decreasing sequences in both studies. The presence of negative contrast, while it may explain some details of performance, does not at first sight seem relevant to the main theoretical issues involved here; however, we argue below that it may be so. (It should be mentioned that Capaldi & Verry, 1981, offer a different account of these phenomena, which is dealt with in the General Discussion.)

The V pattern of running speed also suggests the possibility that the concurrent sequences are at first treated by the experimental subjects not as two distinct sequences 10-1-0 and 0-1-10, but as examples of a single pattern in which Run 1 and Run 3 receive partial reinforcement (50% 10 or 0 pellets) while Run 2 always results in 1 pellet. Leventhal, Morrell, Morgan, and Perkins (1959) have shown in a T-maze study that rats prefer a random sequence of 2 or 0 units of food to a constant 1 unit, so in the present case the partial (10 or 0) schedule must be more rewarding than the constant 1 pellet for Run 2. This would give rise to the tendency to run faster on Runs 1 and 3 than on Run 2. That simple account does not suffice to explain why Run 3 is slower on D than on I trials, but we may combine it with the idea (see above) that the value of the 1-pellet reward is affected by negative contrast, being lower on D than on I trials. This difference might carry over to the third run, in the same way that rewarded and nonrewarded trials in a standard random partial reinforcement sequence tend to be followed by faster and slower running, respectively, on the next trial (Tyler, Wortz, & Bitterman, 1953).

The main characteristic of this type of explanation is that it attributes behavior on Runs 2 and 3 of a trial to simple aftereffects of the rewards—10-1 or 0-1—received on the first two runs, plus the anticipation of *either* 10 or 0 pellets unpredictably on Run 3, but it does not assume that rats correctly anticipate which reward magnitude will occur on Run 3. By employing a different kind of control procedure, Experiment 2 was designed to assess whether such correct anticipation actually occurs under the experimental condition.

EXPERIMENT 2

Here the experimental group, E, received the same treatment as did Group E in Phase 2 of the previous experiment. The sequences 10-1-0 and 0-1-10 were alternated at four trials per day throughout 24 days of training. The control group, C, also received four trials per day, and on each trial the rewards received for Runs 1 and 2 were the same as those received for the corresponding runs by an experimental subject. However, the reward for the third run was 0 or 10 pellets on a quasirandom basis, in-

dependently of whether the preceding pair had been 10-1 or 0-1. If, according to the ideas outlined above, performance on Run 3 was determined simply by aftereffects of Runs 1 and 2, the two groups should behave similarly. If, however, Group E rats could learn to anticipate correctly that 10-1 was followed by 0 and 0-1 by 10, they should show superior discrimination on Run 3.

Method

Subjects. The subjects were 12 experimentally naive female rats, aged 15 weeks at the start of training, procured and maintained in the same way as for Experiment 1. They were assigned six each to Groups E and C.

Apparatus. The runway was the same as that used for Experiment 1. Two photobeams were added, 2.5 cm beyond the start-box door and 7.5 cm before the end of the goalbox, that is, just in front of the food bowl. Timers were connected to these and to the original photobeam 2.5 cm beyond the goalbox entry in order to record separately start time (from start door being raised to start photobeam being broken), run time (from start photobeam to goal entrance beam), and goal time (from goal entrance beam to food bowl).

Procedure. Pretraining was similar to that in Experiment 1. During the main experiment, the general procedure was the same as in Phase 2 of Experiment 1, each rat receiving four trials per day, three runs per trial for 24 days. Any subject spending more than 30 sec in the start or run section was moved to the next section, and a 30-sec time was recorded. Rats were trained in rotation in squads of six (three E and three C), the interval between trials being 30-40 min and between runs within a trial 15-20 sec, during the last four blocks.

Group E received alternating I (0-1-10) and D (10-1-0) trials, starting with I or D on alternate days, half the subjects starting with I each day. Group C's trials began with two runs having the same reward sequence as the corresponding pair of a matched E subject, but the third run ended in 0 or 10 pellets quasirandomly in a counterbalanced manner. For example, a day's trials for an E subject might comprise 10-1-0/0-1-10/10-1-0/0-1-10, while a matched C subject received 10-1-10/0-1-0/10-1-0/0-1-10.

Results and Discussion

The separate start, run, and goal times showed broadly consistent patterns of results, so the data presented as speeds here are the reciprocals of the total times (start + run + goal). The times analyzed in Experiment 1 were the equivalent of start + run.

Figure 4 shows the mean speeds of Group E on Runs 1-3 of I and D sequences pooled over 4-day blocks throughout the experiment. Trials within days are not separated as statistical analyses (below) revealed no significant trials effects. Figure 5 presents the data from Group C in the same form, but here I and D correspond not to consistently increasing and decreasing sequences but to sequences commencing with an increasing or decreasing pair, in line with the matched E subject, and terminating quasirandomly with 0 or 10.

Figure 4 shows that Group E acquired a pattern similar to that shown in Phase 2 of Experiment 1; they ran uniformly fast on Run 1, while Runs 2 and 3 were much slower in D than in I sequences. In par-




Figure 4. Experiment 2: Mean speed of the experimental group over six 4-day blocks of training. I and D trials (four in all) alternated within days, and the data are pooled across trials of all ordinal positions in the day. (Speed is measured from start-door opening to reaching food bowl.)




Figure 5. Experiment 2: Mean speed of the control group during training. Format is as in Figure 4. For interpretation of I and D in this group, see text.

ticular, they showed clear discrimination on Run 3, confirming the principal result of Experiment 1 and those of Capaldi and Verry (1981). In contrast, Group C (Figure 5), while clearly manifesting the V pattern referred to under Experiment 1, showed weaker differentiation between I and D trials, and this tended to appear on Run 2 and not Run 3.

Blocks 4-6 (the last 12 days) were selected for statistical analysis as Group E's behavior seemed stable over that period. Separate within-group analyses, I/D (2) × blocks (3) × trials (4) × runs (3), were first carried out for each group. In the case of Group E, only the following effects were reliable: I/D [$F(1,5) = 88.68, p < .01$], runs [$F(2,10) = 71.32, p < .01$], and I/D × runs [$F(2,10) = 16.75, p < .01$]. No other factor approached significance.

The corresponding analysis on Group C showed significant effects for runs [$F(2,10) = 20.55, p < .01$], I/D × blocks [$F(2,10) = 4.41, p < .05$], and I/D × runs [$F(2,10) = 4.80, p < .05$]. Figure 5 suggests that the I/D × blocks interaction arises because the I-D differentiation was less in Block 4 than in the other two blocks. The I/D × runs interaction, along with the corresponding effects on Group E, is examined further in the analyses reported next.

Since these preliminary within-group analyses do not allow direct comparison of the groups, we also carried out a mixed-design analysis, groups (2) × I/D (2) × runs (3). Data were pooled over the

trials and blocks factors, as these had produced no substantial effects in the foregoing analyses.

The relevant means are summarized in Figure 6. This illustrates the fact already mentioned, that Group E discriminated I and D sequences on Runs 2 and 3, while Group C did so, if at all, only on Run 2. The following effects were reliable: I/D [$F(1,10) = 57.50, p < .01$], groups × I/D [$F(1,10) = 26.87, p < .01$], runs [$F(2,20) = 55.79, p < .01$], I/D × runs [$F(2,20) = 17.70, p < .01$], and groups × I/D × runs [$F(2,20) = 11.08, p < .01$].

The triple interaction was examined first by tests on simple interaction effects (groups × I/D) at each of the three runs. At Run 1, there were no significant effects [$F_s(1,10) \leq 2.11$]. At Run 2, there were significant effects of I/D [$F(1,10) = 38.01, p < .01$] and the groups × I/D interaction [$F(1,10) = 7.56, p < .01$]. Tukey's tests showed the interaction to arise from a significant difference between groups ($p < .05$) in the D but not the I condition. However, both groups were significantly slower in the D than in the I condition ($p < .01$ for Group E, $p < .05$ for Group C). Finally, on Run 3, there were again significant effects of I/D [$F(1,10) = 20.81, p < .01$] and the groups × I/D interaction [$F(1,10) = 32.22, p < .01$]. Tukey's tests showed the latter interaction to have two components. As on Run 2, Group E ran more slowly than Group C ($p < .05$) in the D but not the I condition. Also, Group E showed a significant I-D effect ($p < .01$) but Group C did not. To sum up: both groups ran more slowly on Run 2 of D than of I trials, but Group E showed the stronger discrimination. Only Group E showed an I-D discrimination on Run 3.

Interrun comparisons (Newman-Keuls tests) showed that none of the differences between Runs 1, 2, and 3 in the I condition was significant for either group.




Figure 6. Experiment 2: Mean speeds of experimental and control groups during Blocks 4-6, showing that Group E discriminated I from D on both Run 2 and Run 3, while Group C did so on Run 2 but not on Run 3.

However, in the D condition, Group E ran more slowly on both Run 2 and Run 3 than on Run 1. Group C rats were slower on Run 2 than on Run 1, but Run 3 did not differ from either. Thus, as seen in Experiment 1, Group E "tracked" D sequences rather imperfectly, not showing a monotonic decline in speed. On I sequences, there were slight indications of tracking in that Run 3 tended to be faster than Run 2, though not significantly so; ceiling effects presumably played a part. It must be stressed, however, that the behavior of Group E is complex and cannot be simply characterized as "tracking," even on D sequences. This is clear from a consideration of Group C, which also showed "discrimination" on Run 2, even though they were not exposed to consistent monotonic sequences. We may conclude that at least part of the "discrimination" shown by Group E on Run 2 is a simple aftereffect of the Run 1 reward, presumably negative contrast, which is also shown by the control group. (Other factors might also play a part in Group E; see General Discussion.)

No such ambiguities arise over the discrimination on Run 3, which appeared in the experimental group only. Such discrimination must reflect genuine anticipation of the reward magnitude on Run 3, and is not an artifactual consequence of the rewards received for Runs 1 and 2 and associated contrast effects; if it were, it should have been at least partly evident in the control group. We may conclude that rats are capable of discriminating between the memories of reward sequences such as 10-1 and 0-1, despite the fact that they are identical in respect to the second or most recent reward experienced. The model outlined by Capaldi et al. (1980) needs extension to encompass these data.

GENERAL DISCUSSION

Capaldi and Verry (1981, p. 450) drew a strong conclusion from their data which might possibly also be applied to ours: "rats can anticipate a given hedonic event in a series by remembering all prior events in the series and the order in which they occurred." They characterize the ability to remember a series of events, rather than just the most recent, as "multiple hedonic memory." (By contrast, Hulse, 1978, seems to consider that the series is not remembered in "raw" form but is encoded in terms of higher order characteristics such as monotonicity. The data presently under discussion are not relevant to that issue.)

Capaldi and Verry apply the notion of multiple hedonic memory in attempting to explain "Run 2 discrimination" (our terminology) or the "non-rewarded-trial serial position effects" (their terminology), as well as performance on Run 3. Their

account of the former effects is as follows. It is assumed that ordered sequences can be recalled and, moreover, that generalization can operate between the memories of event sequences and of single events that share common elements, just as it does between single-event memories (see Introduction). Experiment 1 of Capaldi and Verry (1981), which employed the sequences 20-0-0 and 0-0-20, may be used as an illustration. The phenomenon to be explained is that the rats learned sooner to anticipate the 0 pellets (run slowly) on Run 2 of the 20-0-0 sequence than on Run 2 of the 0-0-20 sequence. We attributed such effects to simultaneous contrast, but Capaldi and Verry argue as follows. Let (20) signify the memory of 20 pellets, (20-0) the memory of the sequence 20-0, and so on. So, in the present case, there are three memories which predict 0 pellets on the succeeding run, that is, (20) \rightarrow 0, (20-0) \rightarrow 0, and (0) \rightarrow 0, and one which predicts 20 pellets, that is, (0-0) \rightarrow 20. It should be hard to learn to anticipate 0 pellets following (0) and (20-0) because these memories are similar to (0-0), which is followed by a very different consequence, that is, 20 pellets. By the same argument, it should be relatively easy to anticipate 0 pellets following (20) because there should be little generalization between (20) and (0-0). This explains why it should be easier to run slowly on Run 2 of 20-0-0 than on either of the other 0-pellet runs, that is, Run 3 of 20-0-0 or Run 2 of 0-0-20. Analogous accounts are possible for all Capaldi and Verry's data and for our Experiment 1.

Our explanation of Run 2 behavior in terms of simultaneous contrast has the advantage that it can be applied to the occurrence of "Run 2 discrimination" in the control group of Experiment 2, while it seems Capaldi and Verry's cannot. In their terms, the control group must learn the following: (10) \rightarrow 1, (0) \rightarrow 1, (10-1) \rightarrow 0 or 10, (0-1) \rightarrow 0 or 10. Since the sequence memories (10-1) and (0-1) are followed by the same consequences, there is no reason why (10) and (0) should be differentially affected by generalization from the sequence memories, and thus no basis for predicting differential performance on Run 2, which the control group showed. Our hypothesis clearly does predict such differentiation, which follows simply from the contrast between the 10 or 0 received on Run 1, and the 1 anticipated on Run 2. This explanation does not exclude the possibility that the anticipation-generalization mechanism operates as well in Group E, over and above simultaneous contrast. That would be consistent with our finding that Group E showed stronger Run 2 discrimination than Group C in Experiment 2.

In our view, then, the idea that rats can remember ordered sequences such as 10-1 provides, at best, only a partial explanation for Group E's behavior on Run 2. We would further contend that even when

it seems most necessary, namely in explaining their behavior on Run 3, the concept of multiple hedonic memory has been stated more strongly than the data permit. What is shown in this case, taking our results and those of Capaldi and Verry (1981) together, is that rats can in some way discriminate 10-1 from 0-1, 20-0 from 0-0, 1-0 or 5-0, etc. Such discriminations can be achieved merely by remembering *whether or not* a particular event (e.g., a large 10- or 20-pellet reward) has previously been received during the trial; precise knowledge or reconstruction of the sequence 10-1, 20-0, or whatever, is not logically necessary. It is already known that animals can master tasks of the former kind. For example, D. Gaffan (1979) has shown that rhesus monkeys can learn to use the previous occurrence of reward or nonreward in association with one of a series of objects (independently of its order of occurrence in the series) as a cue for later responding. Richter and Kay (1980) and E. A. Gaffan and Davies (1982) found that rats spontaneously recalled the occurrence of a large, small, or zero reward in association with a particular maze arm even though other different reward events intervened. Performance of that kind indicates that animals can remember whether or not a particular hedonic event occurred some way back in a series, but carries no implication that they remember the entire sequence of events, because they are not required to do so. Similarly, we would argue that the present task does not require complete recall of the sequence. If it is desired to show that rats can recall a sequence such as 10-1 in its entirety, a different kind of discrimination task is needed; for example, 10-1 might be followed by a large reward, but 10-10, 1-10, and 1-1 by none.

Capaldi and Verry (1981) reported a related experiment (No. 5) in which a long series of similar events "branched," that is, changed at a specific point. The rats learned concurrently the sequence 0-0-0-0-20 and a sequence in which 20 was followed by a varying number of 0s. They gradually learned to sharply increase their speed on the fifth run of the 0-0-0-0-20 series, while maintaining a slow speed on all trials that followed an initial 20. We accept Capaldi and Verry's conclusion that in the first type of series the rats must have been capable of "counting" the number of preceding 0s in order to anticipate the terminal 20. However, their behavior in the 20-0-0-0 series requires only that they should recall whether a 20-pellet reward had occurred at some stage during the trial. As we have argued above, it is unnecessary to recall the entire series. Admittedly, even recalling just that 20 has occurred requires discrimination between within-trial and extratrial events, and must be achieved in the face of considerable retroactive interference; but it is not the same as the capacity to recall accurately a series of unlike events. We find it plausible that rats can

recall the number of like events (nonrewards) that have occurred—just as they can estimate the number of responses they have completed in a fixed-ratio schedule (e.g., Blackman & Scruton, 1973; Mechner & Guevrekian, 1962) or the number of sounds heard in a sequence (Fernandes & Church, 1982)—but see no compelling reason to conclude that they can recall the whole of a heterogeneous sequence.

Finally, while we may differ from Capaldi and Verry in our detailed interpretation of the phenomenon of multiple hedonic memory, we agree with their suggestion that "because of interference, it would make sense to employ multiple hedonic memory only as a last resort, after simpler solutions had failed" (Capaldi & Verry, 1981, p. 453). The multiple-memory solution is required in the concurrent-sequence paradigm, so rats learn to use it, albeit with considerable difficulty, as shown by the slow acquisition and inaccurate tracking. However, we see no reason to suppose that they use such a process in the single-sequence experiments reviewed by Hulse (1978). For those, the simpler model, which assumes memory of just one preceding hedonic event (Capaldi et al., 1980), seems to provide an adequate account.

REFERENCES

- BLACKMAN, D., & SCRUTON, P. Conditioned suppression of counting behavior in rats. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 93-100.
- CAPALDI, E. J. A sequential hypothesis of instrumental learning. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 1). New York: Academic Press, 1967.
- CAPALDI, E. J. Memory and learning; a sequential viewpoint. In W. K. Honig & P. H. R. James (Eds.), *Animal memory*. New York: Academic Press, 1971.
- CAPALDI, E. J. Latent discrimination learning under a regular schedule of partial reinforcement. *Animal Learning & Behavior*, 1979, 7, 63-68.
- CAPALDI, E. J., & MOLINA, P. Element discriminability as a determinant of serial pattern learning. *Animal Learning & Behavior*, 1979, 7, 318-322.
- CAPALDI, E. J., & VERRY, D. R. Serial order anticipation learning in rats: Memory for multiple hedonic events and their order. *Animal Learning & Behavior*, 1981, 9, 441-453.
- CAPALDI, E. J., VERRY, D. R., & DAVIDSON, T. L. Memory, serial anticipation pattern learning and transfer in rats. *Animal Learning & Behavior*, 1980, 8, 575-585.
- FERNANDES, D. M., & CHURCH, R. M. Discrimination of the number of sequential events by rats. *Animal Learning & Behavior*, 1982, 10, 171-176.
- FOUNTAIN, S. B., & HULSE, S. H. Extrapolation of serial stimulus patterns by rats. *Animal Learning & Behavior*, 1981, 9, 381-384.
- GAFFAN, D. Acquisition and forgetting in monkeys' memory for informational object-reward associations. *Learning and Motivation*, 1979, 10, 419-444.
- GAFFAN, E. A., & DAVIES, J. Reward, novelty and spontaneous alternation. *Quarterly Journal of Experimental Psychology*, 1982, 34B, 31-47.
- HAGGBLOOM, S. J. Intertrial interval effects on internal stimulus control of behavior in brightness differential conditioning. *Learning and Motivation*, 1978, 9, 347-358.

- HULSE, S. H. Cognitive structure and serial pattern learning by rats. In S. H. Hulse, H. Fowler, & W. K. Honig (Eds.), *Cognitive processes in animal behavior*. Hillsdale, N. J: Erlbaum, 1978.
- HULSE, S. H., & CAMPBELL, C. E. "Thinking ahead" in rat discrimination learning. *Animal Learning & Behavior*, 1975, 4, 305-311.
- HULSE, S. H., & DORSKY, N. P. Structural complexity as a determinant of serial pattern learning. *Learning and Motivation*, 1977, 8, 488-506.
- HULSE, S. H., & DORSKY, N. P. Serial pattern learning by rats: Transfer of a formally defined stimulus relationship and the significance of nonreinforcement. *Animal Learning & Behavior*, 1979, 7, 211-220.
- LEVENTHAL, A. M., MORRELL, R. F., MORGAN, E. F., JR., & PERKINS, C. C., JR. The relation between mean reward and mean reinforcement. *Journal of Experimental Psychology*, 1959, 57, 284-287.
- MACKINTOSH, N. J. *The psychology of animal learning*. London: Academic Press, 1974.
- MECHNER, F., & GUEVREKIAN, L. Effects of deprivation upon counting and timing in rats. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 463-466.
- RICHTER, M. L., & KAY, E. J. The roles of expectancy and after-effects in alternation. *Learning and Motivation*, 1980, 11, 503-521.
- TYLER, D. W., WORTZ, E. C., & BITTERMAN, M. E. The effect of random and alternating partial reinforcement on resistance to extinction in the rat. *American Journal of Psychology*, 1953, 66, 57-65.
- WIKE, E. L., & KING, D. D. Sequences of reward magnitude and runway performance. *Animal Learning & Behavior*, 1973, 1, 175-178.

(Manuscript received August 23, 1982;
revision accepted for publication November 29, 1982.)