

*INCREASING THE VARIABILITY OF RESPONSE
SEQUENCES IN PIGEONS BY ADJUSTING
THE FREQUENCY OF SWITCHING
BETWEEN TWO KEYS*

ARMANDO MACHADO

INDIANA UNIVERSITY

Three experiments compared the amounts of behavioral variability generated with two reinforcement rules. In Experiments 1 and 2 pigeons received food whenever they generated a sequence of eight pecks, distributed over two keys, provided that the sequence contained a certain number of changeovers between the keys. Although no variability was required—the birds could obtain all reinforcers by repeating the same sequence—the pigeons emitted a large number of different sequences. In Experiment 3 pigeons received food whenever they generated a sequence that had not occurred during the last 25 trials. After prolonged training, the birds showed more sequence variability than in the first two experiments. The analysis of the internal structure of the response sequences revealed that, in general, (a) the location of the first peck was highly stereotyped; (b) as the trial advanced, the probability of switching to the initially preferred key decreased whereas the probability of switching to the other key increased; and (c) a first-order Markov chain model with transition probabilities given by a logistic function accounted well for the internal structure of the birds' response sequences. These findings suggest that, to a large extent, the variability of response sequences is an indirect effect of adjustments in changeover frequency.

Key words: behavioral variability, switching behavior, nonstationary Markov chain model, response sequence, key peck, pigeon

When a class of responses is reinforced, the distributions of force, duration, latency, location, and topography across the members of the class typically become less variable (for a review see, e.g., Boulanger, Ingebos, Lahak, Machado, & Richelle, 1987). But reinforcement may also promote and maintain variation within and between response classes. For example, when porpoises received food for generating behavior that experienced trainers had not seen before, they produced new and highly variable behavior (Pryor, Haag, & O'Reilly, 1969). Similarly, when pigeons received food for generating sequences of eight left or right choices, provided that such sequences had not occurred during the last 25 trials, they produced high degrees of sequence variability (Page & Neuringer, 1985). Similar results have been reported with different procedures, animal species, and re-

sponse topographies (e.g., Blough, 1966; Bryant & Church, 1974; Hest, Haaren, & Van de Poll, 1989; Machado, 1989, 1992, 1993; Morgan & Neuringer, 1990; Morris, 1987, 1989; Neuringer, 1991; Schoenfeld, Harris, & Farmer, 1966; Shimp, 1967). Taken together, the two sets of studies suggest the following generalization: When reinforcement is contingent on response variation, response variability increases; when reinforcement is not contingent on response variation, response variability decreases.

However, it is still unclear how animals come to behave variably when reinforcement depends on response variability. According to one viewpoint, animals are directly sensitive to the variability requirements of the task. Thus Page and Neuringer (1985) suggested that when pigeons receive food for generating novel response sequences, an internal variability generator is activated and its output is tuned to the degree of variability demanded by the schedule. In the same vein, Stokes (1995) suggested that rats in a new task learn new classes of responses and how much variability to sustain within each class. A second viewpoint suggests that, at least in some of the studies mentioned above (e.g., Bryant & Church, 1974; Hest et al., 1989; Ma-

Parts of this paper were presented at the Winter Conference on Animal Learning, Winter Park, Colorado, 1996. I am grateful to Kathy Webster and Jennifer Burin for their help running the experiments, and to Francisco Silva, Munire Cevik, and Richard Keen for many helpful comments on early versions of the paper.

Correspondence concerning this article should be sent to Armando Machado, Department of Psychology, Indiana University, Bloomington, Indiana 47405 (E-mail: amachado@indiana.edu).

chado, 1989, 1992, 1993; Morgan & Neuringer, 1990; Page & Neuringer, 1985), response variability may have been a derivative of more fundamental processes. The nature of these processes and how they might engender response variability as one of their by-products is best illustrated by an example.

In Page and Neuringer's (1985) study, pigeons received food whenever they emitted an eight-peck sequence that had not occurred during the last 25 trials. This reinforcement rule instantiates two contingencies, one on sequence variability and the other on the behavior of switching between the two keys. The contingency on switching takes place because reinforced sequences are more likely to contain an intermediate number of changeovers (i.e., three or four) than too many (e.g., seven) or too few (e.g., zero). This bitonic relation between the number of changeovers per sequence and the probability of reinforcement holds because there are more eight-peck sequences with three or four changeovers than with any other number. Hence, the schedule used by Page and Neuringer not only reinforces sequence variability but also could shape the total number of changeovers per sequence towards intermediate values.

But how could response variability emerge from the simple effects of the schedule on changeover frequency? The question is pertinent because the relation between the frequency of switching and sequence variability is asymmetric: To generate variable sequences of choices one must switch between the response alternatives, but the converse is not necessarily true, because one can switch between the alternatives without generating variable sequences. However, the hypothesis I entertain here states that, for reasons identified below, the birds in Page and Neuringer's task may have been unable to reproduce eight-peck sequences with three or four changeovers (the sequences most likely to yield food) without simultaneously varying the location of the changeovers within the sequence and consequently without varying their response sequences.

Several reasons may explain why, in the absence of external cues, pigeons and rats do not switch at the exact same locations in a sequence of eight choices. First, the components of the sequence, left and right key

pecks or lever presses, are similar response events and consequently some generalization between them is likely to occur. Second, if a changeover after the first response (e.g., LRRRRRRR) is easy to learn, a changeover after seven responses (e.g., LLLLLLLR) is not, because the stimulus control function of number of responses, or elapsed time, shows increased variance as number or time increases (for a review see, e.g., Gallistel, 1990). Third, the occasional reinforcement of sequences that contain different numbers of changeovers may reduce the tendency to always produce three or four changeovers per sequence. I refer to these reasons collectively as limitations of stimulus control.

In summary, the approaches outlined above differ in how they conceptualize the relations between sequence variability and switching or changeover behavior. One approach implies that in Page and Neuringer's (1985) and similar studies the animals learned to vary their behavior and, as a consequence, most of their sequences contained three or four changeovers; the other implies that the animals learned to generate sequences with three or four changeovers and, as a consequence of stimulus control limitations, they also generated a high degree of sequence variability. In the former case, variability is fundamental, whereas the pattern of changeovers is a derivative; in the latter case, the frequency of switching and limitations of stimulus control are fundamental, whereas variability is a derivative.

The two viewpoints are not mutually exclusive. For example, pigeons and rats may be sensitive to both the variability requirements of the task, particularly when the response sequences are short and easily discriminable from one another, and the switching requirements, particularly when the sequences are long. In other words, pigeons and rats may discriminate that reinforcers are more likely to follow different sequences as well as sequences that contain an intermediate number of changeovers. If the preceding viewpoint is correct, then only empirical analyses can reveal how much each process contributes to a specific outcome. The present set of experiments initiated these analyses.

In Experiments 1 and 2 I asked how much sequence variability is generated when reinforcement depends exclusively on the num-

ber of changeovers in an eight-peck sequence. In both experiments, all reinforcers could be obtained by emitting only one sequence, which means that response variability, although permitted, was not required. My prediction was that because of the limitations of stimulus control mentioned above, sequence variability would be substantial, particularly when the distribution of switches was shaped towards the intermediate values of three or four. In order to assess the significance of the amounts of behavioral variability obtained in the first two experiments, Experiment 3 reproduced the conditions that in Page and Neuringer's (1985) study generated the highest degree of sequence variability. Experiment 3 provided a baseline against which the results from Experiments 1 and 2 could be compared.

Another goal of the present study was to characterize the serial structure of behavior in schedules that promote response variation. How sequences of behavior are internally organized remains a relatively neglected topic of research, presumably because experimenters have been concerned mainly with global effects (e.g., does variability increase when reinforcement depends on variability? How is the degree of behavioral variability affected by the type of reinforcement schedule?). Yet, the analysis of the serial order of behavior should help us to identify the processes that are responsible for maintaining behavior in a more or less variable condition. For example, if one finds that behavior, although highly variable, departs systematically from randomness (i.e., it contains some structure), then the type of structure will provide clues to the principles of serial order in behavior. Hence, the second part of this study analyzes the internal structure of the response sequences generated in switching-based (Experiment 2) and variability-based (Experiment 3) schedules of reinforcement.

EXPERIMENT 1

In the first experiment, reinforcement was contingent on minimal amounts of switching. Specifically, pigeons received food whenever they generated eight-peck sequences with at least one (Group 1) or two (Group 2) changeovers between two response keys. Three reasons motivated the experiment.

First, when pigeons and rats are trained to vary their sequences of choices and then receive the same amount of food without any variability constraint (yoked control procedure), their initially variable behavior gradually converges to the homogeneous patterns "all choices on the left" or "all choices on the right" (Hunziker, Saldana, & Neuringer, 1996; Machado, 1989, 1992, 1993; Page & Neuringer, 1985). Second, in my previous studies on behavioral variability, I observed that during the first session of variability training pigeons lose most of the available reinforcers because they frequently repeat these two homogeneous patterns. Third, even when behavior is variable, it is not uncommon to see that the patterns with zero or one changeover outnumber the patterns with more changeovers (e.g., Hunziker et al., 1996, Figure 4). Not surprisingly, perhaps, the patterns with zero or one changeover have greater unconditional strength and are more sensitive to reinforcement than other patterns. Hence, the first experiment was arranged to investigate how much sequence variability would be generated if all but these patterns were reinforced. If stimulus control is reduced as the birds switch more frequently between the keys, then sequence variability should be substantial even if all reinforcers could be obtained by repeating only one sequence.

METHOD

Subjects

Five experimentally naive pigeons (*Columba livia*) and 2 with previous experience in an autoshaping study participated in the experiment. The birds were housed in individual home cages, with water and grit continuously available, but with no dark-light cycle in effect. Throughout the experiment the birds were maintained at 80% of their free-feeding body weights.

Apparatus

A standard experimental chamber for pigeons from Med Associates® was used. The front aluminum panel contained three keys centered on the wall, 2 cm in diameter, 22 cm above the floor, and 8 cm apart, center to center. The keys could be illuminated from behind with red light. Because the right key

was not used during the present experiments it was covered with black tape. Directly below the center key and 4 cm from the floor was a hopper opening (6 cm by 7 cm). The bird had access to mixed grain when the hopper was raised and illuminated with a 7.5-W white light. On the back wall of the chamber, another 7.5-W houselight provided general illumination. An outer box equipped with a ventilating fan enclosed the experimental chamber. All events were controlled by a personal computer.

Procedure

The birds learned to peck the keys through a modified autoshaping procedure. After variable intertrial intervals ($M = 60$ s) during which the houselight was on, one of the two keys was randomly selected and lit. If no peck occurred for 6 s, the lit key was turned off and food was presented for 3 s; a peck at the lit key provided food immediately. During food presentations, the keylights and the houselight were turned off and the hopper light was lit. After three or four sessions, all birds pecked the keys reliably.

During the experiment proper, each daily session was divided into trials, and each trial began with the illumination of the houselight and both keylights. A peck at either key turned both keylights off and initiated a 0.4-s interpeck interval. Pecks during this period reset the timer for the interval but had no other scheduled consequences. After 0.4 s without a peck, both keys were illuminated again and the procedure was repeated for a total of eight pecks. After the eighth peck, either a 3-s timeout, during which all lights were off, or 3 s of access to food followed; then, a new trial began. Sessions ended after 100 trials or 50 reinforcers, whichever occurred first.

The 7 birds were assigned randomly to two groups, with the constraint that the two experienced birds had to be in different groups. In Group 1 (Birds 5291, 1782, 10798, and 10405), all eight-peck sequences with more than zero changeovers (i.e., all sequences except LLLLLLLL and RRRRRRRR) were reinforced. In Group 2 (Birds 2186, 25433, and 5269) all sequences containing more than one changeover were reinforced (of the total 256 eight-peck sequences, two contained zero changeovers and 14 contained one). The ex-

periment lasted at least 15 sessions and until the proportion of sequences that were reinforced showed no consistent trend for five consecutive sessions. The total number of sessions varied across birds from 15 to 25.

RESULTS AND DISCUSSION

Most sessions except the first ended before the 60th trial. Therefore, to make the present results directly comparable to those of Experiment 3 and to previous studies, only the data from the first 50 trials of each session were analyzed. In addition, because there were no significant differences between the two experimental groups, the results are presented without reference to the group membership of each bird. Unless otherwise stated, all analyses are based on the last three sessions of training.

The top panel of Figure 1 shows each bird's average proportion of reinforced sequences. For all birds except 10798, more than 95% of the sequences were reinforced, which implies that the number of changeovers per sequence was at least equal to the minimum required by the reinforcement schedule. The middle panel shows the averages of the number of changeovers per sequence. Clearly, all birds switched significantly more frequently than the minimum required to obtain food. Moreover, for 6 birds the average number of changeovers was above the value predicted by random responding. Hence, not reinforcing the patterns with zero changeovers, or with zero and one changeovers, was sufficient to generate a high frequency of switching between the keys.

The bottom panel of Figure 1 shows the average proportion of different sequences generated by each bird. A proportion of .6, for example, means that during the first 50 trials the bird emitted 30 different sequences on average. All birds generated a substantial number of different sequences ($M = 34.5$) even though only one was required to obtain food. However, the amount of sequence variability was consistently below the lower limit of the 95% confidence interval predicted by random responding (the Appendix derives the predictions for random responding).

Figure 2 shows that the high frequency of switching observed at the end of training was a true schedule effect. As predicted, during the first session the birds tended to generate




Fig. 1. Top: the horizontal lines show the proportions of reinforced sequences predicted by random responding for Groups 1 (.99) and 2 (.94). Middle: the horizontal line shows the value of 3.5 switches per sequence predicted by random responding. Bottom: the solid and dotted lines show the average proportion of different sequences and the 95% confidence interval predicted by random responding. The dark bars show averages across birds. All data come from the first 50 trials of the last three sessions of Experiment 1.

sequences with few or no changeovers, but by the last day all distributions had shifted considerably to the right. The average curves in the bottom right panel summarize the effect.

In conclusion, when reinforcement depended on minimal amounts of switching per eight-peck sequence, the birds switched substantially more frequently than required and earned most of the available reinforcers. With the increased frequency of switching came a substantial amount of sequence variability, despite the fact that variability was not required to obtain food. The switching distributions also suggest that sequence variability was not more substantial because the birds switched too much; as the number of changeovers increases from four to seven, the number of different sequences is reduced from 70 to two.

EXPERIMENT 2

The purpose of Experiment 2 was to reproduce more closely the contingencies on switching behavior that are embedded in most variability-inducing schedules. Thus, in Page and Neuringer's (1985) and similar studies, sequences with zero or seven changeovers are differentially extinguished when they occur frequently because only four different sequences can be produced with that number of changeovers. For similar reasons, sequences with one and six changeovers may be reinforced less often than sequences with two and five changeovers. More generally, the reinforcement rule in variability-inducing schedules should shape the distribution of switches towards intermediate values. When the birds emit sequences predominantly with three or four changeovers they are able to generate a substantial number of different sequences and collect most of the available reinforcers.

An earlier experiment by Bryant and Church (1974) also shows how an adequate shaping of switching behavior may generate high levels of response variability. In a two-choice situation, one group of rats received food with a probability of .75 every time they switched levers and with a probability of .25 every time they repeated the preceding choice. In two other groups these probabilities were .5/.5 and 1/0, respectively. The authors found that the rats in the .5/.5 group tended to persevere on one lever, those in the 1/0 group alternated frequently, and those in the .75/.25 group generated random-like behavior. These results suggest that response

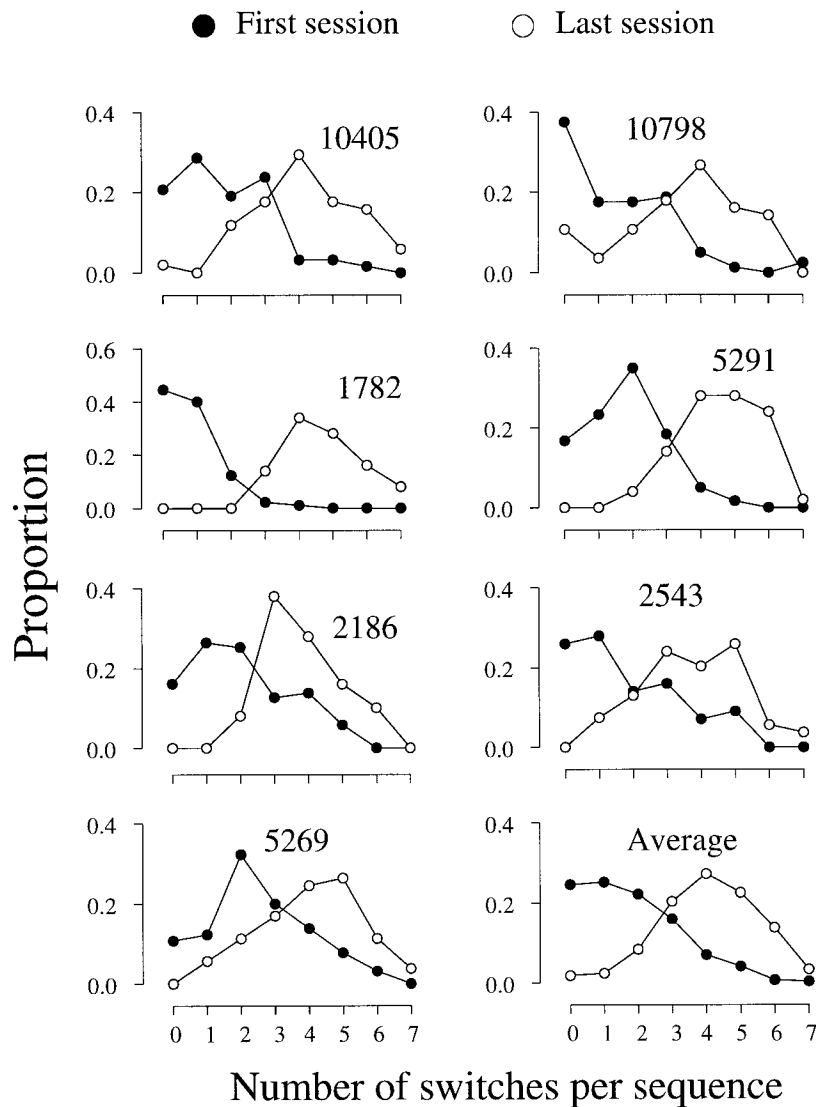


Fig. 2. Relative frequency distributions of the number of switches per sequence for each pigeon. The filled and open circles correspond to the first and last sessions of training, respectively.

variability may be increased if the reinforcement contingencies favor to some extent the rat's nonpreferred response alternative (switching).

Hence, in Experiment 2 I attempted to shape a binomial distribution of switches centered around 3.5, the distribution predicted by random responding. Sequences with too few or too many changeovers were less likely to be reinforced than were sequences with an intermediate number of changeovers. Hence, as in Experiment 1, reinforcement depended exclusively on the number of changeovers

per eight-peck sequence and not on variability. If sequence variability is indeed a by-product of switching, then it should increase as the number of sequences with three and four switches increases.

METHOD

Subjects and Apparatus

The birds, housing conditions, and experimental chamber were the same as in Experiment 1.

Procedure

All procedural details remained the same except that the probabilities of reinforcement varied with the number of changeovers per sequence according to a binomial distribution: Sequences with zero or seven changeovers were reinforced with a probability of .03, sequences with one or six were reinforced with a probability of .2, sequences with two or five were reinforced with a probability of .6, and sequences with three or four were reinforced with a probability of 1. These values are proportional to the total number of sequences that can be generated with a given number of switches. Thus, a total of 2, 14, 42, and 70 sequences are possible with 0, 1, 2, and 3 changeovers, respectively. With the probability of reinforcement equal to 1 for sequences with three changeovers, the remaining values were automatically set. The probabilities of reinforcement following the sequences with 4, 5, 6, and 7 changeovers preserved the symmetry of the binomial distribution.

Each session ended after 50 trials, and the experiment lasted 20 sessions for Bird 1782 and 25 sessions for the remaining birds. Although the number of reinforced sequences stabilized before 20 or 25 sessions, the additional sessions were conducted to collect the amount of data necessary to analyze the internal structure of the sequences.

RESULTS AND DISCUSSION

Preliminary analyses of the results revealed no significant differences between the data from the last three and the last six sessions. For this reason, the data from the last six sessions of training are reported. The top panel of Figure 3 shows that 5 birds collected almost as many reinforcers as those predicted by random responding (birds' average = 0.74; random = 0.77), but 2 birds earned less than 35% of the total reinforcers. Given this large difference between the two sets of birds, the median of all scores was selected to characterize the birds' performance. The middle panel of Figure 3 shows that, compared to Experiment 1, the average frequency of switching for all birds (except 2543) was reduced. For Birds 10405, 1782, 5291, and 2186 the average was between three and four switches, the optimal values; for Bird 2543 the average remained high at 4.3 switches, and




Fig. 3. Top: average proportion of reinforced sequences for each bird. The horizontal line shows the proportion predicted by random responding. Middle: average number of switches per sequence. Bottom: average proportion of different sequences. The data are from the last six sessions of Experiment 2. The dark bars show the group medians.

for Birds 10798 and 5269 switching decreased markedly to average values less than 1.5. Interestingly, the behavior of these 2 pigeons changed gradually, in one case after nine sessions during which the average switching frequency was 2.5 (Bird 10798), and in another case after eight sessions during which the av-

erage was 3.5 (Bird 5269). Both birds showed the terminal performance for more than 10 consecutive sessions.

The bottom panel of Figure 3 shows the proportion of different sequences generated by each bird. As one might have guessed from the preceding results, Birds 10798 and 5269 generated less than 40% of different sequences (in Experiment 1 both had generated more than 60%); for Birds 2543 and 2186 the proportion of different sequences decreased by about .1, for Birds 10405 and 1782 it increased by about .05, and for Bird 5291 it remained constant. Overall, variability decreased from Experiment 1 to Experiment 2 (the median proportion of different sequences, for example, decreased from .69 to .61).

In conclusion, at the end of Experiment 1 the birds were switching more often than predicted by random responding, and presumably for that reason the variability of their response sequences was not more substantial. Experiment 2 was implemented to increase sequence variability by reshaping the distribution of switches to intermediate values. The results show that 5 pigeons generated sequences with an average number of switches close to 3.5, the value predicted by random responding, but, contrary to my prediction, their behavior did not become more variable. For 2 other birds, the frequency of switching decreased so much that most reinforcers were lost, and the variability of their response sequences was greatly reduced. The reasons for the maladaptive behavior of these birds remain unclear.

In one respect, however, Experiment 2 corroborates the findings of Experiment 1, namely, that a substantial amount of sequence variability may be achieved by reinforcing the behavior of switching. In fact, the birds that switched three or four times per sequence on the average generated more than 30 different sequences for 20 or more consecutive sessions, even though only one sequence was required to obtain food. I will return to the results of Experiment 2 in a subsequent section concerned with the analysis of the internal structure of the response sequences.

EXPERIMENT 3

The first two experiments were predicated on the assumption that the variability ob-

served in some earlier studies (e.g., Bryant & Church, 1974; Hest et al., 1989; Machado, 1989, 1992, 1993; Morgan & Neuringer, 1990; Page & Neuringer, 1985) may have been due largely to the effects of reinforcement on switching, not to the reinforcement of variability per se. The results presented in Figures 1 and 3 show that simply by adjusting the frequency of switching between the two keys, as many as 30 to 35 different sequences may be generated during 50 trials. To assess the significance of these values, Experiment 3 reproduced the conditions that in Page and Neuringer's (1985) study generated the highest degree of sequence variability. If the amount of variability observed when the reinforcement rule is defined in terms of switching frequency (Experiments 1 and 2) matches the amount of variability observed when the reinforcement rule is defined in terms of sequence variability (Experiment 3), then one is more inclined to believe that sequence variability is not directly reinforced, or, equivalently, that sequence variability is a by-product of other processes. On the other hand, if the results show that when variations are explicitly reinforced, sequence variability is greater than when switching is reinforced, then the difference between the amounts of variability observed under the two conditions will help to quantify the specific effects of the contingencies of reinforcement on response variation.

Experiment 3 also provided the data that allowed a comparison of the structure of the response sequences observed in switching-based and variability-based schedules of reinforcement. If that structure proves to be very similar in the two tasks, there is reason to believe that the same behavioral processes underlie both types of performance.

METHOD

Subjects and Apparatus

Five experimentally naive pigeons (*Columba livia*) participated in the experiment. The housing conditions remained as in Experiment 1. Two experimental chambers were used, the one used in Experiments 1 and 2 and another identical one. Birds 9393 and 5320 were studied in the new chamber.

Procedure

All birds learned to peck the keys using the autoshaping training described in Experiment

1. The experiment proper followed an ABA design. During the A, or VAR, phases, reinforcement followed each eight-peck sequence that had not occurred during the last 25 trials. For example, if the bird produced the sequence LLRLLRR on Trial 35, it received food if that sequence had not occurred on Trials 10 to 34. If the sequence repeated any of the last 25 sequences, a timeout occurred. Starting with the second session, the last sequences of the previous session were used to decide whether to reinforce the first sequence of the new session. Thus, the first sequence of each session was compared with Sequences 50, 49, . . . , 26 of the previous session; the second sequence was compared with the first one and with Sequences 50, 49, . . . , 27 of the previous session, and so on. During the very first session the computer generated 25 sequences randomly to simulate the previous session.

Phase B, or NoVAR, was a self-yoked condition wherein variability was permitted but not required to obtain food (Page & Neuringer, 1985). Specifically, if the first VAR phase lasted for, say, 20 sessions, then the first session of NoVAR reproduced the order of reinforced and unreinforced trials of Session 15. That is to say, if during Session 15 reinforcement occurred on Trials 1, 3, and 5 and a timeout occurred on Trials 2, 4, and 6, then the same sequence of outcomes occurred on the first session of NoVAR, regardless of the bird's pattern of behavior. The second session of NoVAR reproduced the sequence of outcomes of Session 16, the third session those of Session 17, and so on, with the seventh session of NoVAR reproducing again the sequence of outcomes of Session 15. Thus, each session during NoVAR reproduced the sequence of outcomes of one of the last six sessions of the preceding VAR phase. After the NoVAR phase, the birds were reexposed to the VAR condition.

Each phase lasted until the proportions of reinforced and different sequences showed no consistent trend for five consecutive sessions. This criterion yielded 57 to 61 sessions during the first VAR phase, 11 to 19 sessions during NoVAR, and 12 to 29 sessions during the second VAR phase. All other procedural details remained as in Experiment 2. Data analyses are based on the last six sessions of each phase.

RESULTS AND DISCUSSION

Figure 4 shows that during the first VAR condition the proportions of reinforced trials ranged from .50 to .69, with an average of .60. The self-yoked condition forced the same values during NoVAR. During the second VAR condition, more trials ended with food and the proportions ranged from .61 to .78, with an average of .72. Despite the improvement in performance in the last phase, the proportion of reinforced trials was always significantly below the value expected from random performance.

Figure 4 also shows that when reinforcement required response variation, switching was more frequent than when reinforcement required no variation. On the other hand, the differences between the two VAR conditions were not large (average difference, .24; range, $-.16$ to $+.57$). In all conditions, the average number of changeovers per sequence was always below the value predicted by random responding.

Figure 4 also shows the proportion of different sequences generated by each bird. Three features of the data are noteworthy. First, sequence variability was more substantial when variability was required (VAR) than when it was simply permitted (NoVAR). This result agrees with previous findings (e.g., Machado, 1989; Page & Neuringer, 1985). Second, variability increased from the first to the second VAR phase, with the size of the effect varying from slight (Birds 10097 and 5320) to substantial (Bird 9393) (average difference, .11; range, .03 to .24). Third, the levels of variability during the VAR phases (.64 and .74) were always below the values predicted by random responding.

Experiment 3 implemented the reinforcement rule that in Page and Neuringer's (1985) study generated the highest degree of sequence variability, but it did not yield the same outcome. In fact, Page and Neuringer obtained an average of 85% of different sequences, whereas Experiment 3 obtained only 64% (first VAR) and 74% (second VAR phase). One reason for this difference may be that in Page and Neuringer's study the high variability requirement—each sequence had to differ from the last 25—was introduced only after the birds had experienced less stringent requirements (initially each se-




Fig. 4. The average proportion of reinforced sequences (top), the average number of switches (middle), and the average number of different sequences (bottom) obtained in Experiment 3. For each bird, the left and right open bars show the results from the first and last VAR phases, respectively; the dark bar shows the data from the NoVAR phase. The three rightmost columns show averages across birds. All data come from the last six sessions of each phase. The horizontal lines show the average (solid) and the 95% confidence interval (dotted) predicted by random responding.

quence had to differ from the last 5, then 10, and then 15 preceding sequences). The hypothesis that a gradual increase of the variability requirement leads to more variation than a nongradual increase remains to be tested.

The degrees of variability observed in Experiment 3 (64% and 74%) were close to the degree observed in Experiment 1 (69%) and were slightly greater than the median value obtained in Experiment 2 (61%). However, the fact that in Experiment 3 the average number of switches per sequence was lower than in the first two experiments suggests that the similarity in molar measures of variation (e.g., proportion of different sequences) may hide differences in the underlying behavioral processes. Whether this was indeed the case can only be ascertained by means of a more refined analysis of the serial structure of behavior.

THE INTERNAL STRUCTURE OF RESPONSE SEQUENCES

Several questions point to the importance of analyzing the serial structure of the response sequences: Why was behavior not more variable in Experiments 2 and 3? How did performance differ between these experiments? Given that most birds produced 30 or more different sequences during 50 trials, what characterizes those sequences? More generally, what behavioral process generated the sequences? The approach I followed to answer these and related questions consisted of hypothesizing a simple process of response generation, deriving its predictions concerning the internal structure of the sequences, and comparing these predictions against the data. If a simple process was rejected, then a slightly more complex one was studied next.

The analysis is restricted to the last six sessions of Experiment 2 and the last six sessions of the first VAR condition of Experiment 3, because the relatively large number of sessions in these two cases increases the chances that the molecular processes that determine the serial organization of the sequences had reached a steady state (behavior may be stable at the molar level after a few sessions but may still undergo molecular changes).

The simplest process of response generation is the Bernoulli process—the bird pecks

the right key with probability p and the left key with the complementary probability $1 - p$. When $p = .5$ this process corresponds to random responding. None of the data sets presented in Figures 1 to 4 support this model because, among other things, a Bernoulli process would always generate fewer than 3.5 changeovers per sequence on the average, whereas some data sets clearly show a higher number (see middle panels of Figures 1 and 3, and Figure 2).

As a variation of the above process, one might let the probability p vary within the sequence. That is, on the first peck of the sequence the bird could choose the right key with probability p_1 , on the second peck with probability p_2 , and so on, with p_1 not necessarily equal to p_2 . This nonstationary Bernoulli process could try to capture the combined effect of variables such as the increasing proximity to reinforcement as the sequence progresses and the effect of the preceding pecks on the location of subsequent pecks. Although plausible, a nonstationary Bernoulli process can also be rejected because, as I show below, it fails to predict the switching probability profiles generated by the birds.

A third candidate is a first-order Markov chain with stationary transition probabilities. Here, the location of the next peck is influenced only by the last peck. That is, the bird's last choice defines its current state (left or right), and the probabilities of choosing the left or right keys on the next peck depend on the state. Again, the profiles of switching probability presented below were inconsistent with this model.

Finally, I arrived at a process that describes the data well: a nonstationary first-order Markov chain. The model is illustrated in Figure 5. The changeover probabilities during the next choice depend on the previous choice (hence a first-order Markov chain) and how far in the sequence the bird is (hence nonstationary transition probabilities). The heart of the model is the description of how the two changeover probabilities, from the left to the right key, $p(n)$, and from the right to the left key, $q(n)$, change with peck number n . I assumed that $p(n)$ changes with n according to the equation

$$\frac{\Delta p}{\Delta n} = \alpha p(\beta - p), \quad (1)$$

where α is a rate parameter and β is one of p 's two fixed points (i.e., when $p = \beta$, no further changes in p take place; $p = 0$ is the other fixed point). A similar equation, possibly with different parameters, holds for $q(n)$. Next, I approximated this difference equation by its continuous version,

$$\begin{aligned} \frac{dp}{dt} &= \alpha p(\beta - p) \\ &= \gamma p \left(1 - \frac{p}{\beta}\right), \quad \gamma = \alpha\beta, \quad (2) \end{aligned}$$

and assumed that the eight pecks occurred at times $t = 0, 1, \dots, 7$.

The bottom panels of Figure 5 plot Equation 2. Depending on the sign of γ , two cases may occur. If $\gamma > 0$, then $dp/dt > 0$ for $p < \beta$ and therefore p will increase (see direction of arrows along the p axis); if $p > \beta$, then $dp/dt < 0$ and p will decrease. Regardless of its initial value, p will approach β , the only stable equilibrium. On the other hand, if $\gamma < 0$ (right panel), β is an unstable equilibrium and, for $p < \beta$, p will approach 0 (the other possibility, $p > \beta$, would eventually yield $p > 1$, which cannot occur if p is a probability).

The solution of Equation 2 is the logistic function

$$p(t) = \frac{\beta}{1 + \frac{\beta - p(0)}{p(0)} \exp(-\gamma t)}, \quad (3)$$

where $p(0)$ is the initial probability of pecking the right key (see Figure 5, top). In summary, the model implies that (a) the two changeover probabilities are the fundamental response units in the situation (in this regard, the model is similar to Myerson and Miezin's, 1980, kinetic model of choice), (b) the two types of changeover, from the left to the right and from the right to the left keys, follow independent courses, and (c) as the trial proceeds, the changeover probabilities change according to a logistic function. In what follows, the model is compared against the data from Experiments 2 and 3.

Response Sequences in Experiment 2

Figure 6 provides the empirical justification of the model. The data points are the switching probabilities observed in Experiment 2 as a function of position in the sequence. There was

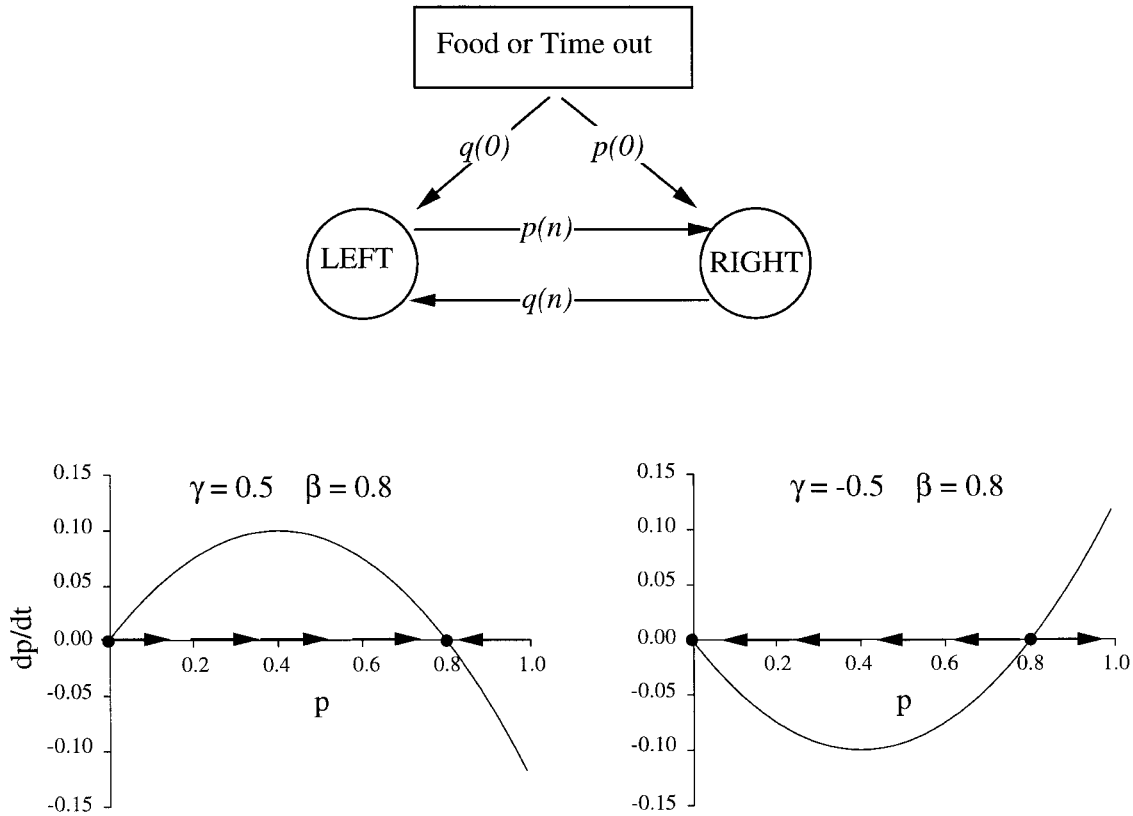


Fig. 5. Top: First-order Markov chain model with nonstationary transition probabilities. After food or a timeout, the bird pecks the right and left keys with probabilities $p(0)$ and $q(0)$, respectively, where $p(0) = 1 - q(0)$. Subsequent choice probabilities are given by $p(n)$ and $q(n)$, where n is the number of preceding pecks. For example, if the second peck was on the left key, the probability of pecking the right key on the third choice would equal $p(2)$; if the second peck was on the right, the probability of pecking the right key on the third choice would equal $1 - q(2)$. Bottom: plot of Equation 3. The growth parameter γ is positive on the left panel and negative on the right.

a remarkable consistency across birds. First, for all birds the two switching probability curves are clearly different, which means that the direction of switching matters. Second, the tendency to switch to one key increased as the trial advanced, whereas the tendency to switch to the other key either decreased (for 6 birds) or stayed constant (Bird 2543). Third, Birds 10798 and 5269 (top panels) fit the same general description even though they lost most of the available reinforcers because they rarely switched between the keys.

Another noteworthy feature of the data is that the birds that varied their sequences the most all showed a strong stereotypy on their first choice (first data points away from .5), but the birds that varied their sequences the least showed greater variability during their first choice (first data points close to .5). The

first peck emerges from this analysis as one of the organizing elements of response sequences, presumably because the event preceding it is markedly different from the events preceding subsequent pecks.

The solid lines in Figure 6 show the fits of Equation 3. To satisfy the constraint that $q(0) = 1 - p(0)$, $p(0)$ and $q(0)$ were set equal to their observed values (i.e., they were not treated as free parameters). Parameters γ and β were estimated by a nonlinear least squares algorithm. Table 1 shows the obtained parameter values. In general the model fitted the data well. The only exceptions came from the birds that switched so little, particularly late in the sequence, that some probabilities were estimated from very small numbers. In fact, the two most discordant points for Bird 10798 were not used to fit the model.

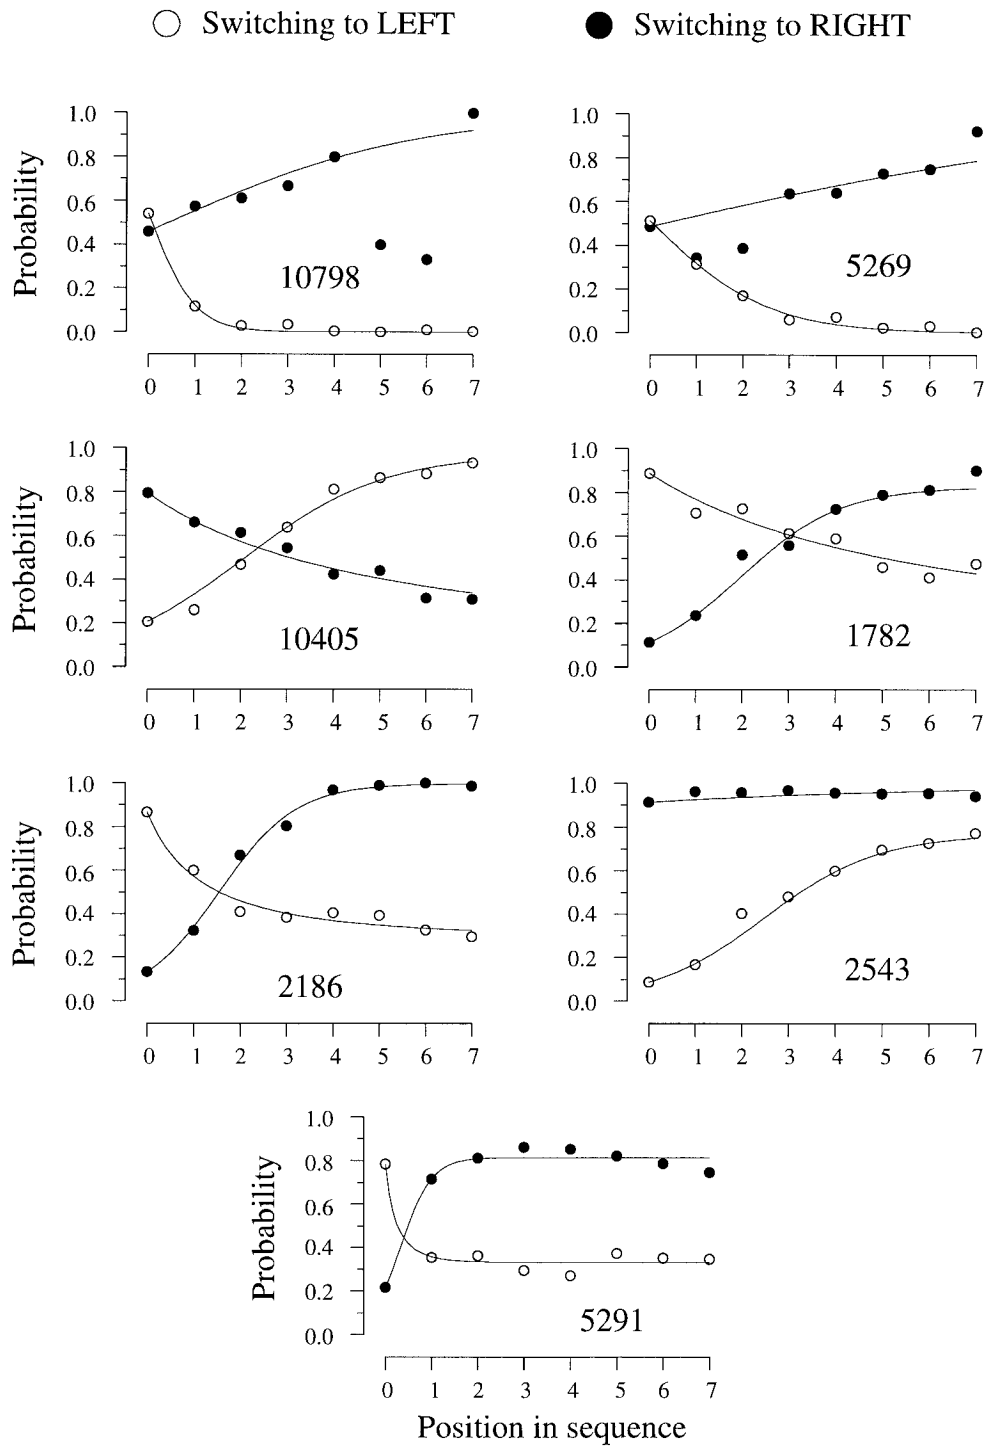


Fig. 6. Probability of switching to the left key (open circles) and to the right key (filled circles) as a function of position in the sequence. The first data point of each curve corresponds to the probability of starting the sequence on the corresponding key. The data are from the last six sessions of Experiment 2. The solid lines show fits based on Equation 3.

Table 1

Parameter values used to fit Equation 3 to the data of Experiments 2 and 3.

Bird		γ	β
Experiment 2			
10405	$p(n)$	0.000690	.00284
	$q(n)$	0.654	.994
2543	$p(n)$	0.160	1.00
	$q(n)$	0.908	.756
1782	$p(n)$	0.938	.853
	$q(n)$	0.00708	.0387
2186	$p(n)$	1.21	.997
	$q(n)$	0.325	.303
5269	$p(n)$	0.196	1.00
	$q(n)$	-0.815	1.00
5291	$p(n)$	3.01	.813
	$q(n)$	2.06	.332
10798	$p(n)$	0.376	1.00
	$q(n)$	-2.14	1.00
Experiment 3			
10097	$p(n)$	1.44	.518
	$q(n)$	0.304	.489
10770	$p(n)$	0.000710	.000790
	$q(n)$	0.452	1.00
5320	$p(n)$	1.97	.252
	$q(n)$	1.89	.973
5841	$p(n)$	-1.35	.382
	$q(n)$	1.11	.620
9393	$p(n)$	0.732	.845
	$q(n)$	-0.759	.855

If Equation 3 accurately describes the process underlying the birds' behavior, then it should fit not only the changeover probability curves displayed in Figure 6 but other properties of the response sequences as well. Furthermore, these additional properties should be fit with the same parameter values. Figure 7 shows that the model fit well the unconditional probabilities of pecking the two keys and the overall probability of switching as a function of position in the sequence. Figure 7 also shows why a nonstationary Bernoulli process or a first-order Markov chain with stationary transition probabilities does not fit the data. In fact, if one tries to predict the switching probabilities on the basis of the probabilities of pecking the left and right keys, the observed values are underestimated systematically. This fact suffices to reject the nonstationary Bernoulli process. Similarly, a first-order Markov chain also fails because it predicts a constant probability of

switching within the sequence, whereas all birds showed systematic deviations from a constant probability.

The preceding fits were all local fits because each data point was derived from one or two consecutive pecks. Figures 8 and 9 illustrate the fit of the model to statistics derived from more pecks. Figure 8 shows the observed and the predicted distributions of the number of switches per sequence. Not surprisingly, two distinct patterns can be observed. The top two panels show a distribution with a mode at one and a high frequency of sequences with zero switches. The remaining panels show more symmetric distributions. Interestingly, 4 birds (10405, 2186, 1782, and 5291) showed modes at three and five, even though the maximum probability of reinforcement occurred after sequences with three and four switches. The relatively high frequency of sequences with an odd number of switches occurred because these birds tended to start and end their sequences on different keys. (The filled circles in Figure 7 clearly show that the probability of pecking the right key was well below .5 on the first choice but was well above .5 on the last one.) The lower right panel of Figure 8 shows that the 5 birds that maintained a high degree of sequence variability switched slightly more than random responding would predict. In general, the model fits all switching distributions reasonably well.

Figure 9 highlights the response structure at the beginning of the trial in terms of run-length distributions. Again, the behavior of Birds 10798 and 5269 differed from the behavior of the remaining birds; their run-length frequency curves decreased monotonically from zero to seven, and the curve for the left key increased because many of these birds' sequences consisted of pecking the left key eight consecutive times. For the other birds the two distributions had different shapes; one peaked at zero because of the strong preference for the opposite key during the first choice; the other peaked at either one (Bird 5291) or three (Birds 2543, 10405, 2186, and 1782), which suggests some tendency to persevere early in the sequence. As the solid lines show, the model fits all data sets reasonably well.

In summary, when exposed to a schedule that differentially reinforced sequences with

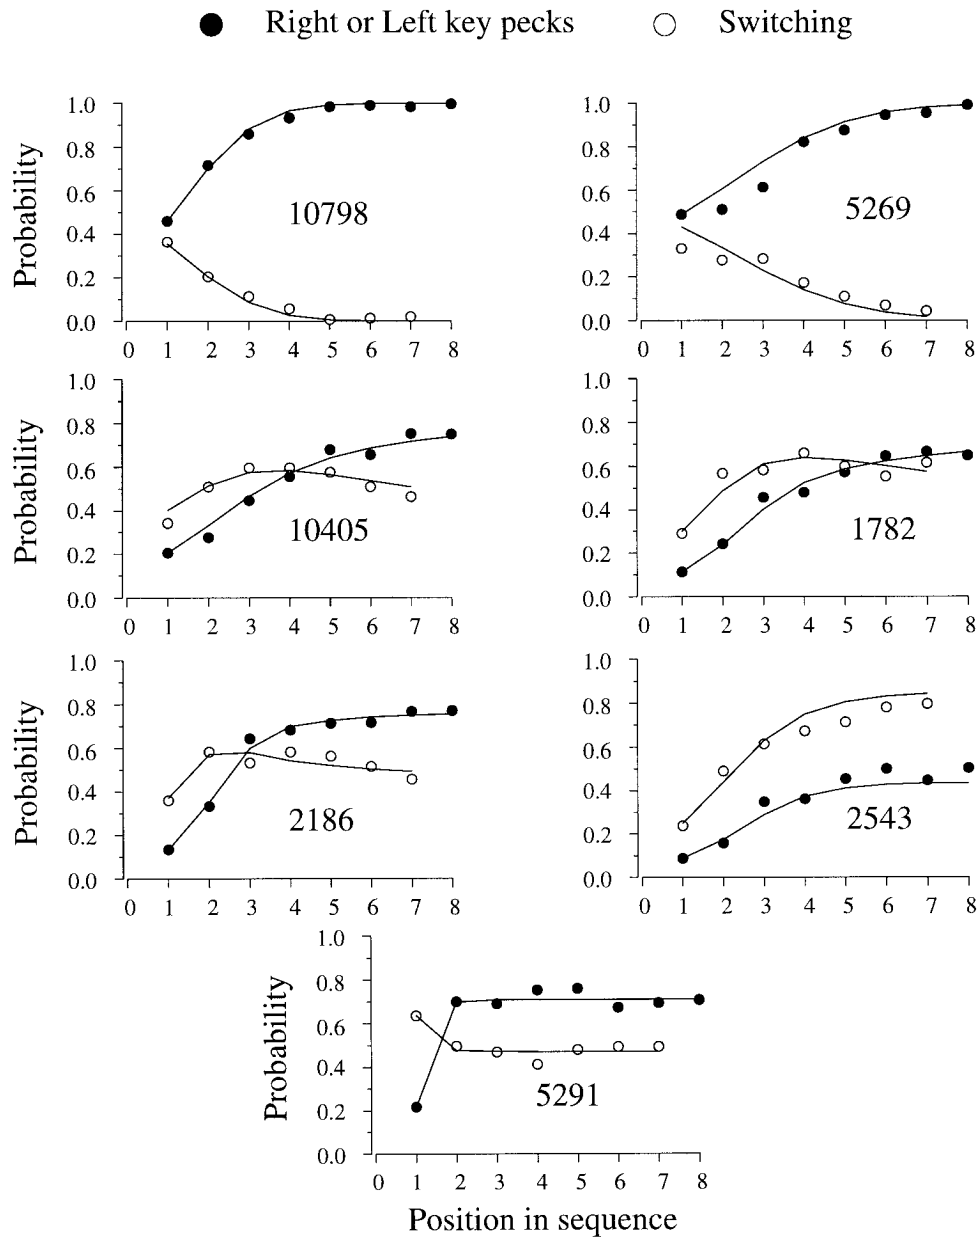


Fig. 7. Probability of right or left key pecks (filled circles) and overall probability of switching (open circles) as a function of position in the sequence. Switching is defined as either a left-right or a right-left sequence of pecks. The solid lines show the predictions of the model. The data are from the last six sessions of Experiment 2.

an intermediate number of changeovers, the majority of the birds adapted to the schedule by (a) starting most of the sequences on one key, (b) increasing the probability of switching to the other key as the trial advanced, and (c) decreasing or maintaining constant the probability of switching to the initially pre-

ferred key. By adjusting the rates at which the switching probabilities changed along the trial, 5 birds managed to generate most of their sequences with the optimal or close to the optimal number of changeovers. In the process they also generated a large number of different sequences.

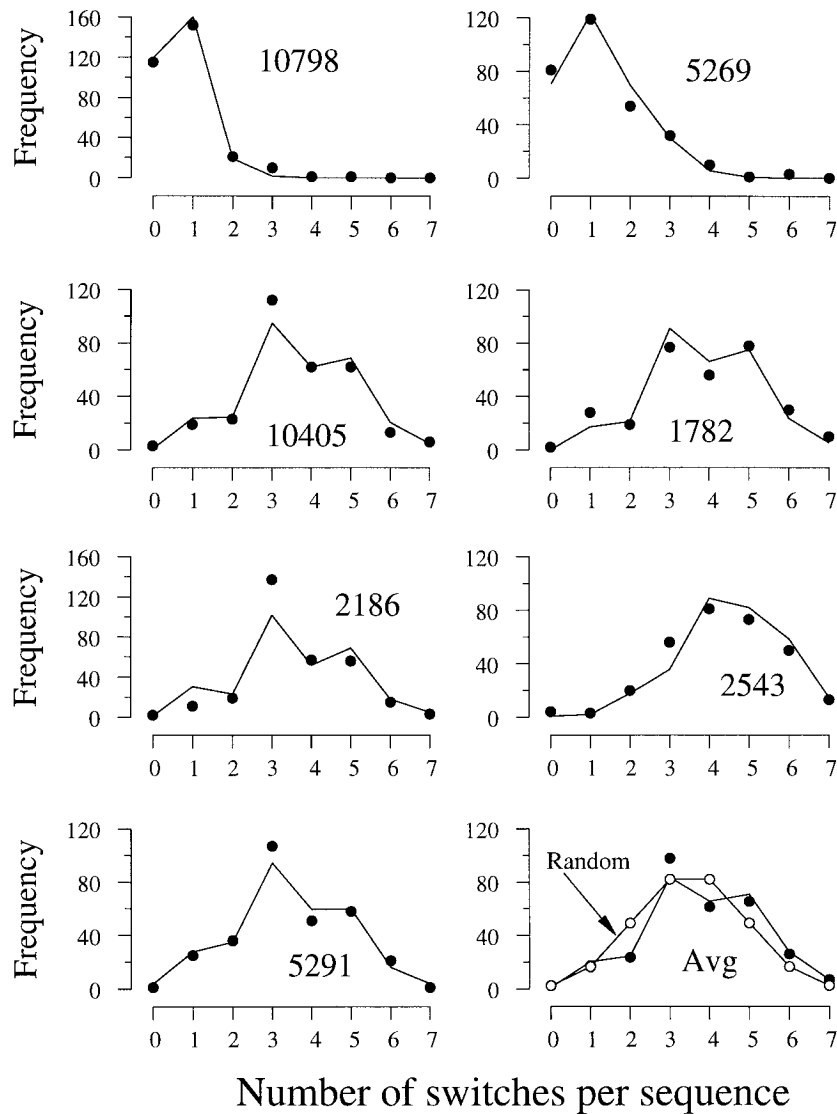


Fig. 8. Observed (circles) and predicted (lines) distributions of the number of switches per sequence. The bottom right panel shows the average of all birds except 10798 and 5269 (filled circles), the corresponding average predicted by the model (solid line), and the binomial distribution predicted by random responding (open circles). The data are from the last six sessions of Experiment 2.

Response Sequences in Experiment 3

To what extent does the previous model account for the data when sequence variability is explicitly reinforced? If the same account holds, then it seems plausible that the same behavioral processes could have been engaged by the two schedules and, therefore, that switching adjustments and limitations of stimulus control may play a large role in what Page and Neuringer (1985) called operant

variability. However, if the model fails, then one is forced to conclude that different processes are at play when switching and variability are the targets of reinforcement. A third alternative is possible: The same model may hold but its parameters, and hence the shape of the curves, are systematically different in the two experiments. In this case one should conclude that the behavioral processes are the same but that the specifics of each

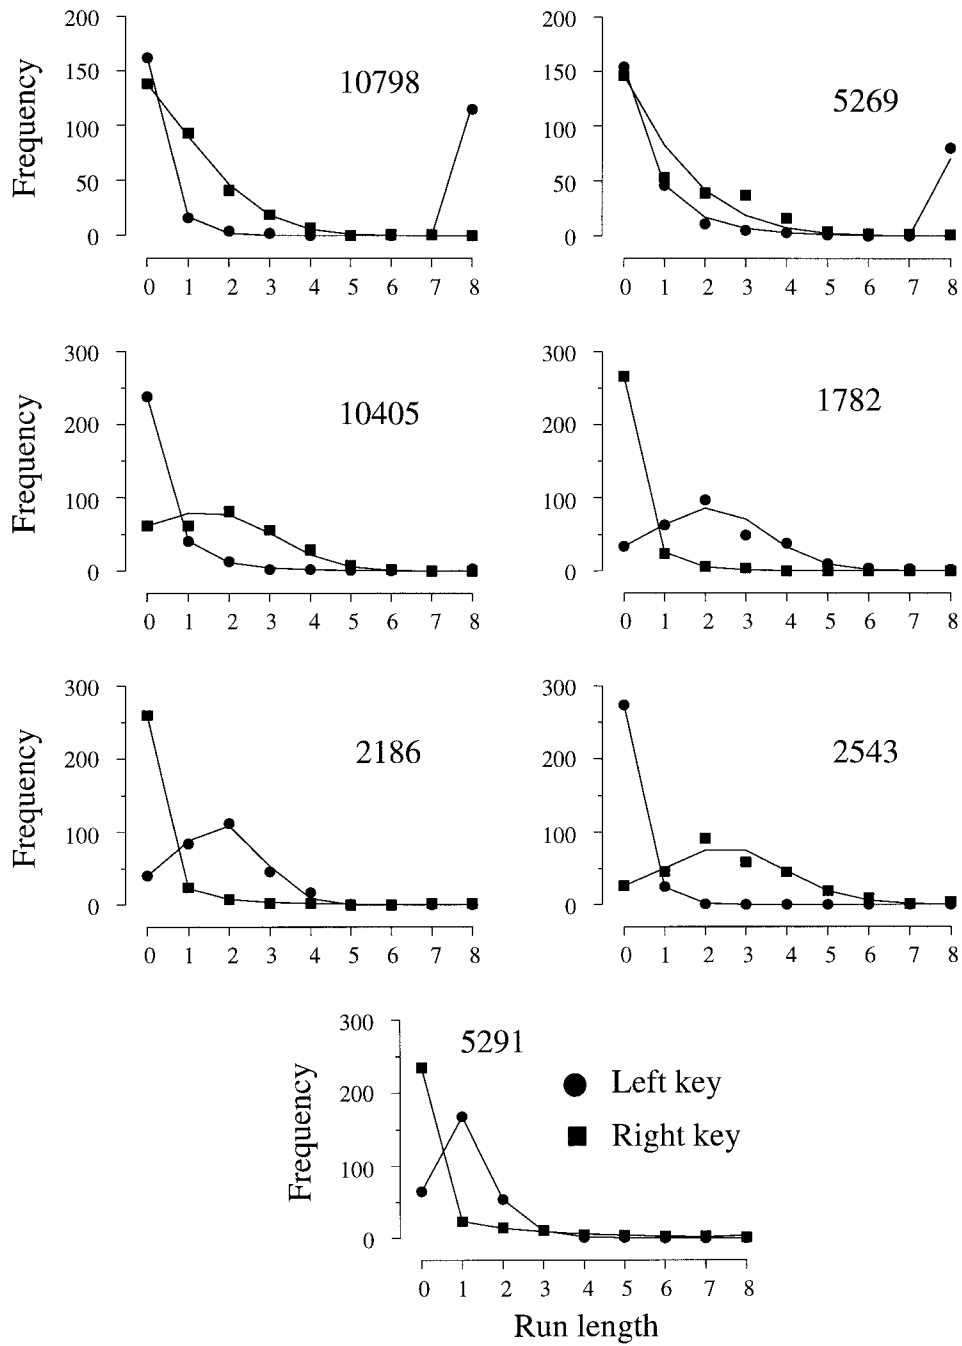


Fig. 9. Observed (circles and squares) and predicted (solid lines) distributions of run lengths at the beginning of a sequence. Circles and squares are for runs of left and right key pecks, respectively. A run length of three on the left key, for example, includes all the sequences that start with the pattern LLLR. The data are from the last six sessions of Experiment 2.

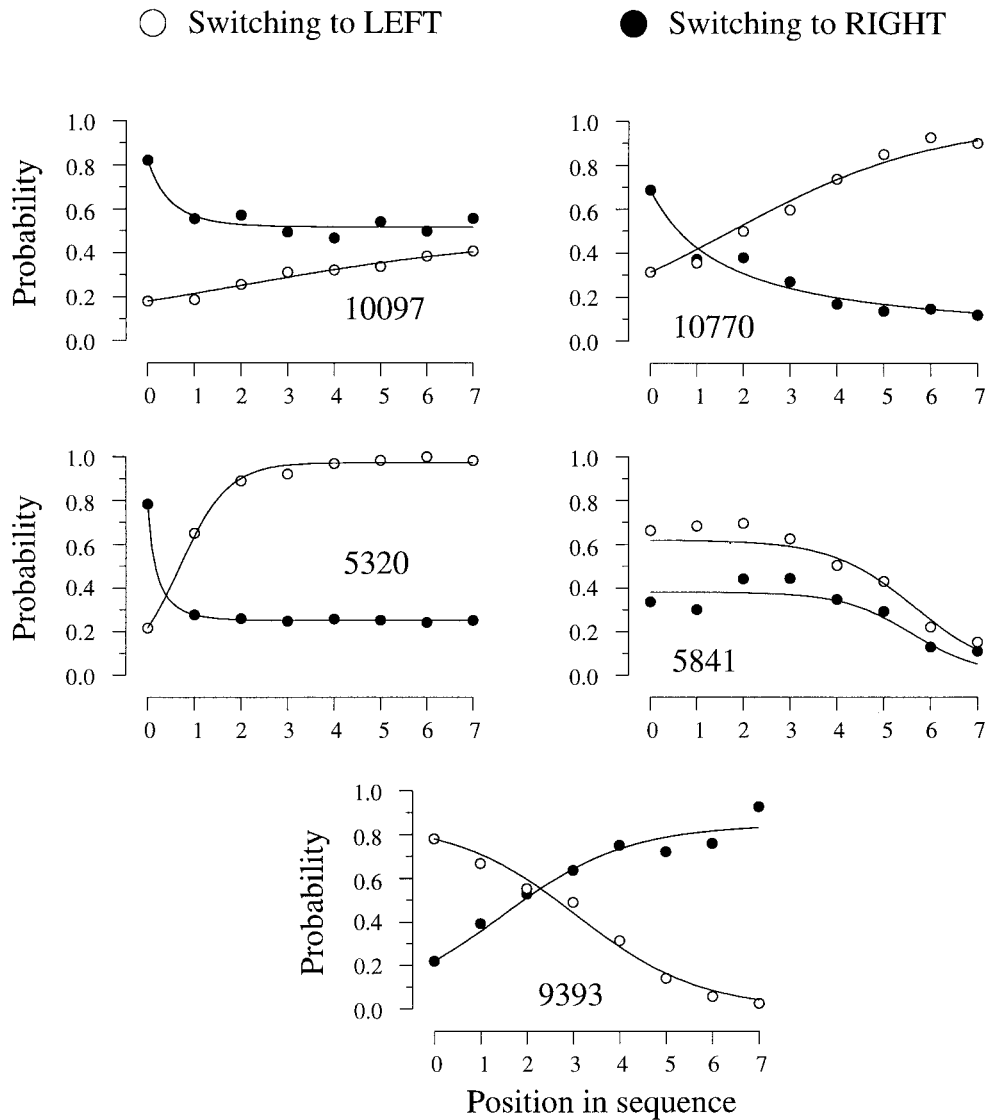


Fig. 10. Probability of switching to the left key (open circles) and to the right key (filled circles) as a function of position in the sequence. The first data point of each curve corresponds to the probability of starting the sequence on the corresponding key. The data are from the last six sessions of the first VAR phase of Experiment 3. The solid lines show fits based on Equation 3.

schedule lead these processes into different directions.

Figure 10 shows the switching probabilities when sequence variability was explicitly reinforced, and Table 1 lists the parameters for each curve fit to the data. For 4 birds, the shape of the curves is similar to that observed in Experiment 2. That is, the birds showed a marked preference for one of the keys on their first choice (first data points away from

.5). Afterwards the probability of switching to the initially preferred key decreased while the other changeover probability increased. Bird 5841 is the exception because both probabilities decreased. [For this bird, goodness of fit is reduced markedly if $p(0)$ and $q(0)$ equal the observed values. Hence, I treated $p(0)$ as a free parameter and then set $q(0) = 1 - p(0)$.]

Figures 11, 12, and 13 test whether the

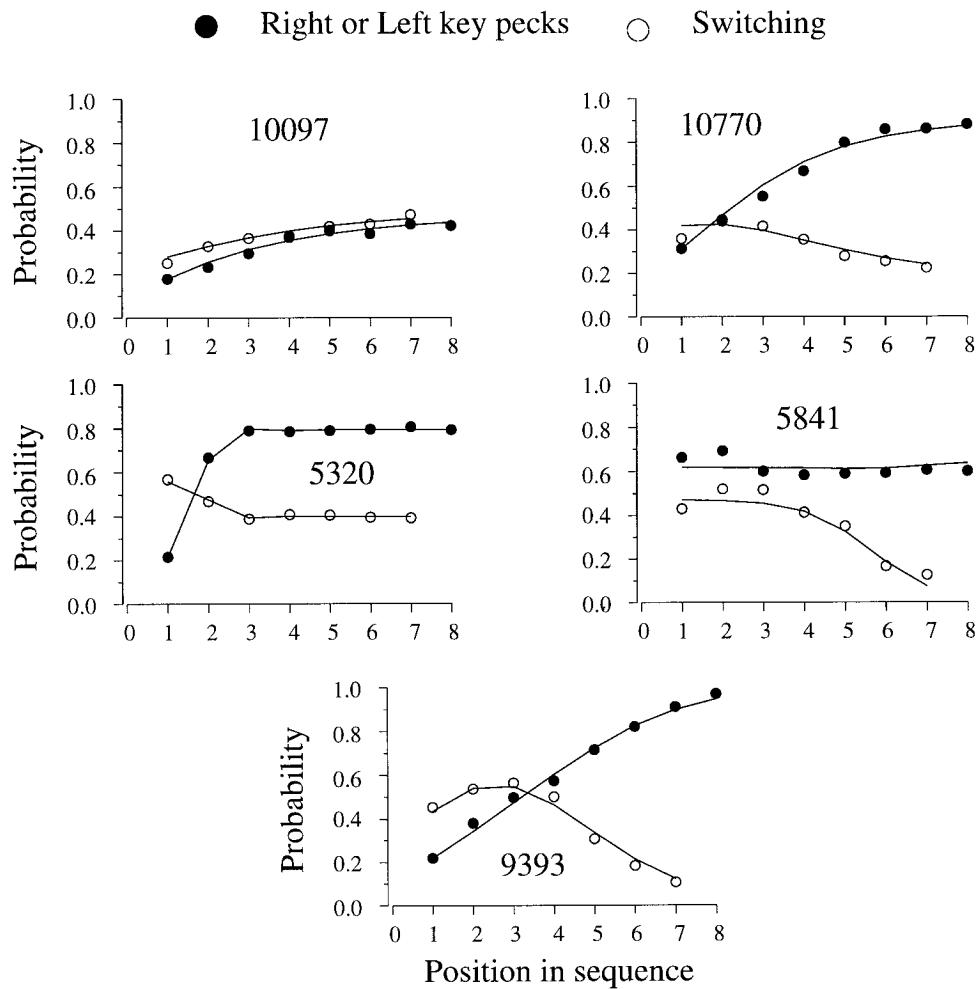


Fig. 11. Probability of right or left key pecks (filled circles) and overall probability of switching (open circles) as a function of position in the sequence. The solid lines show the predictions of the model. The data are from the last six sessions of the first VAR phase of Experiment 3.

model could fit additional properties of the response sequences using the same parameter values. Figure 11 shows that the model fit well the overall probabilities of switching and the unconditional probabilities of pecking the two keys; Figure 12 shows that it fit the switching distributions well, and Figure 13 shows the distributions of run length. Figures 12 and 13 also reveal additional similarities and differences between the data from Experiments 2 and 3. For example, 3 of the 5 birds (10070, 5320, and 9393) tended to repeat sequences with an odd number of switches, particularly with one and three (see Figure 12). This result was also observed in Experiment 2 (see Figure 8), except that the

peaks in Experiment 2 occurred at three and five. On the other hand, the average curve in Figure 12 indicates that when variability was explicitly reinforced, the birds switched less often than expected from random responding, whereas in Experiment 2 (see Figure 8) the birds switched slightly more often than predicted by random choices. Figure 13 reveals that in Experiment 3 the distributions of run length for the right and left keys differed for all birds. The strong preference for one key during the first choice yielded the fast-dropping curves in the figure. The curves for the other key were bitonic in two cases, with a mode at one, or monotonic but with a rather slow rate of decay with increasing run

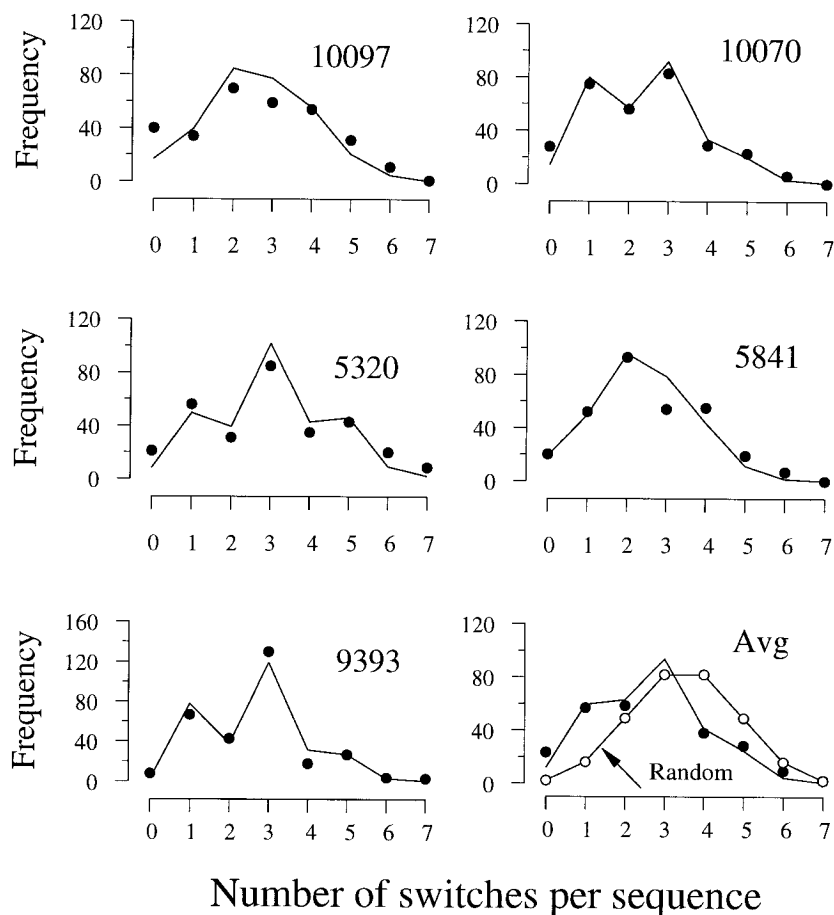


Fig. 12. Observed (circles) and predicted (lines) distributions of the number of switches per sequence. The bottom right panel shows the average of all birds (filled circles), the corresponding average predicted by the model (solid line), and the binomial distribution predicted by random responding (open circles). The data are from the last six sessions of the first VAR phase of Experiment 3.

length. These results are qualitatively similar to those observed in Experiment 2.

In summary, when reinforcement depended on sequence variability, the pigeons adapted by (a) starting most of their sequences on one key; as the trial advanced, (b) the probability of switching to the other key generally increased, whereas (c) the probability of switching to the initially preferred key generally decreased. However, for 1 bird the two changeover probabilities decreased during the trial. The model used to fit the data obtained when switching was explicitly shaped (Experiment 2) also provided a good description of the data obtained when variability was explicitly reinforced (Experiment 3).

GENERAL DISCUSSION

The experimental study of a behavioral phenomenon tends to proceed along two avenues: the analysis of the conditions that are necessary and sufficient to generate the phenomenon, that is, the analysis of its causes, and the synthesis of the phenomenon as a way to test our understanding of its causes. Quite often investigators are able to synthesize a behavioral phenomenon without clearly understanding its causes. Such is the case of behavioral variability, because although several studies have shown that response variation can be increased by making reinforcement contingent on variation, the causes of this phenomenon and how they are interre-

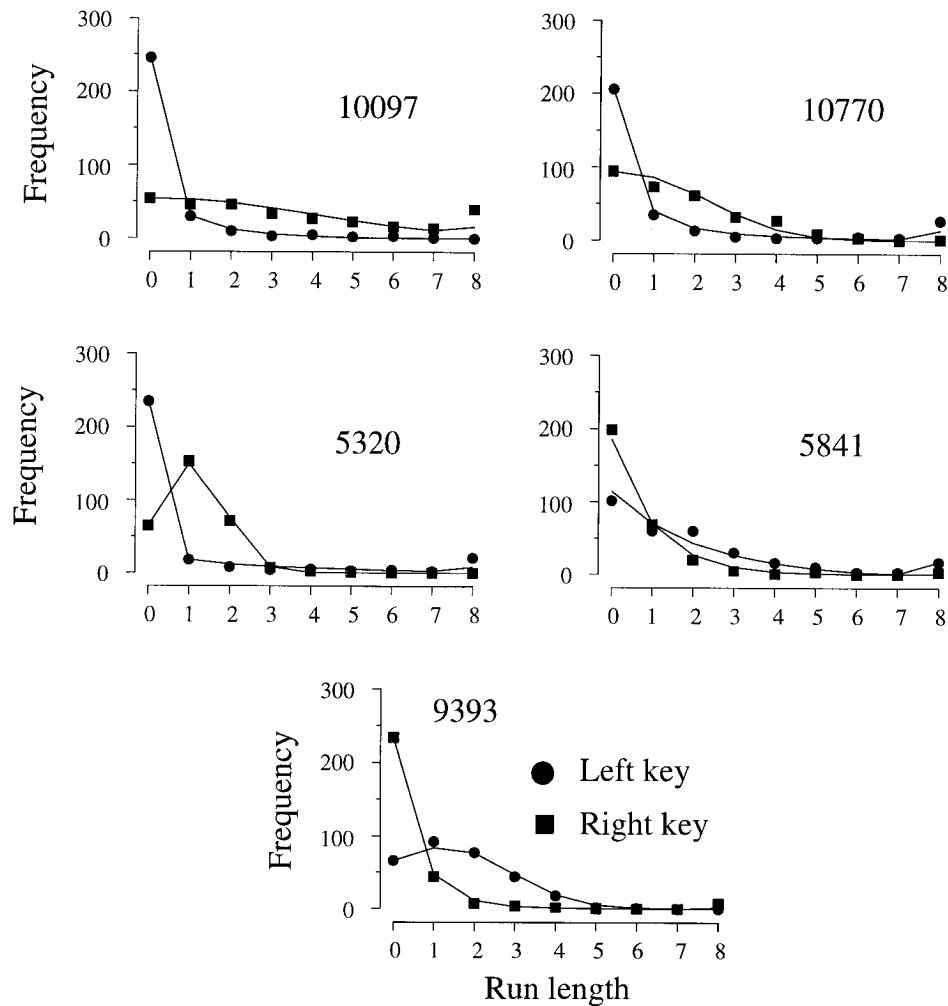


Fig. 13. Observed (circles and squares) and predicted (lines) distributions of run lengths at the beginning of a sequence. Circles and squares are for runs of left and right key pecks, respectively. The data are from the last six sessions of the first VAR phase of Experiment 3.

lated remain unclear. The experiments reported here attempted to fill this gap.

Molar Results

The first two experiments investigated the role of switching in promoting the variability of response sequences. Experiment 1 showed that when reinforcement followed sequences that contained one or more changeovers between the keys, sequence variability was substantial. Similarly, Experiment 2 showed that 5 of 7 birds maintained a relatively high degree of sequence variability when the reinforcement schedule shaped the number of switches per eight-peck sequence towards the intermediate

values of three or four. Although only one sequence was required to obtain all reinforcers in both experiments, most birds generated more than 30 different sequences in the first 50 trials. Experiment 3 showed that when sequence variability was explicitly reinforced, the birds varied their behavior more than when the same pattern of reinforcers and nonreinforcers was delivered regardless of sequence variability (Machado, 1989; Page & Neuringer, 1985). Furthermore, at the end of the experiment, the birds generated a higher number of different sequences, but a lower average number of switches per sequence, than in the first two experiments.

The molar results from the three experiments suggest the following interpretation. Initially the birds tended to repeat sequences with few or no changeovers between the keys (see Figure 2). If the schedule differentially extinguished these sequences, then the birds increased their frequency of switching (see the middle panels of Figures 1, 3, and 4) and generated a relatively large number of different sequences (see the bottom panels of Figures 1 and 3). This large amount of sequence variability does not seem to be due to any direct effects of reinforcement on variability but to limitations of stimulus control. For example, when the sequence RLLRRRRL is reinforced, the bird may not be able to repeat that sequence exactly because it will not be able to switch precisely after the first, third, and seventh pecks. The bird's difficulty in repeating sequences with an intermediate number of changeovers may stem from the similarity of the two responses, the increasing variance in the stimulus control function of the number of previous pecks and the time since the beginning of the sequence, or the occasional reinforcement of sequences with much higher or much lower number of switches.

When variability was explicitly targeted, the birds generated slightly more different sequences than when reinforcement depended only on switching. The reasons for this small difference remain unclear. On the one hand, it could be due to a direct effect on response variability of the reinforcement rule used in Experiment 3. For example, the birds in that experiment may have learned after prolonged training that reinforcement followed only relatively new sequences, that is, sequences that had not been emitted recently. If so, then variability is indeed an operant, as Neuringer and his collaborators have argued. On the other hand, it is also possible that the differences in the degree of variability were due to differences in how the two types of schedules interacted with the switching behavior of the bird. That is to say, although Experiments 1 and 2 attempted to reproduce the contingencies on switching that are embedded in variability-inducing schedules, some additional contingencies may have been omitted. For example, when the number of switches per sequence increased, the probability of reinforcement for the two homogeneous patterns LLLLLLLL

and RRRRRRRR also increased in Experiment 3 but not in Experiments 1 and 2. This dynamic property may explain why the average number of switches per sequence was lower in the last experiment (compare the middle panels of Figures 1, 3, and 4). The general point is that we cannot rule out the hypothesis that such differences underlay the differences in the degree of sequence variability across the three experiments. The hypothesis could be tested by shaping the distribution of switches to different values (two, three, etc.) and measuring the resulting degree of sequence variability. If the hypothesis proves to be true, then switching behavior, not variability, is the operant.

The Internal Structure of Response Sequences

The similarity between the internal structure of the response sequences observed in Experiments 2 and 3 further suggests that the two types of schedules engaged the same processes. In both cases the sequences were characterized by three major attributes: the location of the first peck, the probability of switching in one direction, and the probability of switching in the opposite direction. Specifically, the location of the first peck tended to be highly stereotyped in most birds, even though in Experiment 2 the stereotypy of the first peck did not affect the outcome of the trial, whereas in Experiment 3 such strong stereotypy reduced the chances of reinforcement. Also, the probability of switching to the initially preferred key decreased with the number of pecks within the sequence, whereas the probability of switching in the opposite direction generally increased. However, a few birds in both experiments did not fit the preceding description, because they showed either increased variability on their first choice or a decrease of both changeover probabilities as the trial advanced.

The behavior of all birds was well described by a first-order Markov chain with nonstationary transition probabilities. The good fit of the model further emphasizes the following points. First, because the two changeover probability functions $p(n)$ and $q(n)$ were generally different, a Bernoulli process cannot account for the birds' behavior. Second, because $p(n)$ and $q(n)$ increased or decreased in an orderly fashion with peck number, n , a

first-order Markov chain with stationary transition probabilities can also be ruled out. Third, even though alternative models cannot be ruled out at this stage, the ability of the changeover model to predict a variety of other properties of behavior—the proportion of right key pecks, the overall switching proportions, the distribution of the number of switches per sequence, and the run lengths at the beginning of the sequences—is a powerful indicator that the model captured the major features of the process that structured the birds' behavior in the three experiments.

Equation 2 states that the rate of change of the probability of switching is affected by two factors, an intrinsic growth rate, γ , and a maximum switching probability β . For most birds γ was positive and β was close to one for one response and close to zero for the other response. Together with the information that $p(0)$ was generally far from .5, these parameters indicate that the discriminative stimulus functions of the two keys differed. The birds tended to perseverate on one key for the first few pecks and then switch to the other key. Afterwards, whenever they returned to the initially preferred key, they almost invariably switched back on the next choice. This strategy assured that switching was not so frequent that it lowered the probability of reinforcement and, in addition, that switching was contiguous with the reinforcer (see Figure 7, open circles).

For other birds, γ was negative and p (or q) decreased monotonically to zero. This strategy leads to a concentration of switching early in the sequence (cf. Figure 7, Birds 10798 and 5269, and Figure 11, Birds 5841 and 9393) and therefore it is potentially unstable—reducing the contiguity between switching and the reinforcer could lead to even less frequent switching late in the sequence. However, as switching frequency decreases so does reinforcement rate, which may evoke some switching. The net effect of these two forces, one tending to decrease switching frequency and the other tending to increase it, is likely to depend on each bird's delay-of-reinforcement gradient and resistance to extinction.

The generality of the changeover model cannot be fully assessed because no study of sequence variability has analyzed in detail the internal structure of the animals' choices. On

the other hand, the model does suggest more effective ways to increase response variability. For example, the reinforcement rule should not be based on a composite, such as the total number of changeovers per sequence, that ignores the direction of the changeover. In other words, the schedule should treat the two changeovers independently, and should explicitly attempt to equalize the probabilities $p(n)$ and $q(n)$. Furthermore, to the extent that these changeover probabilities can adjust in a period of 10 s or less (the typical duration of an eight-peck sequence), the schedule should also consider the serial order of the changeovers. To implement this requirement, the schedule could give more weight to the first and the last pecks of the sequence and less weight to the middle ones (Machado & Cevik, 1997).

Explaining Random-Like Behavior

The use of a stochastic model to account for behavioral variability may initially look suspicious. After all, if the model assumes random variables p and q , then does it not explain response variability by fiat? The question raises the more general issue of how we should deal theoretically with highly variable, random-like behavior. For example, assume that in one experiment a pigeon behaved in a way that is practically indistinguishable from random responding (i.e., an overall proportion of right key pecks equal to .5, changeover probabilities always at .5, a binomial distribution of the number of switches centered at 3.5, a geometric distribution of run lengths, etc.). What would constitute an adequate explanation of its behavior? One possibility is to say that the bird did not learn to respond randomly; rather, it unlearned to respond in specific ways. For example, its bias for one key during the first choice was eliminated, the strong tendency to switch in one direction was weakened, and the like. According to this point of view, random-like responding is what is left when all sources of bias are eliminated (Machado, 1993; Neuringer, 1986). In terms of the model, random responding would be approached as the parameter γ goes to 0 and $p(0)$ goes to .5. This negative definition of randomness immediately suggests that random-like responding may be very difficult, if not impossible, to obtain, and, perhaps more important, that our theories should not attempt to explain ran-

dom-like behavior per se, but rather the type and amount of structure still present in the animal's performance. In other words, the theoretical approach suggested here considers random-like behavior as a limit condition not to be directly explained. Instead, our theories should focus on the types of environmental conditions that generate particular types of behavioral structures (e.g., the alternation between the two keys), the conditions that determine the strength of those structures (the rate or probability of alternation) and, by extrapolation, the conditions that may strongly reduce those structures (variability-inducing schedules) and thereby approach random-like performance.

In conclusion, the present study shows that when reinforcement depends on the behavior of switching between two keys, pigeons automatically vary their sequences of choices. The experiments also suggest that reinforcing sequence variability explicitly is slightly more effective at generating response variation than is reinforcing switching only. The molecular analyses of the birds' performance revealed three organizing elements of their response sequences: the location of the first peck and the two probabilities of a changeover. A first-order Markov chain model with transition probabilities given by a logistic function described the birds' behavior well.

REFERENCES

- Blough, D. S. (1966). The reinforcement of least-frequent interresponse times. *Journal of the Experimental Analysis of Behavior*, 9, 581–591.
- Boulanger, B., Ingebos, A., Lahak, M., Machado, A., & Richelle, M. (1987). Variabilité comportementale et conditionnement opérant chez l'animal [Behavioral variability and operant conditioning in animals]. *L'Année Psychologique*, 87, 417–434.
- Bryant, D., & Church, R. M. (1974). The determinants of random choice. *Animal Learning & Behavior*, 2, 245–248.
- Gallistel, R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Hest, A. N., Haaren, F. V., & Van De Poll, N. E. (1989). Operant conditioning of response variability in male and female Wistar rats. *Physiology and Behavior*, 45, 551–555.
- Hunziker, M. H. L., Saldana, R. L., & Neuringer, A. (1996). Behavioral variability in SHR and WKY rats as a function of rearing environment and reinforcement contingency. *Journal of the Experimental Analysis of Behavior*, 65, 129–143.
- Machado, A. (1989). Operant conditioning of behavioral variability using a percentile reinforcement schedule. *Journal of the Experimental Analysis of Behavior*, 52, 155–166.
- Machado, A. (1992). Behavioral variability and frequency-dependent selection. *Journal of the Experimental Analysis of Behavior*, 58, 241–263.
- Machado, A. (1993). Learning variable and stereotypical sequences of responses: Some data and a new model. *Behavioral Processes*, 30, 103–130.
- Machado, A., & Cevik, M. (1997). The discrimination of relative frequency by pigeons. *Journal of the Experimental Analysis of Behavior*, 67, 11–41.
- Morgan, L., & Neuringer, A. (1990). Behavioral variability as a function of response topography and reinforcement contingency. *Animal Learning & Behavior*, 18, 257–263.
- Morris, C. J. (1987). The operant conditioning of response variability: Free-operant versus discrete-response procedures. *Journal of the Experimental Analysis of Behavior*, 47, 273–277.
- Morris, C. J. (1989). The effects of lag value on the operant control of response variability under free-operant and discrete-response procedures. *The Psychological Record*, 39, 263–270.
- Myerson, J., & Miezin, F. M. (1980). The kinetics of choice: An operant systems analysis. *Psychological Review*, 87, 160–174.
- Neuringer, A. (1986). Can people behave “randomly”? The role of feedback. *Journal of Experimental Psychology: General*, 115, 62–75.
- Neuringer, A. (1991). Operant variability and repetition as functions of interresponse time. *Journal of Experimental Psychology: Animal Behavior Processes*, 17, 3–12.
- Page, S., & Neuringer, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes*, 11, 429–452.
- Pryor, K. W., Haag, R., & O'Reilly, J. (1969). The creative porpoise: Training for novel behavior. *Journal of the Experimental Analysis of Behavior*, 12, 653–661.
- Schoenfeld, W. N., Harris, A. H., & Farmer, J. (1966). Conditioning response variability. *Psychological Reports*, 19, 551–557.
- Shimp, C. P. (1967). Reinforcement of least-frequent sequences of choices. *Journal of the Experimental Analysis of Behavior*, 10, 57–65.
- Stokes, P. (1995). Learned variability. *Animal Learning & Behavior*, 23, 164–176.

Received January 5, 1997
Final acceptance April 10, 1997

APPENDIX

To obtain the probability distribution of the number of different sequences predicted by random responding, let $p(N, k)$ denote the probability of the event “ k different sequences are produced during the first N trials.” This event can occur in one of two mutually exclusive ways. During the first $N - 1$ trials, the bird either produced $k - 1$ different sequences and then produced a new sequence,

or it produced k different sequences and then repeated one of them on the last trial. The probabilities of these two events are given by the right side of the following equation:

$$P(N,k) = P(N-1,k-1) \frac{M-(k-1)}{M} + P(N-1,k) \frac{k}{M},$$

where M is the total number of sequences

(256 in the experiments reported here) and $N \geq 1$. The boundary conditions of the preceding equation are

$$P(1,1) = 1$$

$$P(N,k) = 0 \quad \text{for } k = 0 \text{ or } k > N.$$

Once the probabilities $p(N,k)$ are computed recursively, the average and the 95% confidence interval of the number of different sequences can be obtained (see bottom panel of Figure 1).