

DIRECTIONAL SELECTION OF RESPONSE NUMEROSITY: AN EXPLORATORY STUDY

SELEÇÃO DIRECIONAL DE NUMEROSIDADE: UM ESTUDO EXPLORATÓRIO

ARMANDO MACHADO, ANDREIA COSTA, AND SUSANA MAIA

UNIVERSIDADE DE MINHO, PORTUGAL

RESUMO

O presente estudo analisa a diferenciação numérica de padrões de resposta. Em uma caixa de Skinner com duas teclas, sete pombos receberam comida após bicarem pelo menos N vezes na tecla esquerda e depois uma vez na tecla direita (programa “*Fixed Consecutive Number*” de Mechner). Em cada ensaio, o parâmetro N era ajustado por um programa de reforço percentil (uma forma de *shaping* automático). O estudo teve dois objetivos. Primeiro, determinar como é que varia o tamanho das corridas na tecla da esquerda durante o procedimento de modelagem (*shaping*) e durante uma fase de extinção que se seguiu. Segundo, comparar os dados obtidos com as previsões de um modelo teórico de diferenciação da resposta. Os resultados mostraram que, durante a modelagem, o tamanho das corridas na tecla esquerda aumentou e depois, para alguns pombos, estabilizou, enquanto para outros pombos permaneceu variável. Alguns pombos pararam de responder quando o tamanho médio da corrida atingiu valores elevados. Observaram-se ainda variações sistemáticas nos tamanhos das corridas no interior de cada sessão como, por exemplo, o aumento do tamanho da corrida ao longo da sessão. Durante a fase de extinção os pombos produziram distribuições de tamanhos de corrida semelhantes às distribuições produzidas durante as últimas sessões de modelagem com exceção, em alguns sujeitos, do elevado número de corridas de tamanho zero. Estes resultados são interpretados à luz do modelo teórico de diferenciação numérica das respostas.

Palavras-chave: Modelo matemático, numerosidade, esquema percentil, modelagem, pombo

ABSTRACT

The experiment examined how pigeons differentiate response patterns along the dimension of number. Seven pigeons received food after pecking the left key at least N times and then switching to the right key (Mechner's Fixed Consecutive Number schedule). Parameter N was set according to a percentile schedule, which is a form of automatic shaping. Our aim was twofold: on the empirical side to determine how run length on the left key would evolve under this shaping procedure and how it would change during a subsequent extinction phase; and on the theoretical side to compare the data with the predictions of a theoretical model of response differentiation. Results showed that during shaping, run length on the left key increased and then, for some pigeons, it stabilized, whereas for others pigeons it remained variable. Some pigeons ceased to respond when average run length reached a high value. There were substantial within-session trends in run length. In extinction, before the pigeons ceased to respond altogether, they emitted the same distribution of run lengths as during the last sessions of shaping with the exception, in some birds, of a large number of runs of length zero. These results are interpreted at the light of the theoretical model of numerosity differentiation.

Keywords: Mathematical Model, Response Numerosity, Percentile Schedule, Shaping, Pigeon

Number is a dimension of stimulus control and response differentiation. A stimulus presented a given number of times may control an operant response (Alsop & Honig, 1991; Emmerton, 2001; Keen & Machado, 1999) and, similarly, animals can learn to emit a response a given number of times (Laties, 1972; Hobson & Newman, 1981; Mechner, 1958;

Platt & Johnson, 1971). The present study is concerned with the latter ability. Suppose that a pigeon faces two response keys, one on the left (L) and one on the right (R), and that it can get food by pecking the L key at least N times followed by one peck on the R key. Thus, if $N=4$, a run of five pecks on the L key (i.e., LLLLLR) would be reinforced but a run of two (i.e., LLR)

would not. Under this reinforcement rule, known as Mechner's Fixed Consecutive Number (FCN) schedule with parameter N , a pigeon's typical performance at the steady state is a Gaussian distribution of run lengths with mean and standard deviation approximately linear functions of N . These linear functions imply that, as N increases, the ratio of the standard deviation of run length to the mean run length (i.e., the coefficient of variation), is either constant or approaches a constant. A constant coefficient of variation expresses Weber's law in numerosity differentiation (see Gallistel, 1990, and Machado & Rodrigues, 2007).

These psychophysical properties describe performance after the pigeon has learned the task, that is, at the steady state. But how is such performance acquired? More generally, how is run length on the L key differentiated? Machado and Rodrigues (2007) suggested the following learning model: On each trial, the pigeon produces a run whose length, x , is sampled from a Gaussian distribution with mean μ and standard deviation proportional to μ . Having produced a run of length x , the pigeon then experiences an outcome, food (reinforcement) if $x \geq N$ or no food (extinction) if $x < N$. The model states that on the basis of the trial outcome and the sampled value, μ will change. Thus learning is conceived of as a change in mean run length. The specific learning rule is as follows:

1) If the trial ends with food and x is greater than μ , then μ increases. It seems plausible that if a run longer than the average is reinforced, the average run length will increase. The increment, $\Delta\mu$, is a linear function of μ with negative slope, $\Delta\mu = \alpha - \beta\mu$, for some positive constants α and β . This means that when μ is large the change in μ is less than

when μ is small, and that there is a maximum mean run length the pigeon can sustain, which equals $\mu_{\max} = \alpha/\beta$. This maximum may depend on the motivational conditions of the animal, the quantity of food per reinforcement, the force required to peck the keys, and some other variables, but for the present purposes it is assumed constant.

2) If the trial ends with food but x is less than μ , then μ decreases by a constant amount (i.e., $\Delta\mu = -\gamma$). Again, it seems plausible that average run length decreases (i.e., $\Delta\mu < 0$) when the reinforced run length is less than the average run length. Moreover, it is also assumed that reinforcement is more effective at decreasing than increasing μ . Mathematically, this is equivalent to the condition $\gamma \geq \alpha$, which means that, for the same value of μ , the absolute value of a decrement in μ is greater than the value of an increment in μ .

3) If the trial ends without food, then μ does not change ($\Delta\mu = 0$). This assumption was made because its alternatives seemed less plausible. Thus, an increase in μ seemed implausible because it seemed to violate the very concept of extinction. A decrease in μ although plausible at first sight would make it difficult to explain how response differentiation takes place. To illustrate, when an experimenter changes the reinforcement requirement from $N=4$ to $N=8$ in a traditional FCN schedule, there is substantial extinction immediately following the change, but the animal eventually learns the new criterion and emits significantly longer runs (see Machado & Rodrigues, 2007, or Mechner, 1958, for examples). It is unclear how such learning could take place if extinction decreased μ . The assumption that $\Delta\mu = 0$ in extinction is not incompatible with the obvious prediction that in extinction the pigeons will

stop pecking the keys. What it says is that while they peck and emit runs, average run length will not change.

Based on the model's three assumptions, we can divide the session into three sets of trials. The first set (cf. Assumption 1) includes the trials in which the current run length was greater than the criterion, N , and greater than the current value of μ . These trials ended in food and, according to the model, μ increased. The second set (cf. Assumption 2) includes the trials in which the current run length was greater than N , but less than μ . These trials ended in food and μ decreased. The third set (cf. Assumption 3) included the trials in which the current run length was less than the criterion N ; the trial ended without food and μ did not change. Note that whereas the trial outcome depends on the comparison between x and N , learning (i.e., the changes in μ) depends on the trial outcome and on how x compares with μ .

The goal of the present study was to develop a procedure to examine in greater detail the foregoing assumptions. Because two of them stress the relation between the run length sampled on each trial, x , and the internal variable, μ , the procedure should allow the experimenter to control that relation. One such procedure is based on the percentile schedule developed by Platt to study shaping (see Platt, 1973) and later adapted by Galbicka to study numerosity differentiation in the rat (Galbicka, Fowler, & Ritch, 1991; Galbicka, Kautz, & Jagers, 1993; see also Galbicka, 1994). Below, we describe how we changed the percentile schedule to examine the assumptions of Machado and Rodrigues' (2007) model of numerosity differentiation and then report some data obtained with the new procedure.

A percentile schedule is a form of automatic shaping in which the criterion for a correct response changes with the subject's recent behavior. When the subject's behavior moves, as it were, along the shaping dimension, the criterion also moves. In Mechner's FCN task under consideration this means that when run lengths increase across trials, the criterion for a correct run length, N , also increases, and, conversely, when run lengths decrease across trials, the criterion also decreases. In contrast, in Mechner's original FCN schedule, the criterion N remains constant and independent of the subject's behavior. The percentile schedule gets its name from the fact that the criterion corresponds to a percentile of the distribution of the subject's recent run lengths. Therefore, although the absolute value of the criterion may change across trials, its relative value remains constant.

The key feature of the percentile schedule is that by adjusting the criterion N on each trial based on the subject's last run lengths, the experimenter controls directly the probability of a correct response. In fact, that probability is the complement of the percentile. Thus, if the criterion is set to percentile $(1-w) \times 100$, with $0 < w < 1$, then the probability of a correct response equals w .

To illustrate the foregoing ideas and further explain how the percentile schedule works, assume the experimenter has set the response criterion at the 60th percentile of the distribution of the last $m=19$ run lengths. Hence, $1-w=0.6$ and $w=0.4$. Then, to decide whether the next run of length x is correct, and therefore eligible for reinforcement, the computer would first rank order the previous 19 run lengths from lowest to highest and then find the rank of the current run length, x . That

rank can range from 1, if x is less than all 19 run lengths, to $m+1=20$, if x is greater than all 19 run lengths. If the successive runs are independent, then all ranks are equally likely and each has probability $1/(m+1)$. The probability that the current run length will be among the first (lower) k ranks equals $k/(m+1)$ and therefore the probability that it will be among the remaining (higher) ranks equals $1 - k/(m+1)$. Setting the latter probability equal to w and solving for k yields the key equation for percentile schedules (see Platt, 1973):

$$k=(m+1) \times (1-w) \quad (1)$$

If the reinforcement rule states that for the current run to be considered correct its length must exceed k of the last m run lengths, with k given by Equation 1, then two goals will be achieved. First, the probability of a correct response will equal w , and, second, only the relatively longer runs will be eligible for reinforcement. For $m=19$ and $w=0.4$, Equation 1 yields $k=12$ and therefore, according to the rule, the current run would be correct provided x exceeded 12 or more of the last 19 run lengths. (The slight complications introduced by fractional k s or by ties in run length are addressed in the Procedure section.)

To see how the percentile schedule allows the experimenter to control the relation between the run length sampled on each trial, x , and the internal variable, μ , and thereby examine Machado and Rodrigues' (2007) model, consider that the last m runs are a sample of the population from which x is obtained. Therefore, the sample percentiles estimate the corresponding population percentiles. In particular, when the population follows a Gaussian distribution, as is typically the case with run

length (see Laties, 1972; Machado & Rodrigues, 2007; Mechner, 1958; Platt & Johnson, 1971), the 50th sample percentile estimates the mean of the distribution, μ .

If the experimenter sets the criterion of a correct response at, say, the 60th percentile of the last $m=19$ runs, then most run lengths that are reinforced (i.e., x is greater than the current criterion) will also be greater than μ . In fact, we show in the Appendix that when $m=19$, more than 97 percent of the reinforced responses will be greater than μ . In mathematical terminology, the probability that $x > \mu$ given that $x > N$, where N equals the 60th percentile of the last 19 run lengths, is greater than .97. More generally, if $m=19$ and $w=0.4$, then 60% of the trials (i.e., $1-w$) will be extinction trials, about 39% of the trials (i.e., $w \times 0.97$) will be reinforced trials with $x > \mu$ and the remaining 1% of the trials will be reinforced trials with $x < \mu$.

As the preceding example shows, the percentile schedule allows the experimenter to control the proportion of trials on which each of the three model assumptions applies. Thus, the percentile value $1-w$ sets the percentage of trials on which the response will be below the criterion (i.e., extinguished) and therefore on which Assumption 3 applies. The complement of the percentile, w , determines the percentage of trials on which the response will be above the criterion (i.e., reinforced) and therefore on which Assumptions 1 or 2 apply. We show in the Appendix that w and m determine also the proportion of reinforced trials on which the response will be greater than μ , and therefore on which Assumption 1 applies. For the w and m values used in the present experiment, that percentage was always greater than 97. Thus, by setting the two schedule parameters, the

percentile ($1-w$) and the sample size parameter (m), the experimenter sets the proportion of trials on which the model assumptions hold and thereby, provided the assumptions are correct, determines how response differentiation occurs. The experiment reported below examined these assumptions.

The experiment was divided into two phases. During the first phase, run length was shaped toward higher values in order to examine the model's first assumption. As the pigeon varied its run lengths, the percentile schedule reinforced the subset of those runs with the greatest lengths and extinguished the subset with the shortest lengths. We expected that mean run length would increase until it reached its maximum sustainable value under the circumstances and then stabilize at that value. During the second phase, reinforcers were discontinued to examine the model's third assumption, that μ would not change in extinction.

METHOD

Subjects

Seven pigeons (*Columba livia*) were maintained at 80 per cent of their free-feeding body weight with water and grit continuously available in their home cages. A 12:12 h light-dark cycle was in effect with the lights on at 8:00 am. Four pigeons (P569, P685, P320, and P873) had experience with variable interval and fixed interval reinforcement schedules. The other three birds (P199, P79, and P18) were experimentally naïve.

Apparatus

The pigeons were studied in three identical operant chambers. Each chamber was 34 cm wide, 30.5 cm long and 34 cm high. The walls and ceiling were made of aluminum

and the floor was wire mesh. The front panel was equipped with three circular keys, 2.5 cm in diameter, centered on the wall, 23 cm above the floor, and 14 cm apart, center to center. Only the left and right keys were used in the experiment. The key on the right could be illuminated with yellow light and the key on the left could be illuminated with green light. On the back wall of the chamber, 3.5 cm below the ceiling, a houselight provided general illumination. Reinforcement consisted of mixed grain delivered through a food hopper. The hopper was accessible through a 6x5 cm opening that was centered on the intelligence panel 8.5 cm above the floor. The operant chamber was enclosed in an outer box. On the back wall of the outer box, a ventilation fan circulated air through the chamber and helped to mask extraneous noises. A personal computer programmed in C++ controlled all experimental events and recorded the data.

Procedure

Pre-training. The birds learned to peck the keys through manual shaping. Next, they were exposed to trials in which after an inter trial interval (ITI) during which the houselight was illuminated, the two side keys also were illuminated. The reinforcement rule was as follows (see Mechner, 1958): With probability p , a reinforcer was delivered after the N th peck on the left key (Fixed Ratio N schedule); with probability $1-p$, a reinforcer was delivered after one peck on the right key provided that peck was preceded by at least N pecks on the left key (FCN). Notice that in both cases the bird had to peck the left key a minimum of N times. If the bird pecked the right key before it completed N pecks on the left key, the trial was cancelled and a 5-s timeout occurred

followed by a new ITI. Across five sessions, parameter p decreased from 1 (100% FR trials) to 0 (100% FCN trials) whereas parameter N always equaled 4.

During the first sessions the reinforcement duration was adjusted for each bird in order to minimize extra-session feeding. Final values ranged from 2 to 4 seconds. During reinforcement, the houselight and the key lights were turned off and the hopper light was illuminated.

Experiment proper. At the beginning of each trial, a percentile reinforcement schedule recalculated the criterion for a correct response. The criterion was set to a fixed percentile (e.g., 60th) of the distribution of run lengths produced on the previous m trials, where m was a schedule parameter. After the bird emitted a run of, say, $x \geq 0$ pecks on the left key followed by one peck on the right key, reinforcement or a timeout followed depending on whether x exceeded the criterion or not. To exceed the criterion the current run length had to be greater than k of the last m run lengths, with k given by Equation 1.

To implement the percentile schedule, three difficulties must be eliminated, a) how to handle ties in run length, b) how to deal with a fractional value of k , and c) how to compute the criterion during the first m trials of each session. When the response attribute is discrete, as is the case with run length, it may happen that because of ties the rank of a score is not well defined. To illustrate, consider $m=3$ run lengths, 10, 11, and 12; the rank of the next run length can range from 1 (if the new run length equals e.g. 5) to 4 (if the new run length equals e.g. 15). However, if the new run length equals 11, its rank is not defined precisely. All we can say is that its rank is greater

than 1 and less than 4. One way to solve the problem of ties (see also Platt, 1973) is to add to each run length a random number uniformly distributed between 0 and 1. This random component changes run length from discrete to continuous but does not affect the workings of the percentile schedule.

As for fractional values of k , consider a memory of size $m=20$ and the 60th percentile. In this case, $k=12.6$, which means that to be correct the current run length x must exceed 12.6 of the last 20 run lengths. The problem is that the number of run lengths exceeded by x is always an integer. To solve the problem, we consider three cases: if x is greater than at least 13 of the last 20 run lengths, then the response is always correct; if x is greater than at most 11 of the last 20 run lengths, then the response is always incorrect; finally, if x is greater than exactly 12 of the last 20 run lengths, then the response is correct with probability w (see Platt, 1973). This algorithm allows the experimenter to use any combination of the two schedule parameters, the percentile value $(1-w)$ and the memory size, m .

The final difficulty is how to determine the criterion during the first m trials of each session (note that the decision rule based on Equation 1 requires m previous run lengths). There are several ways to overcome this difficulty, including the following four: a) Use the last m runs of the preceding session, as if the pigeons were exposed to one uninterrupted session. This solution makes the reinforcement conditions at the beginning of each session similar to those at the end of the previous session; b) use the first m runs of the preceding session making the reinforcement conditions at the beginning of each session similar to those at the beginning of the previous session; c) pre-

sume that before the first trial there were m run lengths, all equal to a specific value (e.g., 4); or d) let m increase with each additional trial or, more specifically, reinforce the first run and then until the m^{th} trial, base Equation 1 on all previous runs. Initially, we used solution a), but then, for reasons explained below, we used a mixture of solutions b) and d).

The experiment was divided into two phases, a Reinforcement Phase and an Extinction Phase. Initially, the ITI was 10 s long, correct responses were reinforced and incorrect responses led to a 3 s timeout. The criterion matched the 60th percentile of the last $m=20$ runs. During the first trials of each session, the computation of the criterion was based on the last 20 runs of the preceding session. To be more specific, on the first trial, the criterion was based on the last 20 runs of the previous session; on the second trial, the criterion was based on the last 19 runs of the preceding session plus the run emitted on the first trial of the current session; on the third trial, the criterion was based on the last 18 runs of the preceding session plus the two runs emitted on the current session. After the 20th trial, the criterion was based exclusively on the last 20 runs of the current session.

Each trial ended when a peck on the right key completed a run and a reinforcer or a timeout was delivered, or when 180 s elapsed without a peck on the right key and a timeout was delivered. We refer to trials without a completed run as empty trials. Each session ended after 40 reinforcers had been collected, 100 trials had been completed, or the pigeon had ceased to respond for at least 3 consecutive empty trials (9 minutes), whichever occurred first.

Only the four experienced pigeons were used with the preceding conditions. After 20

sessions, the three naïve pigeons were added to the experiment and three procedural details were changed to try to optimize shaping (see results of these first sessions below). First, the ITI was eliminated and the timeout duration increased to 5 s. Thus the next trial started immediately after food or after a 5-s timeout. This change was made to increase the overall reinforcement rate. Second, runs of zero length ceased to be reinforced even when they were considered criterional according to the percentile reinforcement rule (i.e., before, if the last 20 run lengths equaled 0, a run with length 0 was reinforced on the next trial with probability $w=0.4$). And third, the run length distribution used at the beginning of each session to compute the criterion was set to the *first* 8 (as opposed to the *last* 20) runs of the previous session. Then, with each additional trial, the distribution increased by one until it reached the final valued of 20. This change was motivated by the fact that for some pigeons the run lengths at the end of each session were much higher than the run lengths at the beginning of the next session. This difference caused a relatively large number of extinction trials at the beginning of a session and it meant that the sample of run lengths produced at the end of a session did not represent well the distribution of run lengths at the beginning of the next session. We reasoned that, because of warm-up effects or distinctive stimulus properties, the beginning of each session might be more similar to the beginning than to the end of the preceding session—hence, the change from a criterion based on the last to a criterion based on the first trials of the previous session. Additionally, by setting the

initial sample size equal to 8, and not 20, we intended to give more weight to the current run lengths even during the first trials.

Table 1 shows the number of sessions and the percentile value used for each pigeon during the Reinforcement and Extinction Phases. Initially, the 60th percentile was applied to all pigeons, but then for two naïve pigeons the percentile changed to 75 because run length either did not increase (P79) or increased but then decreased again (P199). For pigeon P79, after run length started to increase with the 75th percentile, the 60th percentile was reinstated. For pigeon P199, the 75th percentile was maintained until the end of the Reinforcement Phase.

During the Extinction Phase, which lasted five sessions, after 1 to 5 reinforcers were collected, all remaining trials were in extinction. Specifically, in Sessions 1 and 2, all trials after the 5th reinforcer was collected ended in extinction; in Sessions 3 and 4, extinction started after 3 reinforcers had been collected; and in Session 5, extinction started after one reinforcer was collected. The sessions ended after 100 trials had been completed or at least 3 consecutive empty trials (9 minutes) had elapsed, whichever occurred first.

RESULTS

Figure 1 shows the percentage of reinforced runs on each session of the Reinforcement Phase. Sessions with fewer than 25 completed runs were excluded from the analysis but are represented in the figure by a missing data point (see vertical lines on the x-axis, e.g., P79). Consider the experienced pigeons in the top four panels. Because the 60th percentile was used, we expected that 40% of the runs would be reinforced (horizontal lines). For pigeons P873 and P569 the percentage of reinforced runs was close to 40, but for pigeons P320 and P685 that percentage was significantly above 40, particularly during the last sessions. A similar pattern holds for the naïve pigeons (three bottom panels): For pigeons P199 and P18 the observed and expected values were reasonably close, but for pigeon P79 the obtained value was clearly above 40 during the last sessions. We also note several missing data points during the last session of P18 and especially of P79, which means that in several sessions these pigeons stopped responding before they had completed 25 runs.

These results indicate that the main goal of the percentile schedule – to control the

Table 1

For each pigeon, number of sessions and, in parentheses, the value of the percentile holding during those sessions (e.g., P60 = 60th percentile). For the Experienced pigeons, some procedural details changed after 20 sessions (hence the asterisk in P*60).

Pigeon	History	Reinforcement Phase	Extinction Phase
P873	Experienced	20 (P*60) → 77 (P60)	5
P320	Experienced	20 (P*60) → 79 (P60)	5
P569	Experienced	20 (P*60) → 72 (P60)	5
P685	Experienced	20 (P*60) → 77 (P60)	5
P199	Naïve	52 (P60) → 33 (P75)	5
P79	Naïve	17 (P60) → 17 (P75) → 37 (P60)	5
P18	Naïve	71 (P60)	5

DIRECTIONAL SELECTION OF NUMEROSITY

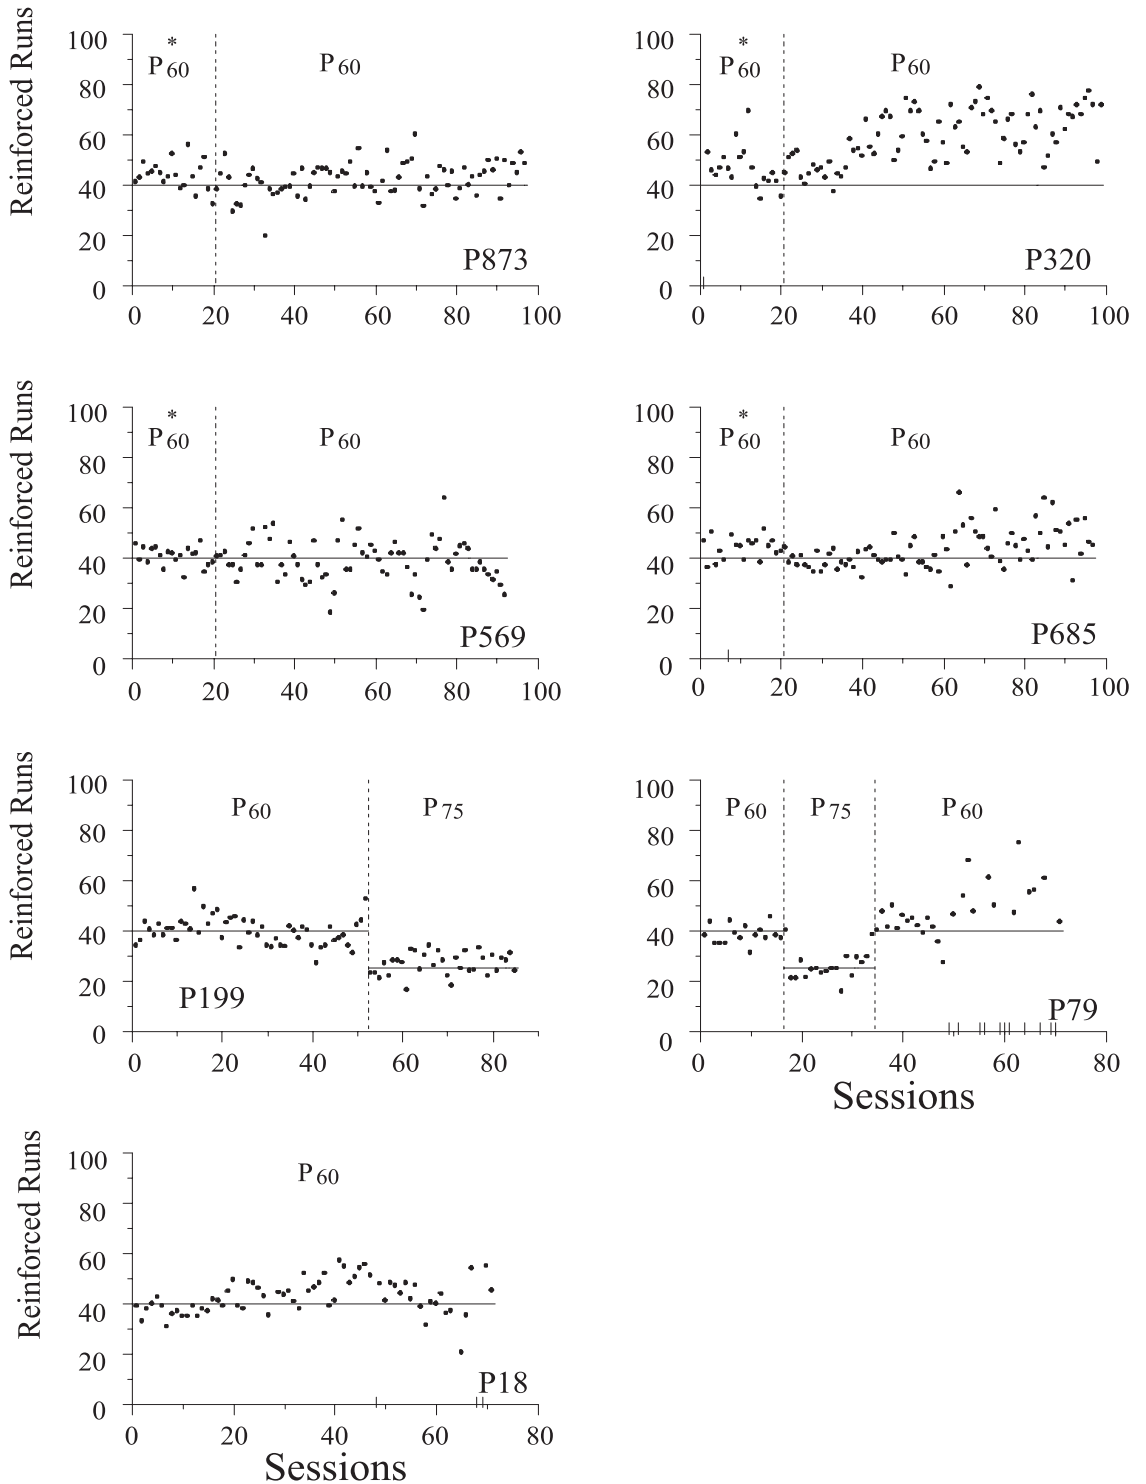


Figure 1. The points show the percentage of reinforced runs of each session of the Reinforcement Phase. The horizontal lines indicate the percentage expected according to the percentile schedule. The vertical lines indicate the moments when procedural changes took place. The top four panels show the results for the experienced pigeons and the three bottom panels show the results for the naïve pigeons. Sessions with fewer than 25 runs are represented by missing data points and by vertical lines above the x-axis.

probability of a correct response – was only partly met. If that probability had always equaled the predicted value, w , then the symbols in Figure 1 would not have deviated systematically from the horizontal lines. For reasons described below, some birds, particularly at the end of the Reinforcement Phase, were able to “beat” the percentile and earn more than the predicted percentage of reinforcers per session.

Figure 2 shows two curves for each pigeon, the average run length (top) and the percentage of runs of length zero (bottom) for each session of the Reinforcement Phase. For the experienced pigeons (top row) average run length increased during the first sessions and then it decreased. After session 20 – when some procedural changes took place —, average run

length increased for pigeon P320, but for the other three pigeons it increased and then decreased again. In some cases average run length decreased gradually (e.g., last sessions of P873 and P569), but in other cases it decreased and increased abruptly (e.g., P569 session 50). The bottom curves in each panel help to explain these abrupt changes. On several sessions, runs of length 0 dominated, accounting for as many as 70% of all runs.

The bottom row shows the corresponding data for the naïve pigeons. Either with the 60th (P18) or the 75th percentiles (P199 and P79), average run length increased across sessions and stabilized at a relatively high value. The percentage of runs with length zero never reached the same percentages as observed with the experienced pigeons.

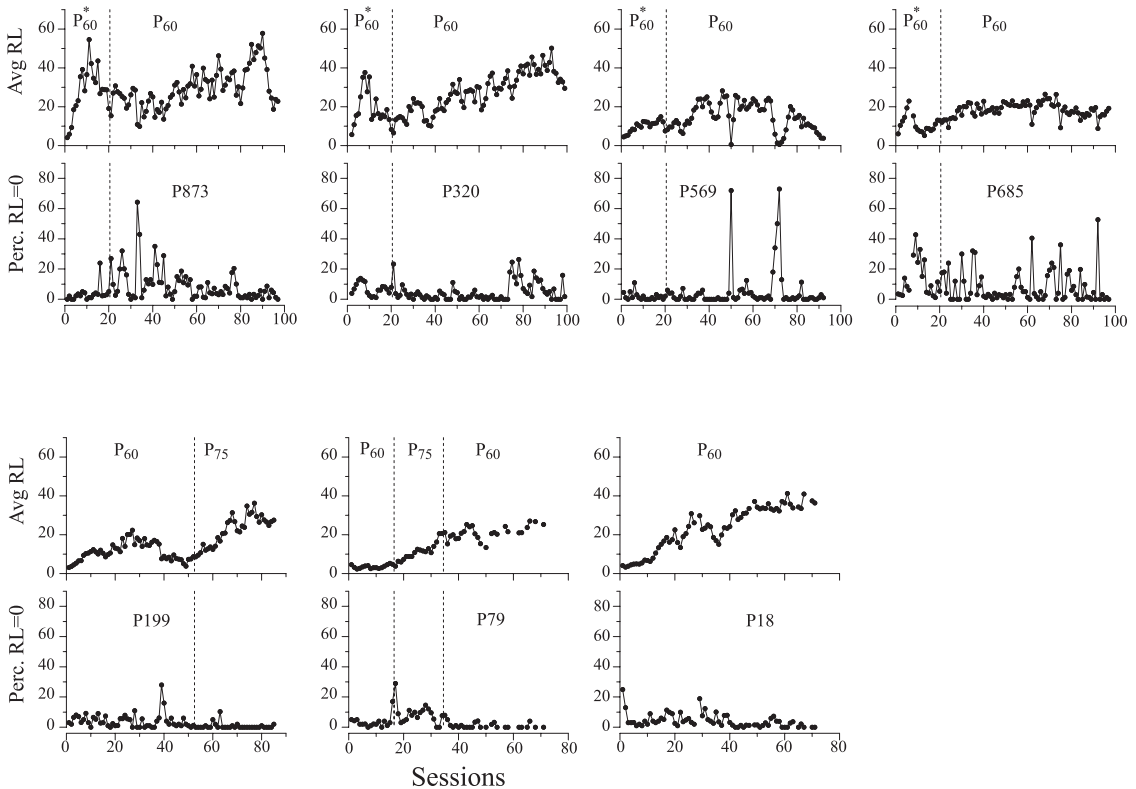


Figure 2. In each panel, the top and bottom curves show the average run length and the percentage of runs of length zero for each session of the Reinforcement Phase. The vertical lines indicate the moments when procedural changes took place. The upper and bottom rows show the results for the experienced and naïve pigeons, respectively.

During the last sessions of the Reinforcement Phase, three pigeons showed signs of extinction. For example, in each of the last five sessions, pigeons P18 and P79 ceased to respond before completing 100 trials or earning all available rewards (number of completed runs ranged from 12 to 53 in P18 and from 20 to 30 in P79). Pigeon P685 never ceased to respond completely, but it produced from 3 to 9 nonconsecutive empty trials during three of the last five sessions. These results suggest that as run length increased, reinforcement rate decreased and reached a value that, on some sessions, did not sustain key pecking for these birds. We will return to these signs of extinction in the Discussion.

Figure 3 shows within-session performance averaged over the last 10 sessions of the Reinforcement Phase. Specifically, the first data point in each panel shows the average of the run lengths emitted on the first trial of the last 10 sessions; the second data point shows the average of the run lengths emitted on the second trial of these same sessions, and similarly for the remaining data points. However, a data point is shown only when there were at least five sessions with a run at the corresponding trial. The lines represent the best-fitting regression lines. The individual cases seem to fall into four patterns. Average run length increased (top four cases), decreased slightly (P569), remained constant (P18), or fluctuated (P873) during the session. The regression line accounted for a substantial percentage of the data variance in the four top cases (range: 56% — 94%) but not in the other cases (range: 0% (P18) — 12%).

Some of these within-session curves help to explain why the obtained and predicted percentages of reinforced trials sometimes differed appreciably. The clearest example is

pigeon P320. Because run length increased steadily during the session, the 20 run lengths used to set the reinforcement criterion for the next trial systematically underestimated the length of the next run. In other words, the percentile schedule assumed that the last 20 run lengths represented well the current tendency to respond and, in particular, that the sample percentiles would estimate well the population percentiles. However, this assumption requires no within-session trends in mean run length. When these trends exist the percentiles estimated from the last 20 run lengths will be systematically biased. If the trend is pronounced and upward, as in pigeon P320, the sample percentiles will underestimate the population percentiles and the pigeon will earn more than the expected percentage of reinforcers. One might say that the pigeon was able to “beat” the percentile schedule and collect significantly more reinforcers than expected. The data from pigeon P569 provide the opposite example. Because run length decreased during the session, the last 20 runs overestimated the length of the next run and, as a consequence, the pigeon earned less than 40% of the available reinforcers (see the last five sessions of Figure 1).

The within-session pattern for pigeon P873 is strongly non-linear. During the last sessions, run length increased approximately 2.5 times during the first 30 odd trials, but then it decreased precipitously and then increased again toward the end of the session. Although we do not know the source of this strong instability, we believe it may be due to the changes in reinforcement rate induced by changes in run length (see Discussion).

The next two figures compare performance during the last 5 sessions of the Reinforcement

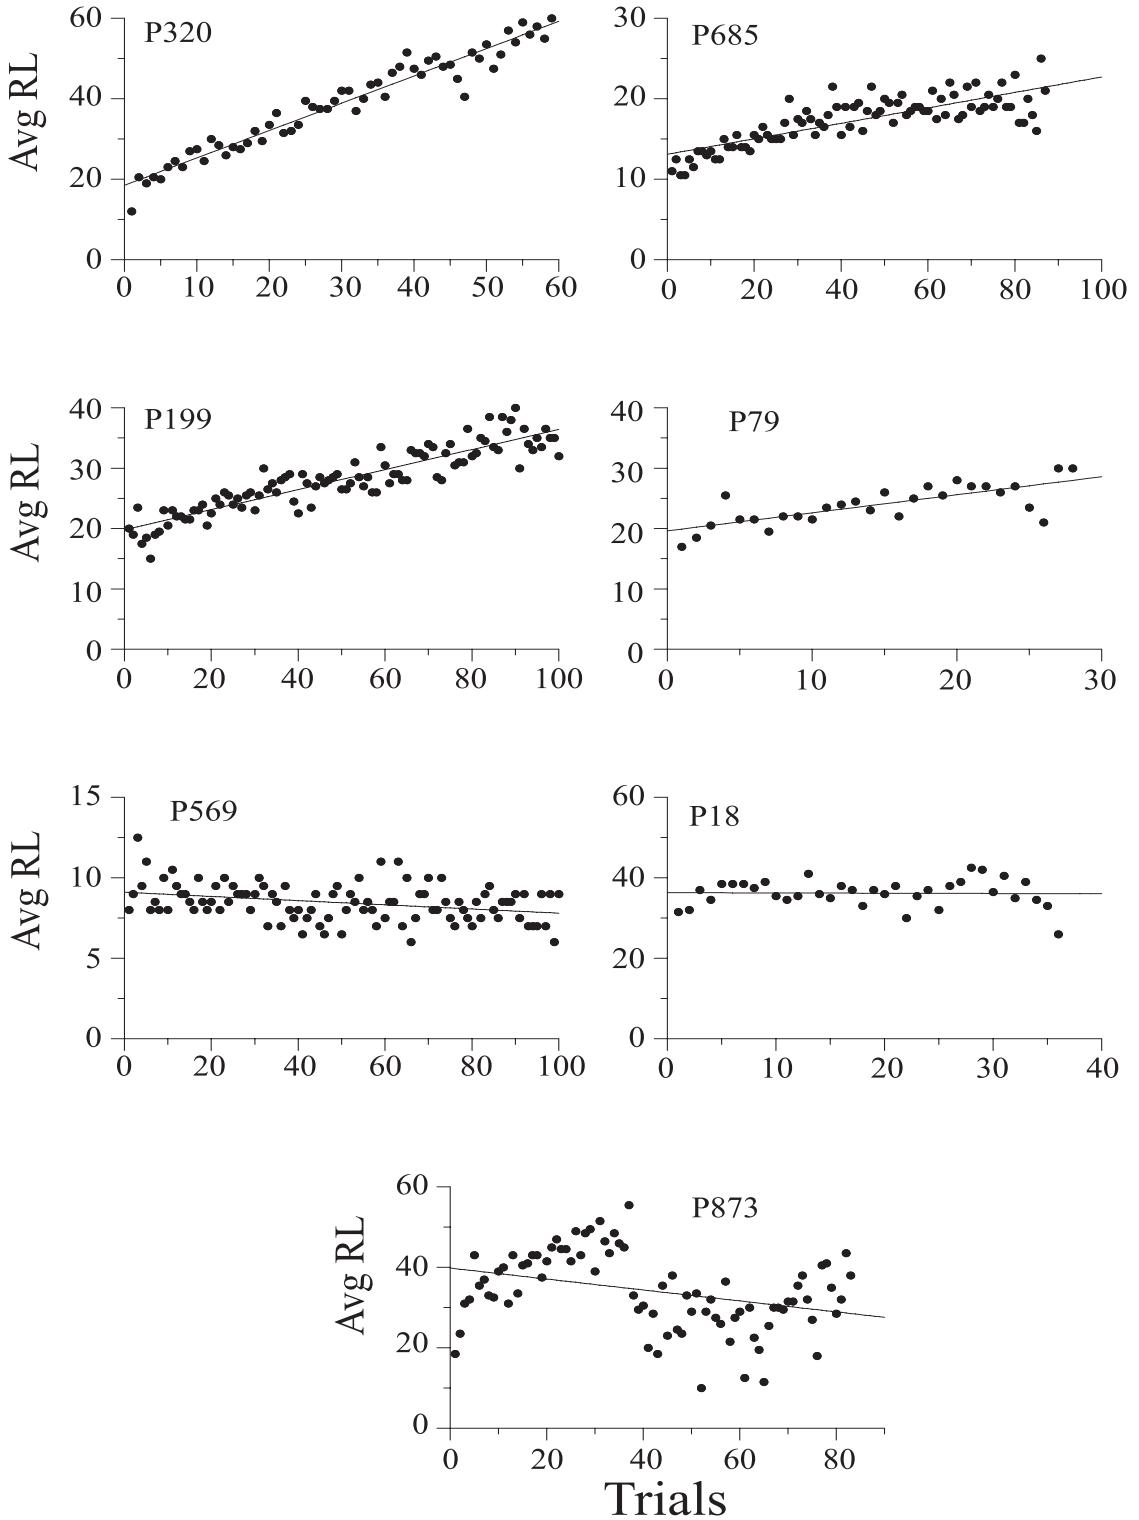


Figure 3. Within-session run lengths averaged over the last 10 sessions of the Reinforcement Phase. A data point represents the average of at least five sessions with a run at the corresponding trial. The lines are the best-fitting regression lines.

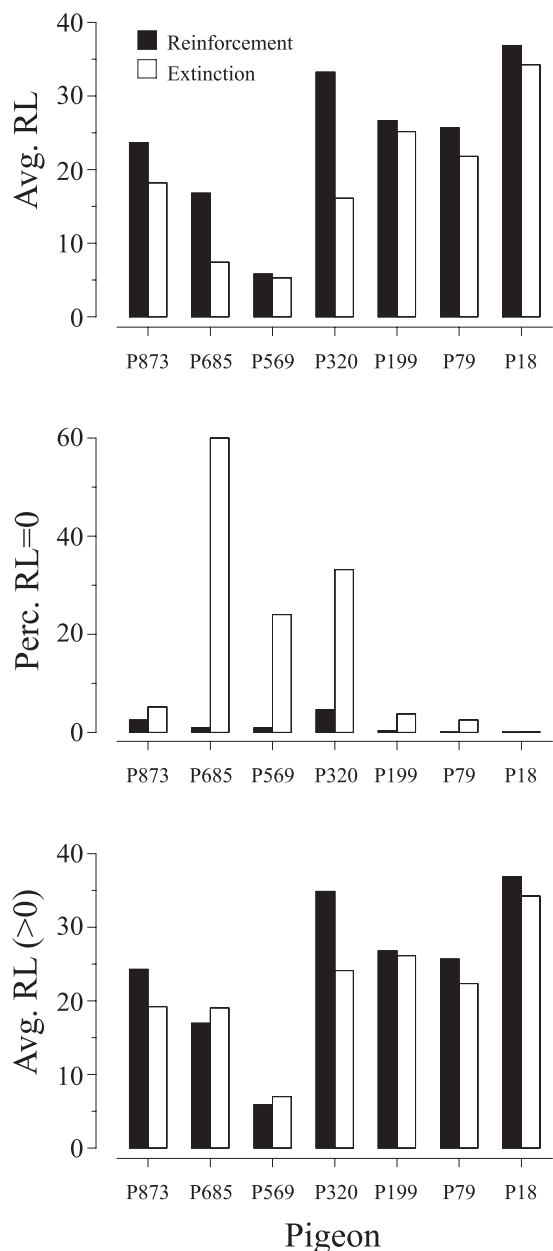


Figure 4. The bars show performance during the last five sessions of the Reinforcement Phase (black bars) and the five sessions of the Extinction Phase (white bars). The top panel shows average run length; the middle panel shows percentage of runs of length zero; and the bottom panel shows average length of runs greater than zero.

Phase with the 5 sessions of the Extinction Phase. The top panel of Figure 4 shows that in extinction average run length decreased for all birds (paired t -test, two-sided: $t(6)=2.6$,

$p<.05$). However, the magnitude of the decrement varied substantially across pigeons, from less than 10% for P569, P199, and P18, to more than 50% for P320 and P685.

The middle panel shows the percentage of runs of length zero. For three pigeons, P685, P569, and P320, this percentage increased substantially during extinction; for three others, P873, P199, and P79, it increased slightly; and for one pigeon, P18, it did not increase. If we exclude runs of length zero and recalculate the average run length in the two phases, then we get the results shown in the bottom panel of Figure 4. With the exception of pigeon P320, all changes were small. The differences in average run length were not statistically significant (paired t -test, two-sided: $t(6)=1.68$, $p=.14$).

Figure 5 shows the distributions of run lengths greater than zero during the last five sessions of the Reinforcement and Extinction Phases. There were no major and systematic differences in the variability of these distributions. In four pigeons, the standard deviation increased in extinction (P685, P569, P320, and P199) but in other three it decreased ($t(6)=0.73$, ns.) The coefficient of variation (i.e., the ratio of the standard deviation to the mean) increased for the same four pigeons but remained roughly constant for the other three ($t(6)=2.15$, $p=0.07$). In summary, the distribution of positive run lengths did not differ systematically between the last sessions of the Reinforcement Phase and the five sessions of the Extinction Phase.

However, the distribution of run lengths does not fully represent the results from the Extinction Phase because in several sessions the pigeons simply ceased to respond. In fact, no pigeon completed 100 trials during the last two extinction sessions — the average number of

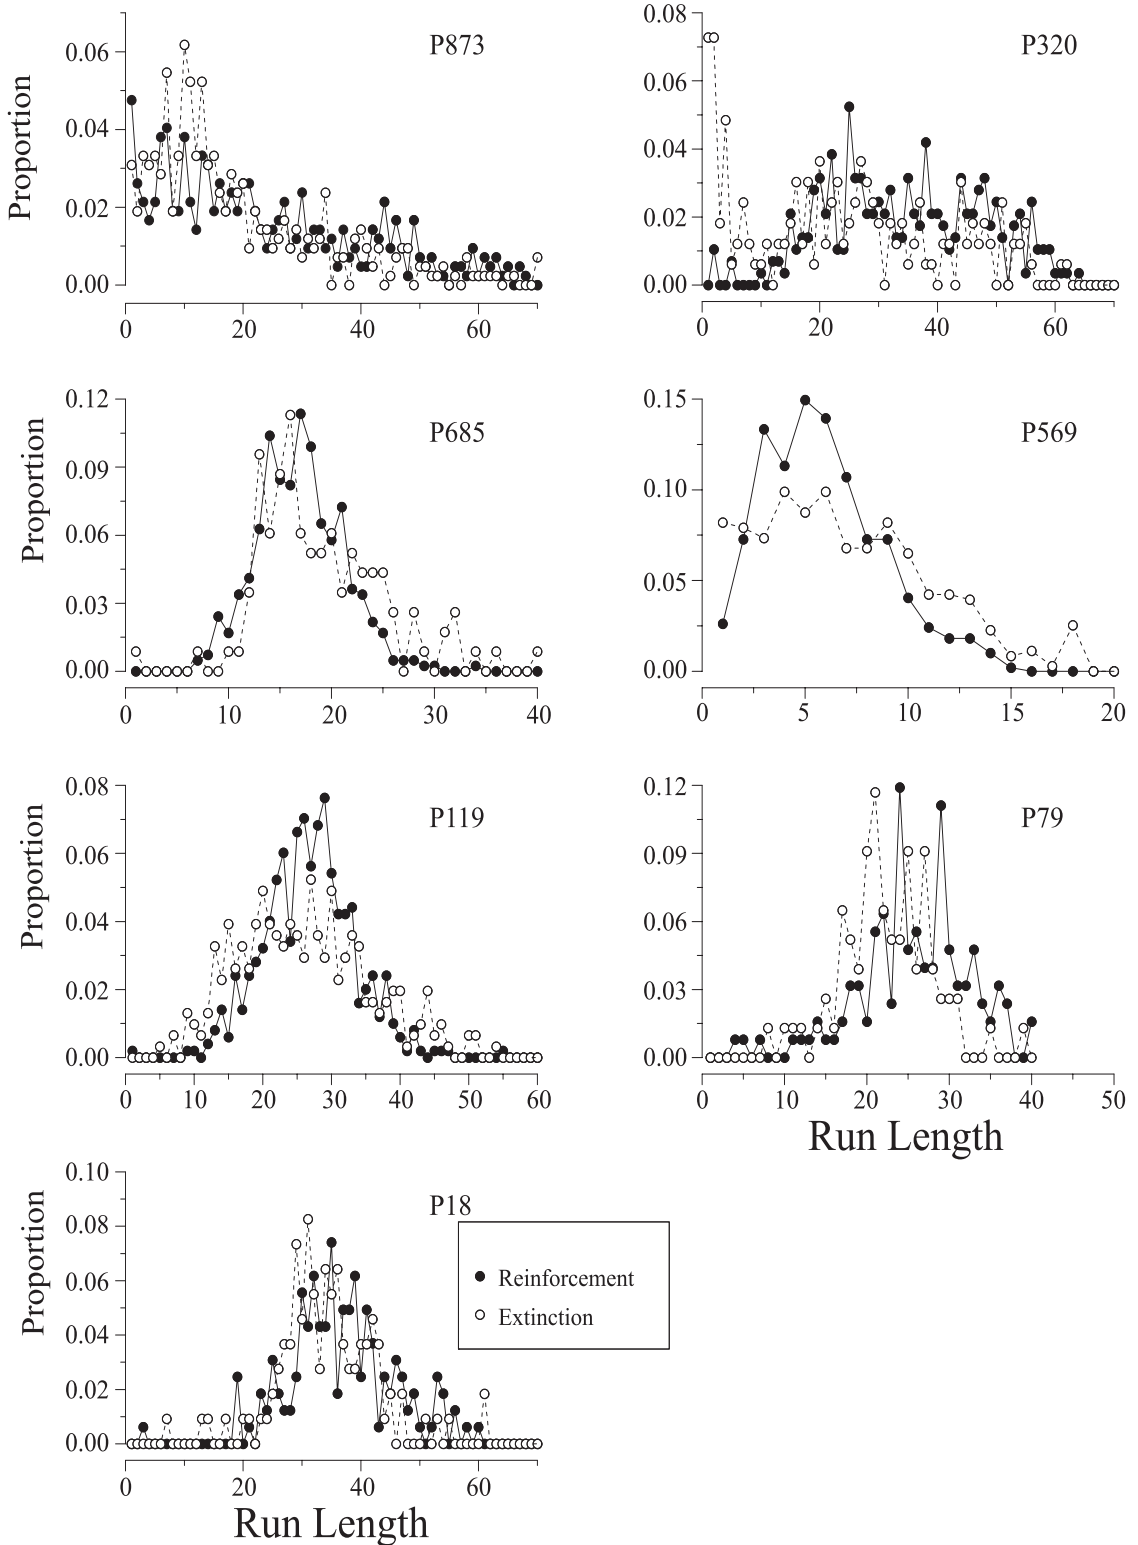


Figure 5. The points show the relative frequency of run lengths greater than zero during the last five sessions of the Reinforcement Phase and the five sessions of the Extinction Phase.

trials with a run equaled 48 in the fourth session (range: 13 — 96) and 49 in the fifth session (range: 17 — 97).

Taken together, the results support the following description of behavior in extinction. Not surprisingly, the pigeons eventually ceased to respond, but while they responded, they emitted the same behavior as during reinforcement (i.e., similar run length distributions) with the exception, in some pigeons, of a large number of runs of length zero. Although we do not know why only some pigeons emitted a significant number of runs of length zero, we note that these pigeons (also pigeon P873) had already shown a significant number of runs of length zero during some sessions of the Reinforcement Phase (see Figure 2). It may be that, at least for some birds, runs of length zero increase significantly when either reinforcement is completely removed (Extinction Phase) or long runs cause a substantial decrease in obtained reinforcement rate (Reinforcement Phase). The two cases are functionally equivalent forms of extinction.

DISCUSSION

Numerosity is a response dimension that can be shaped towards specific values. Using a FCN task and a percentile schedule, Galbicka et al. (1993) shaped run length towards the target value of 12. When the current run length was greater than 12, it was reinforced only if it was below a criterion value — shaping proceeded in the downward direction; when the current run length was less than 12, it was reinforced only if it was greater than a criterion value — shaping proceeded in the upward direction. The schedule implemented a form

of centripetal selection, shaping the mean of the distribution of run lengths towards the target value of 12.

In the present study, we attempted to shape run length towards increasingly larger values. Instead of Galbicka et al. (1993)'s centripetal selection, our procedure implemented directional selection¹. The goal was to examine two assumptions of Machado and Rodrigues' (2007) model of numerosity differentiation. The first assumption stated that if the trial ended with food and the current run length, x , was greater than the mean run length, μ , then μ increased. To examine this assumption, the procedure had to meet one condition, namely, restrict reinforcement to run lengths greater than μ . When this condition is met, the model predicts that average run length will increase steadily until it stabilizes at a maximum value. But the condition is difficult to meet because μ is not directly measurable and, in addition, it is presumed to change during training.

Percentile schedules seemed a reasonable procedure to try for several reasons. First, the criterion for a correct or reinforceable response is based on the last responses the subject has produced. These responses are used as a sample from which the population parameters (e.g., μ) can be estimated. As long as the sequential dependencies between consecutive responses are weak (see Platt, 1973, and Galbicka, 1994, for a discussion of this issue) and the population parameters do not change too quickly (where "quickly" is defined with respect to the sample size m), the sample statistics such as a specific percentile will estimate well the corresponding population parameters. Second, because the proportion of correct responses greater than μ

¹ Alternatively, one could say that in directional selection the target value of selection is at infinity.

depends on the two schedule parameters (w and m), it is under the experimenter's control. And third, by reinforcing only responses greater than the sample percentile, directional shaping is achieved under constant reinforcement probability per trial.

The results only partly supported the assumptions of the percentile schedule. Reinforcement probability per trial remained constant and reasonably close to the predicted value for some pigeons, but for others the obtained and predicted values differed significantly because these pigeons showed strong within-session trends in run length. These strong trends violate the assumption of statistical independence (or at least weak dependence) between consecutive runs and the assumption that the sample percentiles represent the current population percentiles. When average run length increases (decreases) significantly during the session, the sample percentiles underestimate (overestimate) the population percentiles and the probability of a correct response is no longer constant.

For six of the seven birds, average run length, an estimate of μ , increased across sessions, although not systematically (see Figure 2). Some acquisition curves revealed non-monotonic trends – run length increased and then decreased, sometimes abruptly and sometimes gradually. These variations were observed more in the experienced than in the naïve pigeons and they may reflect the complex effects of the animal's learning histories. With respect to Machado and Rodrigues's (2007) model, these results are mixed: For some pigeons, the acquisition curves increased and stabilized as the model predicted, but for other pigeons the acquisition curves were non-monotonic and more variable than the model

predicted. In particular, the model cannot account for the decrease in average run length when reinforcement is following the relatively longer runs (e.g., last sessions of Reinforcement Phase for P873 and P569). Clearly, the model ignored some of the processes underlying the pigeons' behavior. The nature of some of these processes is explored below.

The third assumption of Machado and Rodrigues' (2007) model stated that μ did not change during extinction and the data were somewhat more consistent with this assumption. In fact, provided we exclude runs of length zero, the distribution of run length did not change consistently during extinction. That is, all pigeons eventually stopped responding, but while they responded their runs with positive length had the same distribution as during reinforcement; in particular there was no evidence that run length decreased *gradually*.

Other researchers obtained similar results concerning the changes in average run length from reinforcement to extinction. Using rats in FCN 5 and FCN 10 schedules, Platt and Day (1979) found a slight increase in average run length in Experiment 2 but a slight decrease in Experiment 3. That is, as in our experiment, changes in average run length from reinforcement to extinction were small in magnitude and inconsistent in direction. In contrast, results concerning the changes in run length variability are less clear. Whereas we found no systematic changes, Mechner (1958b) reported greater variability using rats in a FCN 8 schedule. Unfortunately, Mechner's article shows the data for only one rat and reports no descriptive or inferential statistics. Platt and Day (1979, Experiment 3) analyzed the changes in run length during the course of extinction and

found that average run length decreased but the standard deviation did not change consistently. When we exclude runs of length zero (which could not occur in Platt & Day's study), we find that the average run length did not change consistently or substantially during the five extinction sessions (the average across birds equaled 21.5, 21.3, 23.5, 21, 19.4). Differences in species (rats versus pigeons) and procedure (percentile versus FCN schedules) may account for these inconsistencies.

Some birds produced a large number of runs of length zero during extinction. For example, pigeon P685 completed 97 trials during the last extinction session, but of these 81 were runs of length zero. These runs were emitted on consecutive trials even though each peck on the right key was followed by a 5-s timeout. Similar behavior had occurred during the reinforcement period (see Figure 2: P685, session 92; P569, sessions 49 and 72; and P873, session 33). Closer inspection of the data revealed that zero length runs were often preceded and followed by *significantly longer runs*. In other words, zero length runs were not typically surrounded by runs of length 1, 2, or 3, for example, but by runs of length 20, 40 or even 50. The pigeon's performance on these occasions was impressive because the large number of consecutive pecks on the R key seemed quite insensitive to their immediate (and presumably negative) consequence.

The preceding results suggest that extinction may have two effects, the pigeons cease pecking the keys, or the two-link response chain learned during acquisition (i.e., peck the left key several times and then peck the right key) breaks down and the pigeons emit only the second link. Both effects were observed also during the Reinforcement Phase. For example,

pigeons P79 and P18 ceased to respond during each of the last five sessions of the Reinforcement Phase and all pigeons produced at least 20% of runs of length zero during one or more sessions of that phase.

Shaping towards greater response numerosities poses the pigeon a difficult problem: The local reinforcement contingencies promote longer runs because these runs are differentially reinforced, but the global reinforcement contingencies promote shorter runs because these runs correlate with higher reinforcement rates (see also Galbicka et al., 1993). A pigeon exclusively sensitive to the local contingencies would increase run length until its maximum sustainable value; a pigeon exclusively sensitive to the global contingencies would decrease run length to its minimum allowable value. Pigeons show sensitivity to both types of contingencies. On the one hand, the fact that they increase run length with training shows clearly they are sensitive to local contingencies of reinforcement. Similarly, the fact that when average run length is high they start pausing and even cease responding, or they start emitting runs of length zero at significantly higher levels also suggests sensitivity to local contingencies of extinction. On the other hand, the fact that average run length may decrease for several consecutive sessions and even within sessions, despite differential reinforcement of longer runs, suggests sensitivity to global contingencies. Similarly, fast increments in run length across trials (see Figure 3) suggest sensitivity to molar contingencies because by increasing run length during the session a pigeon could obtain significantly more than the scheduled reinforcers. Thus, in the last 10 sessions, pigeon P320 increased and almost tripled run length during each session and for

that reason it obtained on the average 68% of reinforced trials instead of the predicted 40% (see Figures 1 and 3).

Complex interactions between global and local processes may also be responsible for the relatively strange behavior of pigeon P873 during the last sessions of the Reinforcement Phase (see Figure 3). At the beginning of the session, run length averaged about 20 pecks. Then it increased approximately 2.5 times over the next 30 odd trials and, as a consequence, reinforcement rate decreased by the same factor. On several sessions, the bird then emitted a significant number of shorter runs including many runs of length 0, which caused the dip in Figure 3 around trial 40. These runs, which may have been elicited by the lower reinforcement rate, had two interrelated effects on subsequent trials, to decrease the criterion of a correct run and thereby to increase reinforcement rate. These reasons may explain the recurrence for several sessions of the pattern described above (i.e., run length increased during the first trials, then it decreased abruptly for the next few trials, and then it increased gradually until the end of the session), yielding the sinusoidal curve shown in Figure 3.

To conclude, the present study makes two major contributions to the study of numerosity differentiation, one methodological and the other theoretical. On the methodological side, the study extended percentile schedules to directional shaping of response numerosity and identified some of the difficulties in using these schedules (e.g., within-session trends). On the theoretical side, the study showed that Machado and Rodrigues' (2007) model captures some of the processes involved in numerosity differentiation, in particular the unchanging distribution of positive run lengths in extinction. It also showed

that to the local processes suggested by the model's three assumptions, one must add the global processes that seem to come into play when reinforcement rate changes substantially with the animal's behavior (e.g., negative correlation between run length and overall reinforcement rates, or extinction-elicited effects on the integrity of the response chain). How the local and global processes interact in numerosity differentiation remains to be investigated.

REFERENCES

- Alsop, B. & Honig, W. K. (1991). Sequential stimuli and relative numerosity discriminations in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *17*, 386-395.
- Emmerton, J. (2001). Birds' judgments of number and quantity. In: R. G. Cook (Ed.), *Avian visual cognition* [On-line]. Available: www.pigeon.psy.tufts.edu/avclemmerton/
- Galbicka, G. (1994). Shaping in the 21st century: moving percentile schedules into applied settings. *Journal of the Experimental Analysis of behavior*, *27*, 739-760.
- Galbicka, G., Fowler, K. P., & Ritch, Z. J. (1991). Control over response number by a targeted percentile schedule: reinforcement loss and the acute effects of d-amphetamine. *Journal of the Experimental Analysis of Behavior*, *56*, 205-215.
- Galbicka, G., Kautz, M. A., & Jagers, T. (1993). Response acquisition under targeted percentile schedules: a continuing quandary for molar models of operant behavior. *Journal of the Experimental Analysis of behavior*, *60*, 171-184.
- Gallistel, C. R., (1990). *The organization of learning*. Cambridge, MA: Bradford Books/MIT Press.
- Hobson, S. L., & Newman, F. (1981). Fixed-ratio counting schedules: Response and time measures considered. In M. L. Commons, & J. A. Nevin

- (Eds.), *Quantitative analysis of behavior: Vol. 1. Discriminative properties of reinforcement schedules* (pp. 193–224). Cambridge, MA: Ballinger.
- Keen, R. & Machado, A. (1999). How pigeons discriminate the relative frequency of events. *Journal of the Experimental Analysis of Behavior*, 72, 151-175.
- Laties, V. G. (1972). The modification of drug effects on behavior by external discriminative stimuli. *Journal of Pharmacology and Experimental Therapeutics*, 183, 1-13.
- Machado, A., & Rodrigues, P. (2007). The differentiation of response numerosities in the pigeon. *Journal of the Experimental Analysis of Behavior*, 88, 153–178.
- Mechner, F. (1958). Probability relations within response sequences under ratio reinforcement. *Journal of the Experimental Analysis of Behavior*, 1, 109-122.
- Mechner, F. (1958b). Sequential dependencies of the lengths of consecutive response runs. *Journal of the Experimental Analysis of Behavior*, 1, 229-233.
- Platt, J. R. (1973). Percentile reinforcement: paradigms for experimental analysis of response shaping. In G. H. Bower (Ed.), *The psychology of learning and motivation: Vol. 7. Advances in theory and research* (pp. 271-196). New York: Academic Press.
- Platt, J. R., & Johnson, D. M. (1971). Localization of position within a homogeneous behavior chain: Effects of error contingencies. *Learning & Motivation*, 2, 386-414.
- Platt, J. R., & Day, R. B. (1979). A hierarchical response-unit analysis of resistance to extinction following fixed-number and fixed-consecutive-number reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, 5, 307-320.

Received January 10, 2008

Final acceptance August 24, 2008

APPENDIX

We assume that run length is a Gaussian random variable with mean μ and standard deviation σ . Let $F(\epsilon)$ and $f(\epsilon)$ represent its distribution and density functions, respectively. The percentile schedule works as follows. From the last m run lengths, the $(1-w) \times 100$ percentile is calculated and set as the criterion, N . The next run length will be correct (hence, reinforced) provided its length exceeds N . In terms of ranks, N equals the value of the k^{th} lowest of the last m run lengths and, therefore, to be reinforced the rank of the next run length must exceed k . Below we derive the probability that the next run length, x , will be greater than μ provided that it is greater than N . In symbols, we calculate $P(x > \mu | x > N)$ knowing that x comes from a Gaussian distribution and N is the $k^{\text{th}} = (1-w)(m+1)$ lowest of the last m run lengths.

First, note that $P(x > \mu | x > N) = P(x > \mu \text{ and } x > N) / P(x > N)$. But because $P(x > N)$ equals the complement of the percentile, w , we have

$$P(x > \mu | x > N) = P(x > \mu \text{ and } x > N) / w$$

We calculate the numerator in the last expression by conditioning on the value of N .

$$\begin{aligned} &P(x > \mu \text{ and } x > N) \\ &= \int_{-\infty}^{+\infty} P(x > \mu \text{ and } x > N | N \approx \epsilon) P(N \approx \epsilon) d\epsilon \\ &= \int_{-\infty}^{\mu} P(x > \mu) P(N \approx \epsilon) d\epsilon + \int_{\mu}^{+\infty} P(x > \epsilon) P(N \approx \epsilon) d\epsilon \\ &= \frac{1}{2} \int_{-\infty}^{\mu} P(N \approx \epsilon) d\epsilon + \int_{\mu}^{+\infty} (1 - F(\epsilon)) P(N \approx \epsilon) d\epsilon \\ &= \frac{1}{2} \int_{-\infty}^{\mu} P(N \approx \epsilon) d\epsilon + \int_{\mu}^{+\infty} P(N \approx \epsilon) d\epsilon - \int_{\mu}^{+\infty} F(\epsilon) P(N \approx \epsilon) d\epsilon \end{aligned}$$

We are left with the calculation of $P(N \approx \epsilon)$, that is, the density function of the sample percentile. We denote it by $g_k(\epsilon)$ and the corresponding distribution function by $G_k(\epsilon)$. For N to be equal to ϵ , it must be the case that $k-1$ of the last m runs had lengths less than ϵ , 1 had length equal to ϵ , and the remaining $m-k$ had lengths greater than ϵ . Because the m run lengths (RL) are independent and identically distributed random variables, the probability of the preceding event equals

$$g_k(\epsilon) = \underbrace{mf(\epsilon)}_{\text{one RL} \approx \epsilon} \times \underbrace{\binom{m-1}{k-1} [F(\epsilon)]^{k-1}}_{k-1 \text{ RLs} < \epsilon} \times \underbrace{[1-F(\epsilon)]^{m-k}}_{m-k \text{ RLs} > \epsilon}$$

and the corresponding distribution function equals

$$G_k(\epsilon) = \sum_{n=0}^{m-k} \binom{m}{n} [F(\epsilon)]^n [1-F(\epsilon)]^{m-n}$$

Therefore,

$$P(x > \mu \text{ and } x > N) = 1 - \frac{1}{2} G_k(\mu) - \int_{\mu}^{+\infty} F(\epsilon) g_k(\epsilon) d\epsilon$$

With some algebra, one can re-write the preceding expression as follows

$$P(x > \mu \text{ and } x > N) = \int_{\mu}^{+\infty} f(\epsilon) G_k(\epsilon) d\epsilon$$

Finally,

$$P(x > \mu | x > N) = \left[\int_{\mu}^{+\infty} f(\epsilon) G_k(\epsilon) d\epsilon \right] / w$$

By changing to z-scores, one can show that this probability does not depend on either μ or σ . Moreover, for the values used in the experiment (i.e., $m=19$ and $w=0.4$ or 0.25), this probability was greater than .97.