

Acquisition and extinction under periodic reinforcement

Armando Machado *, Munire Cevik

Department of Psychology, Indiana University, Bloomington, IN 47405-1301, USA

Received 31 January 1998; received in revised form 21 April 1998; accepted 24 April 1998

Abstract

This study reexamined the processes of acquisition and extinction under periodic reinforcement. During the first phase of the experiment, pigeons were exposed to a fixed-interval schedule either 40 or 80 s long. During the second phase, each session started with the fixed-interval schedule but changed to extinction at an unpredictable moment. The results showed that during phase 1 the curve for the average rate of pecking along the interval rotated across sessions, that is, the rate immediately after food decreased, whereas the rate at the end of the interval increased. The initial and terminal rates approached their steady state at different speeds. During the extinction trials of phase 2, behavior was characterized by pause-peck oscillations with a period slightly longer than the fixed-interval duration. These findings concerning acquisition and extinction under periodic reinforcement were then compared with the predictions of some current theories of timing. © 1998 Elsevier Science B.V. All rights reserved.

Keywords: Periodic Reinforcement; Acquisition; Extinction; Oscillation; Pigeon; Key peck

1. Acquisition and extinction under periodic reinforcement

What do animals learn when they obtain food periodically? For example, consider a hungry pigeon that receives food whenever it pecks a key 60 s after the last food. After many trials with this procedure, known as a fixed-interval (FI) 60 s schedule, what has the pigeon learned? One strategy to answer this question is to analyze how the

bird's behavior changes from the first to the last days of training, that is, from a moment in which behavior is changing appreciably (acquisition) to a moment when behavior does not change (steady state). The behavioral trajectory followed by the bird may provide clues to the processes and products of temporal learning. Another strategy is to analyze the pigeons behavior when food is removed from the situation, that is, during extinction. For example, it is well known that after some exposure to an FI schedule, pigeons pause immediately after food and then respond at a constant, or slightly increasing, rate until the end

* Corresponding author. Tel.: +1 812 8554642; fax: +1 812 8554691; e-mail: amachado@indiana.edu

of the interval. If each food episode functions as the sole discriminative stimulus for the ensuing pause, then no pauses should be observed in the absence of food (Kello, 1972; Staddon, 1972). Similarly, if some temporal properties of responding (e.g. the acceleration of response rate once pecking starts) are controlled not only by food but by other events, then removing food may show more clearly the behavioral effects of these other events. Behavior during extinction also allows one to assess the effects of the FI training on responding during intervals much longer than the FI. Thus, what a pigeon has learned during an FI 60 s schedule may influence its behavior for an interval much longer than 60 s. The present study followed the two strategies mentioned above—the analysis of acquisition and the analysis of extinction, to answer the question of what pigeons learn under periodic reinforcement.

Basic as the preceding issues are, it may be surprising to know that most studies on temporal control performed during the last 30 years or so have not addressed them. Instead, and with few exceptions (Higa et al., 1991), they have focused on steady-state behavior and its quantitative properties, on the statics of temporal control rather than on its dynamics and transient phases. We have to go back 40 years to find studies on acquisition and extinction under periodic reinforcement. For example, after an extensive analysis of cumulative records of pigeons and rats exposed to a wide range of FI schedules, Ferster and Skinner (1957) concluded that during acquisition the profile of the cumulative record changes from inverted to normal scallops. That is, early during the FI training, response rate is high immediately after food and low before the next food (inverted scallop), but later in training the pattern is reversed, and response rate is low after food and high before the next food (scallop). Concerning extinction, Ferster and Skinner concluded that 'the most obvious interpretation [of extinction following FI training] is that the interval performance [i.e. the scallop] follows whenever the bird has been pausing for any length of time,' (p. 197).

Ferster and Skinner's analyses, however, were too unsystematic, fragmentary, and vague to yield solid generalizations. Presumably because of tech-

nological limitations, the authors have left unanswered many critical questions, some concerning acquisition, most concerning extinction. For example, if, as Ferster and Skinner stated, the response rate immediately after food decreases across sessions whereas the response rate at the end of the interval increases, then one may ask how the rates of change of these two response rates are related. The answer may help us decide whether one or more processes are at work during FI training, for if the initial and terminal rates change across sessions with similar speeds, then one single process may be involved, but if the rates change at different speeds, then two or more processes are likely. In the same vein, if during extinction the birds peck for a while, then pause, then peck for another while and pause again, as Ferster and Skinner reported, then one may ask how long these peck–pause cycles are, and how their duration relates to the FI parameter. Data on these issues are critical to characterize what animals learn during periodic reinforcement.

The study of behavior during acquisition and extinction is also important for what it can tell us about current models and theories of timing. Although, as we mentioned before, these models were initially conceived, and have since been used, mostly to deal with static conditions, they nonetheless make assumptions about, or allow predictions of, behavior during transient conditions. For example, as we show later, according to one interpretation of the scalar expectancy theory (SET), responding should decline relatively early in extinction following FI training, but according to another responding should persist for a long period of time. Similarly, the behavioral theory of time predicts that, under some conditions, response rates early and late in the FI should approach their steady state at different speeds. Hence, another goal of the present study was to compare some empirical findings concerning acquisition and extinction with the assumptions and predictions made by current models of timing.

The experiment reported below was divided into two phases, the first dealing with acquisition, the second with extinction. During phase 1 pigeons were exposed to one of two FI schedules for many sessions. In this phase, we were inter-

ested in characterizing how the average response rate profile changes from acquisition to the steady state. During phase 2, each session started with reinforced trials and ended with extinction. That is to say, somewhere in the middle of the session, and unpredictably from the pigeon's standpoint, food was removed. This extinction procedure seemed preferable to the usual alternative—one long extinction session—because it allowed repeated measurements of response rate at specific times since the last reinforcer. If, as we reasoned, the effects of reconditioning at the beginning, and reextinction at the end, of each session are not strong, at least during the first sessions of phase 2, then the new procedure should yield more reliable response rate data during the initial phases of extinction. The results from both phases were then compared with current theories of timing.

2. Method

2.1. Animals

Six pigeons (*Columbia livia*) participated in the experiment. The birds were housed in individual home cages with water and grit continuously available. Throughout the experiment the birds were maintained at 80% of their free-feeding body weight. Two birds were experimentally naive (birds 5259N and 2775N) and four other birds had been used in an experiment not related to temporal discrimination. The pigeon colony remained always illuminated (no light/dark cycle was in effect).

2.2. Apparatus

Two standard Med Associates® operant chambers for pigeons (29 × 24 × 30 cm) were used. The front and back walls of the boxes were aluminum, the side walls Plexiglas, and the floor wire mesh. The front wall contained three response keys in a linear arrangement, each 2.5 cm in diameter and 8 cm apart, center to center. The keys were centered on the front wall, 24 cm above the floor. Because the side keys were not used in the experiment they were covered with black tape. The center key

could be illuminated from behind with a red light. Directly below the center key and 4 cm from the floor was the 6 × 7 cm hopper opening. A 7.5 W white light illuminated the mixed grain when the hopper was activated. Another 7.5 W white light located on the top of the back wall provided general illumination. The operant chamber was enclosed in an outer box that contained a fan to provide air circulation and mask extraneous noises. A personal computer controlled all experimental events and recorded the data.

2.3. Procedure

All pigeons were trained to peck the key using a modified autoshaping procedure. After variable intertrial intervals ($\bar{x} = 60$ s), the key was illuminated with red light. If the pigeon did not peck the key for 6 s, the keylight was turned off and the hopper was raised and illuminated for 3 s. A peck at the illuminated key turned the key off and produced food immediately. The houselight was always illuminated except during food episodes. Each session ended after 50 trials. Autoshaping lasted two sessions for the experienced birds and four or five sessions for the naive birds.

2.3.1. Phase 1

The pigeons were then assigned randomly to one of two groups with the restriction that each group contained one naive bird. One group was trained on an FI 40 s schedule and the other on an FI 80 s schedule. At the beginning of each trial, the houselight was turned on and the key was illuminated. The first peck after the fixed interval elapsed turned off the key and house lights and raised the food hopper for 3 s. After the first five sessions, the hopper duration was adjusted for each bird to minimize extra-session feeding. The final values ranged from 2 to 5 s. Sessions ended after 50 trials and the number of sessions ranged from 55 to 60.

2.3.2. Phase 2

Following the FI training, all birds received 15 sessions with the new extinction procedure. Each session, still 50 trials long, started with FI (or reinforced) trials and ended with extinction (or

non-reinforced) trials. The exact number of FI trials varied randomly across sessions from a minimum of 10 to a maximum of 40, which means that extinction began unpredictably between trials 11 and 41 inclusive. During extinction trials the key and house lights remained illuminated and, therefore, no cue signalled either the beginning or the end of a trial. It is important to note that reinforced trials ended with the presentation of food contingent on a key peck. Hence, their duration was always longer than the FI. The extinction trials, however, ended when an interval exactly equal to the FI elapsed.

3. Results

3.1. Phase 1—acquisition

For each session, we divided the FI trials into ten equal bins and then computed the average response rate during each bin. This analysis yielded one response rate curve per bird per session, and all analyses of phase 1 were based on these curves. Although across-trials average curves do not always represent performance during single trials, our purpose at this stage was to describe the major trends in acquisition, the molar regularities, and for that purpose average rate curves seemed appropriate (see Baron and Leinenweber (1994), for a discussion of this point).

Fig. 1 shows average response rate curves during the first and the last 25 trials of the first session of training. During the first half of the session, and despite some variation in the absolute rate of responding, five birds showed a bitonic curve with response rate first increasing and then decreasing, for bird 10413 response rate decreased throughout the interval. The major change that occurred from the first to the second half of the session was the decrease in the response rate curve during the initial segment of the interval. For bird 10417, we also observed an increase in the terminal rate of pecking.

To compare more directly the behavior of the two groups during the first session, the two response rate curves displayed in each panel of Fig. 1 were divided by the maximum rate observed

during the first 25 trials. The resulting individual curves were then averaged to obtain group data. Fig. 2 shows the results. The top panel reveals no appreciable difference between the two groups during the first 25 trials. The two curves overlap considerably and show a peak of responding at ≈ 10 s into the trial. The bottom panel shows that the initial segment of the response rate curve was reduced proportionately more for the FI 80 s than the FI 40 s birds.

Fig. 3 shows how the response rate curves changed from the first to the last sessions of training. The four sessions selected for each bird represent qualitatively distinct stages in the development of temporal control. The general trend was always the same—the response rate curves rotated across sessions such that the rate during the early part of the trial decreased, whereas the terminal rate increased. Specifically, during the first stage (filled circles), the rate curves increased quickly and then either decreased or remained roughly constant. In the second stage (unfilled circles), the curves increased less rapidly and tended to remain constant for a large part of the interval. For four of the six birds, the terminal rate was greater than during the first stage. In the third stage (filled squares), the rate curves increased slowly as in stage two, but the terminal rate was higher than before. Finally, at the steady state (unfilled squares, fourth stage), the rate curves showed their typical sigmoidal shape, with the 40 s birds reaching higher terminal rates than the 80 s birds. Moreover, there was some indication that the 40 s birds may have reached an asymptotic rate of pecking because the terminal segments of their curves are either concave down (bird 8554) or flat (birds 10413 and 2775N).

The rotation of the curves meant that there was a segment in the interval during which the rate did not change appreciably across sessions. The double arrows in Fig. 3 identify the approximate locations of these segments. To their left the downward movement of the curves predominated, to their right, the upward movement did, in their vicinity, the movement was small and its direction inconsistent. This finding is particularly striking in bird 8554, but it can also be seen in the curves of the other birds. The location of the arrows along

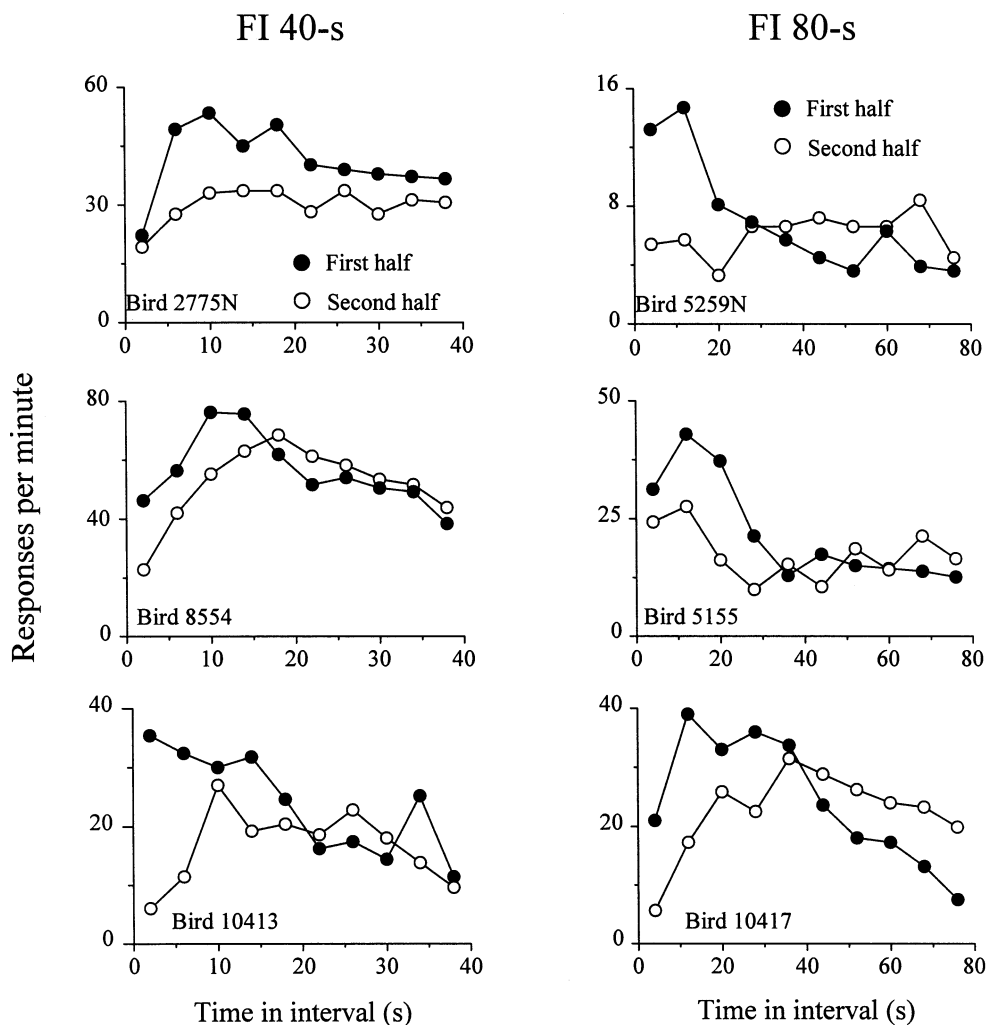


Fig. 1. Average response rates of individual birds during the first session of phase 1. Filled circles show data from the first 25 trials, and empty circles show data from the last 25 trials. The time bins were 4 s long for the FI 40 s birds, and 8 s long for the FI 80 s birds.

the time axis was also consistent with the results of linear regressions of response rate during each bin on session number. For the bins to the right of the arrows, the slopes of the best-fitting lines were generally positive because the terminal rate increased across sessions; for the bins to the left, the slopes were generally negative because the initial rate decreased; at or close to the location of the arrows, the slopes were not significantly different from 0. We return to this finding and its significance later.

The curves shown in Fig. 3 do not reveal clearly how fast the initial and terminal rates approached their steady state because, in most cases, their total amount of change differed. For example, the terminal rate for some birds (e.g. 5259N and 10413) was substantially lower during the first session than at the steady state, but the initial rate was already very close to the steady state. Hence, to compare speeds of approach to the steady state one must take into account the difference in starting and ending points of the initial and terminal

segments of the rate curves. To that end, we normalized the rate curve for each session by using the following formula:

$$R^*(n, t) = \frac{|R(n, t) - R(1, t)|}{|R(\infty, t) - R(1, t)|}$$

where $R^*(n, t)$ is the normalized rate at time t and for session n , $R(n, t)$ is the corresponding, non-normalized rate, $R(1, t)$ is the rate observed during the first session, and $R(\infty, t)$ is the average of the rates observed during the last ten sessions of phase 1. The average over the last ten sessions was used to increase the reliability of the estimated rate curve at the steady state. (A similar averaging for the initial sessions would be inappropriate because most of the changes in response rate that we are trying to characterize occurred during the first sessions, see Fig. 3 and, later Fig. 4). The absolute value was used so that increases and decreases in response rate could be compared directly. From the normalized rate curve for each one of the first 15 sessions we selected as represen-

tatives of the initial and terminal rates the first and the last three bins of the interval, respectively. The remaining sessions were not included in the analysis because they showed only the stabilization of the rates at 1.

A 3-way, between-within Anova (2 groups \times 6 bins \times 15 sessions) revealed no significant effect of group [$F(1, 4) = 0.51$, $P = 0.52$], but strong effects of bins [$F(5, 20) = 9.56$, $P = 0.001$], sessions [$F(14, 56) = 33.58$, $P < 0.0001$], and the bins \times sessions interaction [$F(70, 280) = 2.33$, $P < 0.0001$]. Fig. 4, which shows how the normalized rates changed during the first 15 sessions, helps to interpret these findings. The top and bottom panels are for the FI 40 s and the FI 80 s groups, respectively. Only averages across birds are shown because there were no major individual differences.

First, the two groups showed similar results (no group effect). Second, not surprisingly, across sessions all rates came closer to their steady-state values (the sessions effect). Third, the across-session average rates decreased from bins 1–10 (the bin effect). Fourth, the initial rates approached the steady state faster than the terminal rates (the bin \times sessions effect). In particular, the rate during the first bin (i.e. 4 and 8 s for the FI 40 s and the FI 80 s groups, respectively) approached the steady state faster than the rate during the second and third bins. And whereas the differences between the last three bins were not large, there was some indication in the FI 80 s group that the rate during the eight bin changed faster than the rate during the ninth and tenth bins.

In summary, the average response rate curves rotated across sessions. The initial rate decreased while the terminal rate increased. Furthermore, the initial rate approached the steady state faster than the terminal rate. For all birds there was also a moment in the interval during which the rate remained roughly constant across sessions.

3.2. Phase 2—extinction

Fig. 5 shows the response rates observed immediately after the last food trial and averaged over the first five (circles) and last five (squares) sessions of phase 2. The unfilled symbols represent the rates during the first trial without food, and

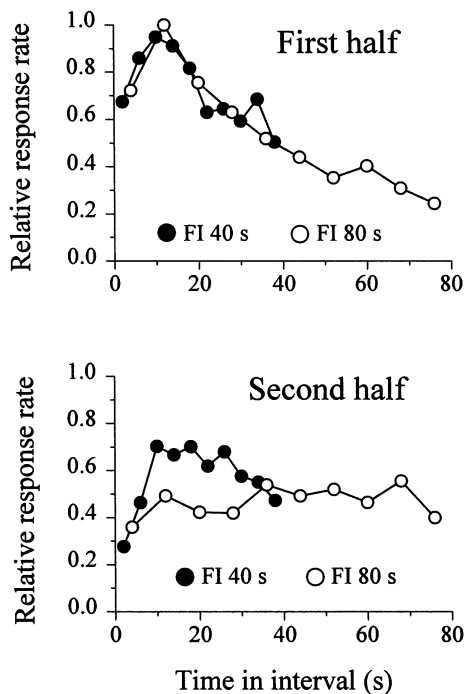


Fig. 2. Average of the relative response rate curves obtained during the first (top) and the last (bottom) 25 trials of the first session of phase 1.

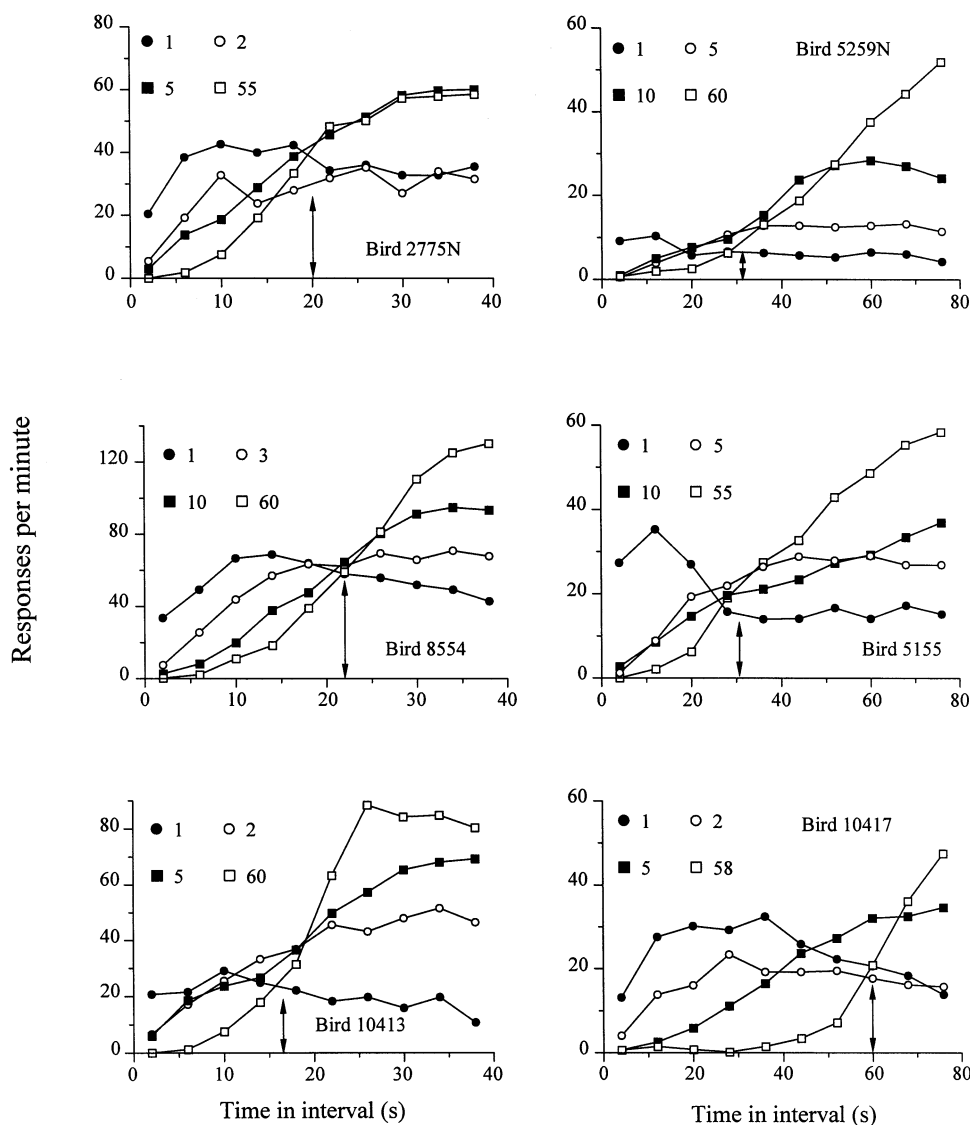


Fig. 3. Response rate curves from selected sessions of phase 1. The numbers identify the sessions. The double arrows locate the segment within the trial during which response rate did not change appreciably across sessions. Response rates decreased during the interval to the left of the arrows, and increased during the interval to the right. The time bins were 4 and 8 s long for the FI 40 s and the FI 80 s birds, respectively.

the filled symbols represent the rates during the next ten trials (remember that an extinction trial is defined as an interval of time equal to the FI parameter). With regard to the first five sessions, three aspects of the results are noteworthy. First, because the birds typically paused after each food episode, the response rates during the first extinc-

tion trial were lower than during the subsequent trials (cf. unfilled vs. first filled circles). Second, the rates during the next ten trials changed little for all birds except 10417. Third, the two best-fitting regression lines through the log-transform of the group data (see bottom panels, filled circles) had similar slopes: -0.0012 responses per s^2

for the FI 40 s birds, and -0.0011 responses per s^2 for the FI 80 s birds. This last result means that the relative rate of decay of response rate during the first five sessions was roughly similar in the two groups.

Consider now the last five sessions. For all FI 40 s birds and for bird 5155 response rate decreased faster (i.e. the regression lines were steeper) than in the first sessions, but for birds 10417 and 5259N response rate did not change appreciably. Overall, the FI 40 s group showed greater deceleration than the FI 80 s group (see bottom panels, filled squares). An Anova based

on the slopes of the regression lines for the two FI groups (between-subjects) during the first and the last 5 sessions (within-subjects) approached significance for the main effects [group: $F(1, 4) = 4.95$, $P = 0.09$, first vs. last sessions: $F(1, 4) = 5.33$, $P = 0.08$]. These findings should be interpreted with caution, however, because not only was the number of birds small but, more importantly, the data became noisier during the last sessions of phase 2 and some curves showed major discontinuities (see, e.g. Bird 2775N at 200 s and bird 8554 at 160 s). The source of these discontinuities was revealed by the analysis of the data from individual sessions, which we describe next.

The cumulative records for the 15 sessions of phase 2 revealed on many occasions a reliable pattern of sustained pecking that alternated with pauses. The next three figures illustrate this finding. Fig. 6 shows the behavior of bird 2775N (FI 40 s group) during five different sessions. Within each panel, the left curve corresponds to the last ten reinforced trials and the right curve to the first trials in extinction, the two periods of observation have equal duration, ≈ 400 s.

The figure shows that the response profile during the reinforced trials, the typical FI scallops, did not change appreciably across sessions, but the same was not true during extinction. In the first session, response rate after the last food remained high and the bird did not pause for an interval ten times greater than the FI value. However, during the next sessions, the behavior during extinction changed from uninterrupted pecking to alternation between pecking and pausing. The remaining four panels in Fig. 6 reveal this oscillatory pattern. In all cases, the pattern of pecking and pausing observed during the extinction trials was similar to the pattern observed during the reinforced trials, except that the period of the oscillation, the interval from one pause to the next, was slightly longer. In the span of ≈ 400 s, there were ten pause to pause cycles during reinforcement, but only seven during extinction.

Fig. 7 shows the behavior of bird 10417 (FI 80 s group). For session one, we have also included the last extinction trials. The results were qualitatively similar to those of bird 2775N except that

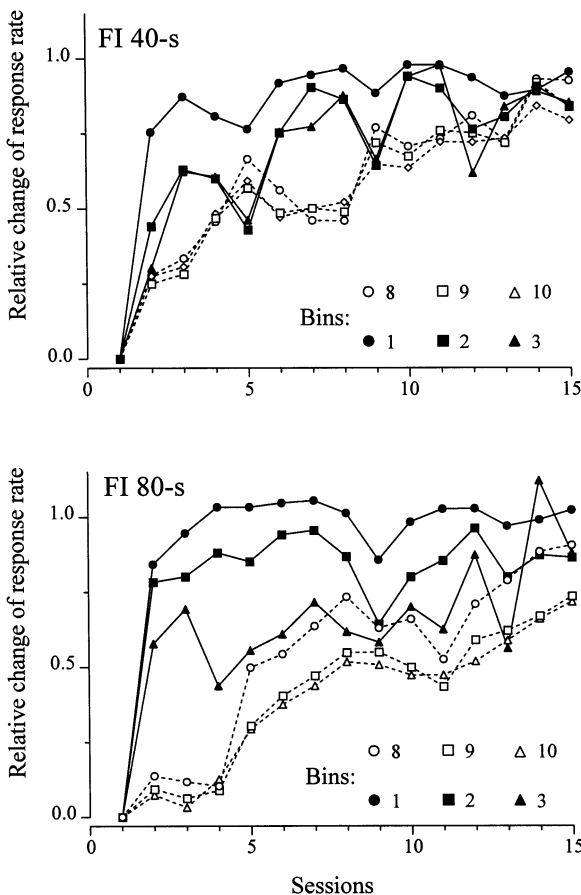


Fig. 4. Rates of approach to the steady state of the initial and terminal rates. In each panel, the data points show the normalized response rates during the first and the last three bins of the first 15 sessions. The bins were 4 s long for the FI 40 s birds (top panel), and 8 s long the FI 80 s birds (bottom panel).

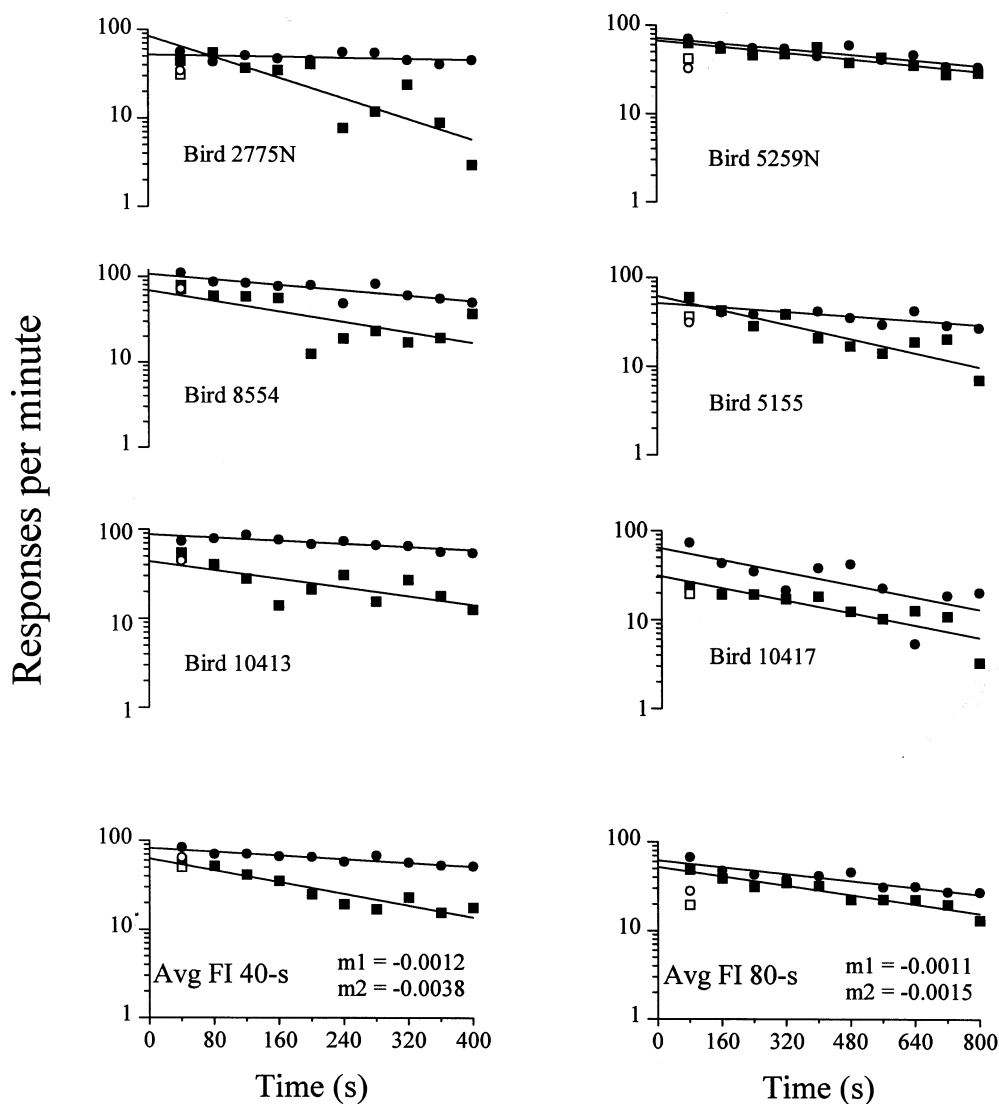


Fig. 5. Response rates during the first moments in extinction after the last reinforced trial (phase 2). The circles and the squares show the averages of the first and the last five sessions of phase 2, respectively. The empty symbols show the response rates during the first extinction trial (i.e. immediately after the last food episode) and the filled symbols show the rate during the next ten trials. The lines are the best-fitting regression lines. The bottom panels show the average data for each group and the slopes of the corresponding regression lines. Note the logarithmic scale on the Y axis.

the pause peck pattern started earlier. During the first session, the first pause (excluding the post-reinforcement pause) occurred after 340 s of sustained pecking, and the second pause occurred after an additional 260 s of pecking. Thereafter, however, pauses became more frequent and by the end of the session the period of the peck-pause

cycles approached 80 s, the interreinforcement interval (we return to this finding below). In subsequent sessions, the oscillatory pattern started earlier and earlier after the last reinforcer. As for bird 2775N, the periods of the oscillations were longer during extinction than during reinforcement.

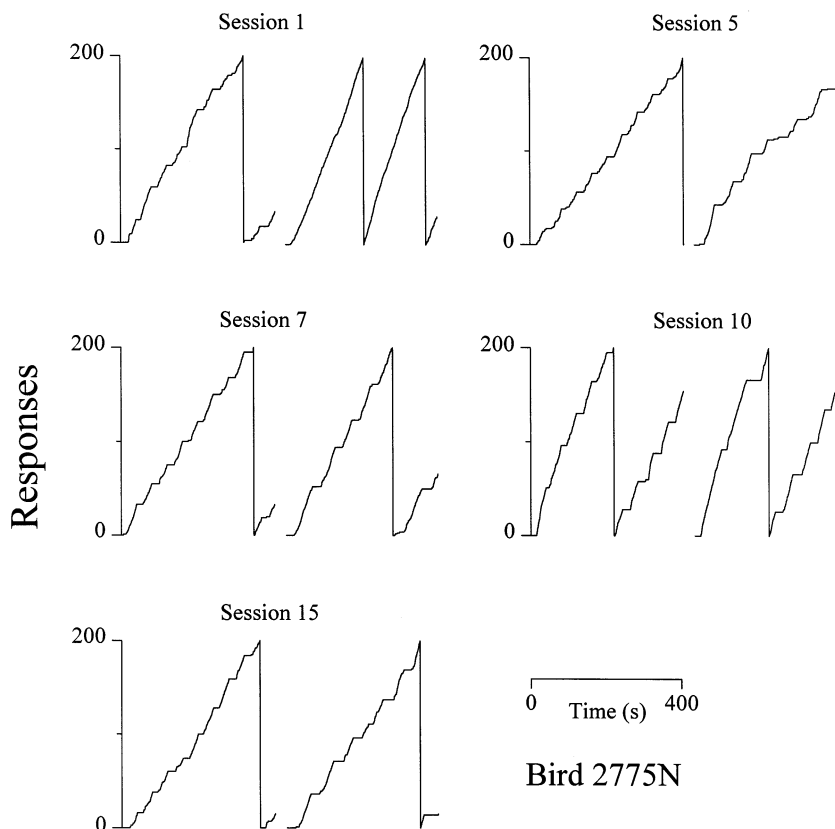


Fig. 6. Cumulative records of bird 2775N (FI 40 s) during phase 2. For each session, the left curve shows performance during the last ten FI trials, and the right panel shows performance during the first extinction trials. The two curves have the same total duration, ≈ 400 s.

The results for the remaining birds were qualitatively similar and are displayed in Fig. 8. Because the response patterns during food trials were the typical scallops, only behavior during extinction is shown. In addition, unlike Figs. 6 and 7, the cumulative records include all trials without food. The top two rows (FI 40 s birds) show, again, almost uninterrupted pecking during the first session and reliable oscillations during subsequent sessions. The bottom two rows (FI 80 s birds) also show sustained pecking during the first moments of extinction in session one, but for bird 5155 there was some indication of pause peck cycles by the end of the session. In subsequent sessions, and despite rough grain in some records, the pause peck oscillations were visible during extended periods of time (e.g. beginning of session

8 for bird 5155, end of sessions 8 and 15 for bird 5259N).

To summarize, the cumulative records of phase 2 showed spontaneous pause–peck oscillations in all birds. For two of them (birds 5155 and 10417) these oscillations occurred at the end of the first extinction session, but for the remaining birds they occurred only after one to three sessions. Thereafter, the oscillations started progressively earlier after the last food trial, and the duration of the pauses increased. Although the pause–peck oscillations were usually clear, there were some sessions, or segments within a session, during which they did not occur reliably.

To further characterize the pause–peck oscillations and relate them to the FI durations, we performed two additional types of analyses. The

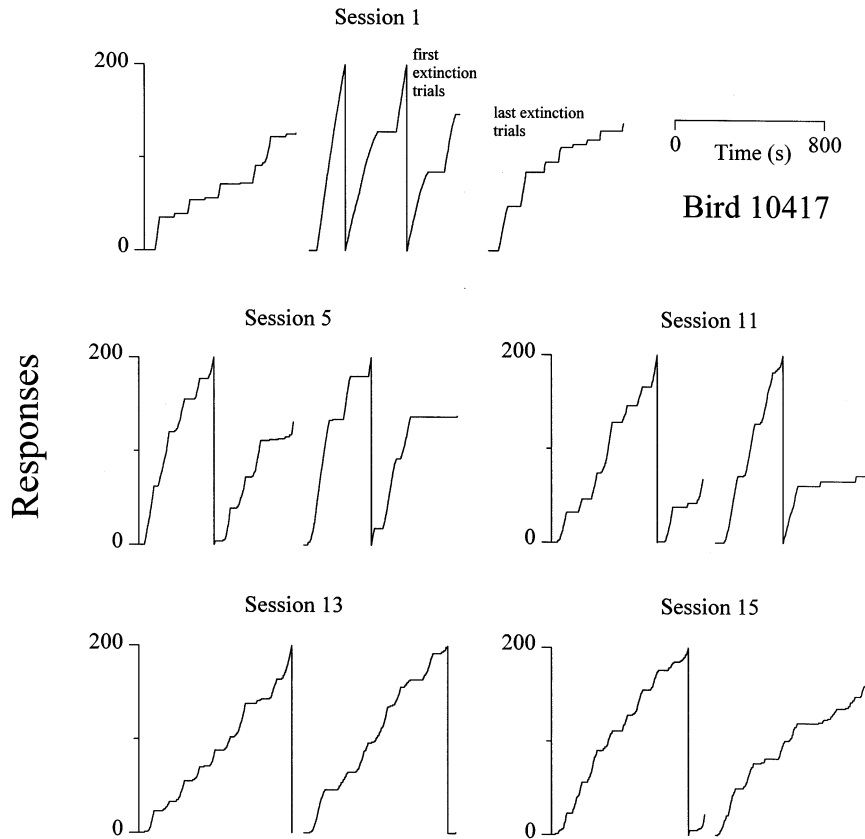


Fig. 7. Cumulative records of bird 10417 (FI 80 s) during phase 2. For session 1, the left curve shows performance during the last ten FI trials, the middle curve performance during the first extinction trials, and the right curve performance during the last extinction trials. For the remaining panels, only the last reinforced trials and the first extinction trials are shown. The curves for each session have the same total duration, ≈ 800 s.

first, a Fourier analysis, attempted to identify the frequencies that were present in the rate data of those segments of session for which the cumulative record indicated, at least by visual inspection, a reliable oscillatory pattern. The second analysis, based on pause duration, attempted to quantify the periodicities of the pause–peck oscillations by using the data from all 15 sessions.

The session segments illustrated in Figs. 6–8 were submitted to a discrete fourier transform (DFT) routine. The input to this routine consisted of the number of responses produced during successive 4 s bins (FI 40 s group) or 8 s bins (FI 80 s group), for a total of 103 bins. That is to say, the input samples were always $4 \times 103 = 412$ s long for the FI 40 s group, and $8 \times 103 = 824$ s

long for the FI 80 s group, in both cases the approximate duration of ten trials. The output of the DFT routine yielded a vector of complex numbers, c_j , related to the amplitudes of the (different frequency) sine and cosine waves present in the response rate data. Because we were interested only in how the total energy of the input signal was divided among the various frequencies, we present the results in terms of the magnitudes $|c_j|$, focus our attention on the relative, not the absolute, values of these magnitudes, and disregard the DC component of the signal (frequency 0, for further details on DFT analysis, consult any book on digital signal processing).

Fig. 9 shows the power spectra for birds 2775N (top) and 10417 (bottom). The length of each

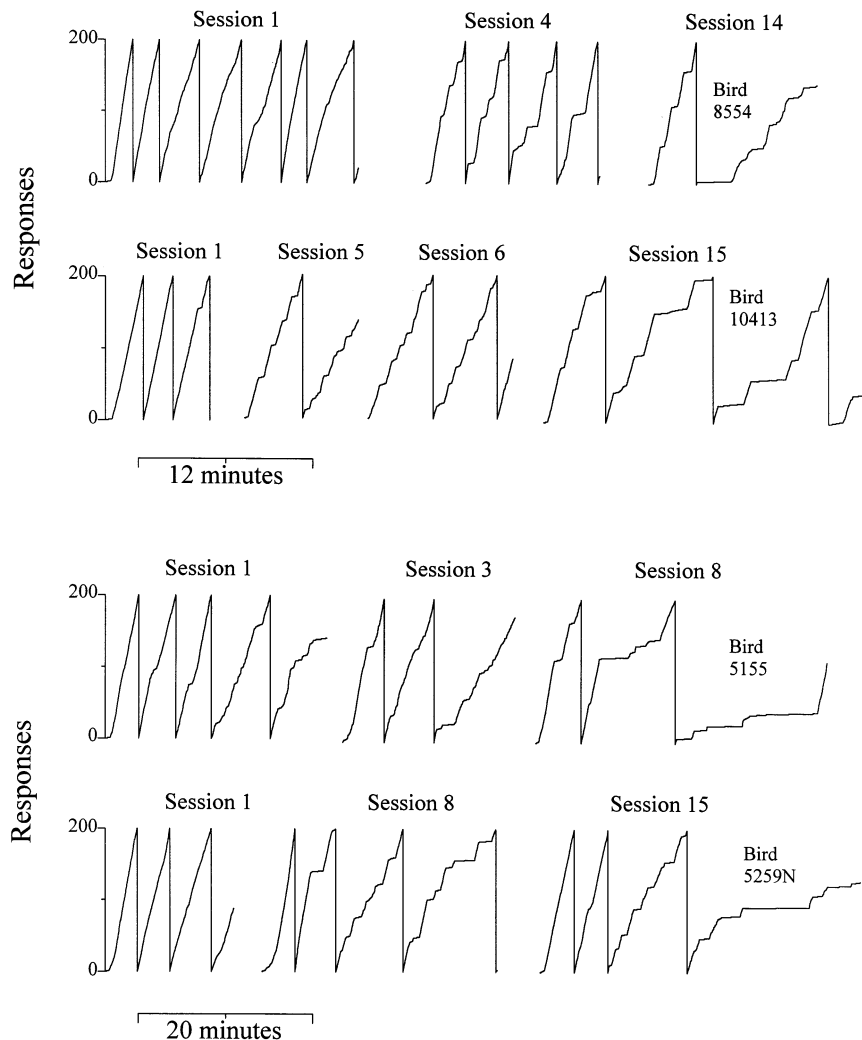


Fig. 8. Cumulative records of birds 8554 and 10413 (FI 40 s group, upper panels) and 5155 and 5259N (FI 80 group, lower panels) during phase 2. Each curve shows the complete record obtained during the extinction portion of the session.

vertical line shows the 'amount' of the corresponding frequency present in the signal. The top and bottom rows of each panel correspond to the reinforcement and extinction trials, respectively. Consider the top row first. Given that food occurred approximately every 40 s, and that the bird paused immediately after food and then started to peck roughly midway through the interval, we expect most of the energy to be near the frequency $1/40$ s, or 0.025 Hz, all other frequencies should have low and similar energies. That this was indeed the case is evidenced by the single spike at the

predicted frequency on the top row graphs. Similarly, if during extinction the pause peck patterns resembled the patterns shown during reinforcement, possibly with slightly longer periods as the cumulative records suggest, then we would also expect high energy in the 0.025 Hz region and in adjacent regions to the left (i.e. regions corresponding to lower frequencies or, equivalently, to periods longer than 40 s). The second row shows that after the first session this was indeed the case, for the main frequencies in the signal were close to, but slightly less than, the reinforcement frequency.

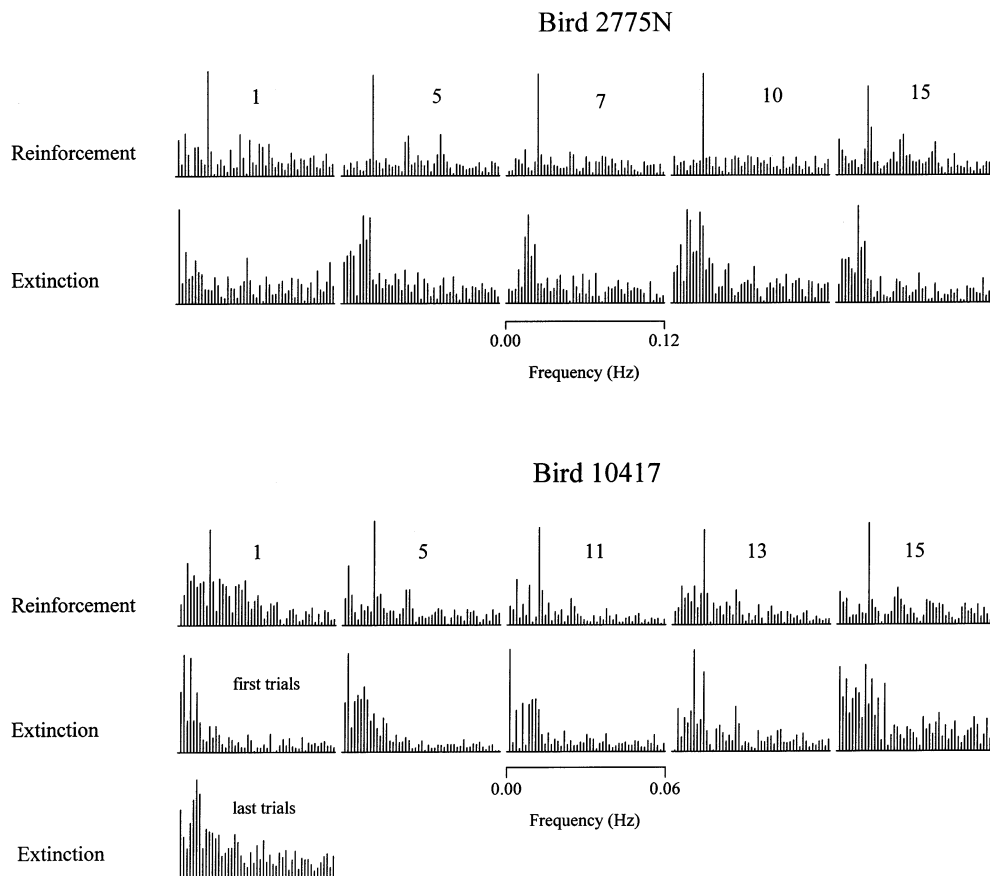


Fig. 9. Power spectra for birds 2775N (upper panel, FI 40 s) and 10417 (lower panel, FI 80 s). The numbers identify the sessions. For each panel, the top and bottom rows correspond to last FI and the first extinction trials, respectively. For session 1 of bird 10417, the spectrum for the last extinction trials is also included. The large spikes during the FI trials correspond to the reinforcement frequency (0.025 Hz in the top, and 0.0125 Hz in the bottom). Figs. 6 and 7 show the corresponding cumulative records.

The bottom panel shows similar results for bird 10417 who was exposed to an FI 80 s schedule. During the FI trials the energy was concentrated mostly near the 1/80 s, or 0.0125 Hz, region, even though in some sessions (e.g. 1, 5, 11, and 13) the signal also contained a substantial amount of lower frequencies. During the extinction trials, and starting with the last segment of the first session, we observe relatively high energies in frequencies close to, but slightly less than, the reinforcement frequency.

Fig. 10 shows the spectra for the remaining birds. These spectra correspond to the initial segments of the cumulative records displayed in Fig.

8 with two exceptions. First, session 1 was excluded from the analysis because in most cases the corresponding cumulative records showed only uninterrupted pecking. Second, the spectra for bird 5259N correspond to the middle segments of the cumulative records, close to the second resetting of the pen, because these were the segments where oscillations occurred. The spikes in the top panels locate the reinforcement frequency. In general, the same picture emerged: The main frequencies were concentrated in the reinforcement region and slightly to its left.

Fig. 11 summarizes the DFT data from the extinction trials. The top panels show the location

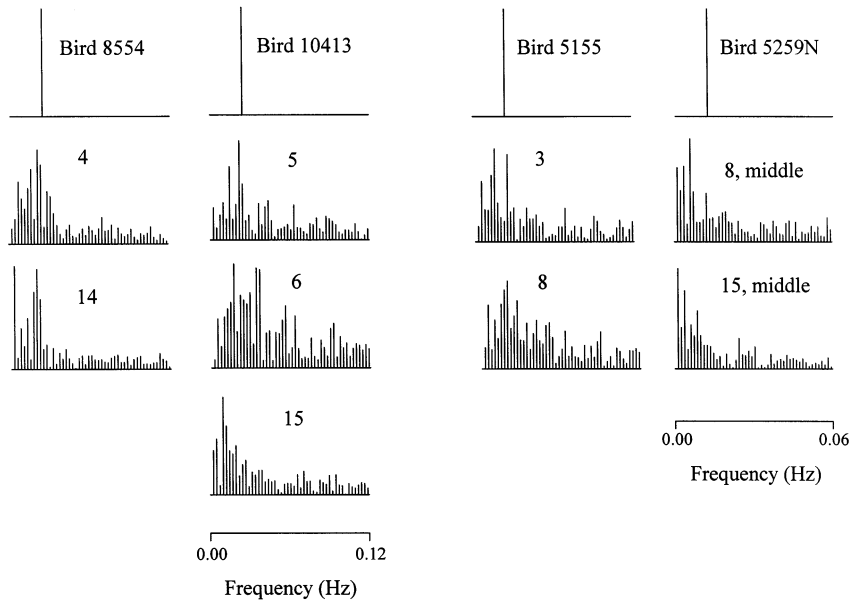


Fig. 10. Power spectra for birds 8554 and 10413 (FI 40 s group), and 5155 and 5259N (FI 80 s group). The numbers identify the sessions. The topmost panels locate the appropriate reinforcement frequency. The spectra correspond to the beginning of the corresponding cumulative records shown in Fig. 8 except that the spectra for bird 5259N correspond to the middle segments, i.e. after the second resetting of the cumulative record.

of the reinforcement frequencies, the middle panels plot the average of the spectra shown in Figs. 9 and 10 excluding the first sessions, and the bottom panel replots the data in relative frequencies, with 0.025 and 0.0125 Hz as the standards. The average spectra confirm that the main frequencies were clustered slightly to the left of the reinforcement frequency, and the normalized plot shows a high superposition of the spectra except at the lowest frequencies. In summary, the DFT results agreed with our interpretation of the cumulative records, namely, that in extinction the birds displayed pause–peck oscillations whose periods were close to, but slightly longer than, the inter-reinforcement interval.

Our second analysis used all extinction data from the 15 sessions, and it attempted to quantify the time between consecutive pauses without being influenced—as the DFT is—by the number of intervening responses. Unfortunately, the definition of a pause is not straightforward because any interval without a peck can be interpreted as either a true pause from pecking, a pause similar

to the post-reinforcement pause, or as an interresponse time, a ‘pause’ while pecking. Any criterion used to classify intervals without pecks as either pauses or interresponse times will occasionally fail and an interresponse time will be classified as a pause, or a pause as an interresponse time. To minimize the first type of error, the criterion for a pause should be greater than most interresponse times (themselves a function of overall response rate and hence of overall reinforcement rate), to minimize the second type of error, however, the criterion should be smaller than most post-reinforcement pauses, presumably our best guides to true pauses during extinction. But given the large number of responses generated during extinction, the minimization of the first type of error should receive priority because most intervals without pecks are in fact interresponse times. These considerations led us to the following pause criteria: For the FI 40 s group, a pause was defined as any period of 10 s or more without a peck, for the FI 80 s group, a pause required a minimum of 20 s without a peck. These values

were significantly larger than the interresponse times but less than the average post-reinforcement pauses.

Fig. 12 illustrates the analysis with the data from the first session of bird 10417. For comparison purposes, we have included also the FI trials. In the top panel the vertical dotted lines show when food occurred and the trial ended. Because the bird paused and then pecked continuously until it received food, one might say that, during reinforced trials, cycle duration equalled trial duration. The middle panel corresponds to the extinction trials. The vertical dotted lines separate, not the 80 s trials, but the pause-peck cycles defined by the 20 s pause criterion. Bird 10417 paused only twice (excluding the post-reinforcement pause) during the first 750 s, but thereafter it paused more frequently and the duration of its pause-peck cycles approached the inter-reinforcement interval. This result is seen more clearly in

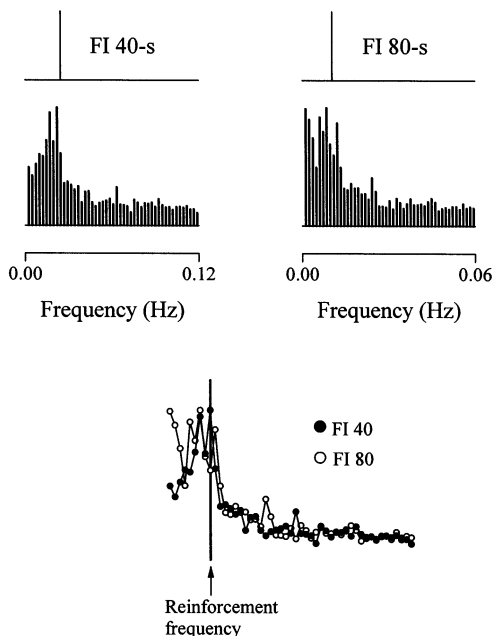


Fig. 11. Summary of the DFT analysis of response rates during extinction trials. The topmost panels locate the reinforcement frequencies. The middle panels plot the average of the spectra shown in Figs. 9 and 10 excluding the first extinction sessions. The bottom panel replots the average data in relative frequencies, with 0.025 and 0.0125 Hz as the standards.

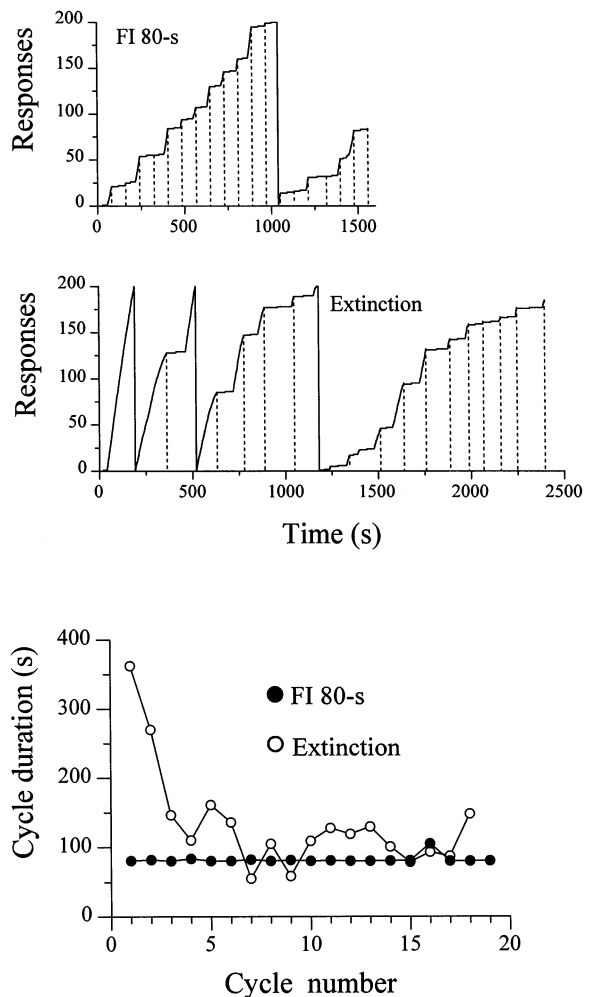


Fig. 12. A detailed analysis of the cycle durations produced by bird 10417 during the first session of phase 2. The upper panel shows the cumulative record during the FI trials. The dashed vertical lines indicate food deliveries (end of a trial or cycle). The middle panel shows the cumulative record during the extinction trials. The dashed vertical lines indicate the beginning of pauses at least 20 s long. The interval between two consecutive vertical lines is the cycle duration. The bottom panel plots the cycle durations produced during the FI (filled circles) and the extinction (empty circles) segments of the session.

the bottom panel, which plots the duration of successive cycles both during reinforcement and extinction. Whereas the reinforcement cycles remained close to 80 s, the extinction cycles decreased and seemed to stabilize around 100 s. The

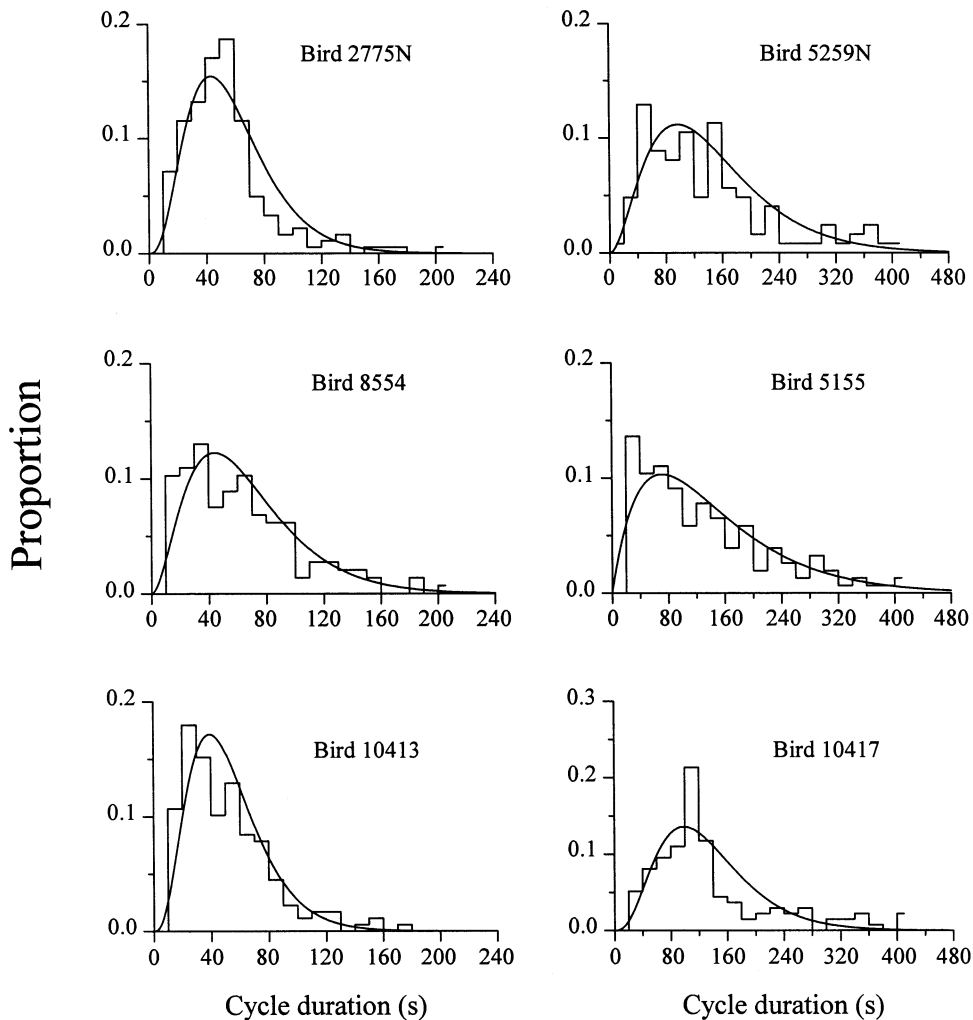


Fig. 13. Relative frequency distributions of the cycle durations produced by individual birds during extinction. The data come from all 15 sessions of phase 2. The smooth curves are the best-fitting Erlang distributions.

same cycles were obtained when the pause criterion equalled any value from 10 to 40 s, and changed only slightly when the criterion was set at 5 s. Interestingly, the number of responses per cycle also seemed to converge to the number observed during the reinforced trials—the average number of responses during all but the first two extinction cycles equalled 16, whereas during reinforcement it equalled 15.

The analysis illustrated in Fig. 12 was conducted for all birds across all sessions and the

obtained distributions of cycle durations are shown in Fig. 13. For five birds, the distribution first increased and then decreased. For the FI 40 s birds, most cycle durations were clustered around 40 s; for two of the FI 80 s birds (except 5155) most cycles were clustered around 100 s. The distributions for the FI 80 s birds also showed greater variance than the distributions for the FI 40 s birds.

The shape of the various distributions was reasonably well fit by the Erlang probability density

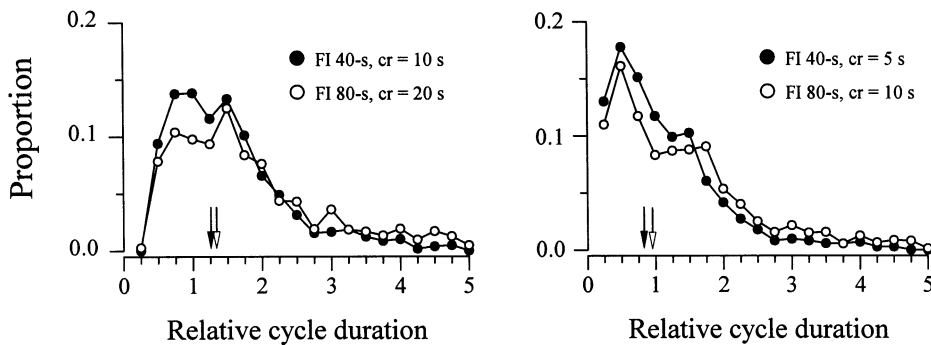


Fig. 14. Relative frequency distributions of normalized cycle durations with 40 s (filled circles) and 80 s (empty circles) as standards. The left panel shows the distributions of cycle durations obtained with a 10 s (FI 40 s) or a 20 s (FI 80 s) pause criterion. The right panel shows the same data with a 5 s (FI 40 s) or a 10 s (FI 80 s) pause criterion. The arrows indicate the mean of the medians of the individual distributions.

function f :

$$f(t) = \frac{s^m}{(m-1)!} e^{-st} t^{m-1} \quad (t \geq 0, s > 0, m > 0)$$

where t is the duration of a cycle and s and m are scale and shape parameters, respectively. m is constrained to take only integer values. The average and variance of the distribution equal m/s and m/s^2 , respectively. The Erlang distribution is the cornerstone of some models of timing (see Section 4) but its use here was motivated only by the shape of the histograms in Fig. 13. The modes of the fitted curves were close to the corresponding FI parameter: 39, 43 and 44 s for birds 10413, 2775N, and 8554, and 100, 95 and 71 s for birds 10417, 5259N and 5155, respectively.

The similarities of the graphs in Fig. 13 as well as the DFT data shown in Fig. 11 suggested that the two average curves might be scale transforms. The left panel of Fig. 14 shows this was approximately true. The two average curves are plotted as a function of relative time, with 40 and 80 s as the standards. To a large extent the curves overlapped and showed a concentration of cycle durations around values close to the FI parameter. Thus, the average of the medians of the individual distributions equalled 50 s for the FI 40 s group (see filled arrow) and 106 s for the FI 80 s group (unfilled arrow). However, the curve for the FI 80 s birds had a longer tail and therefore proportionately greater spread than the curve for the other birds.

The foregoing conclusions were not strongly affected by the criteria used to separate pauses from interresponse times. The right panel of Fig. 14 shows the average results when the pause criteria equalled 5 and 10 s for the FI 40 s and the FI 80 s groups, respectively. Not surprisingly, the curves shifted to the left, but they still remained close to each other and the curve for the FI 80 s group still showed proportionately greater spread. The average of the medians of the individual distributions was 33 s for the FI 40 s group and 77 s for the FI 80 s group (see arrows).

4. Discussion

The present study attempted to characterize what pigeons learn when they receive food periodically by pecking a key. We suggested that this general issue subsumes at least two distinct questions, how the pattern of key pecking changes during the course of training, and how responses occur when food is removed from the situation. Concerning the first question, our findings showed that during the first trials of the first session, response rate increased until ≈ 10 s after food and then decreased. This result was observed for both groups of birds (see Figs. 1 and 2), and it illustrates the first stage of FI training, the inverted scallops in the cumulative record. Next, response rate immediately after food decreased to

a low level and the rate curve quickly reached a steady, terminal value. In the cumulative record, this profile approximates a straight line, the second stage of FI training. During subsequent sessions, the terminal rate continued to increase while the initial rate decreased still further, and by the end of training the typical sigmoidal rate curve was obtained. This final stage corresponds to the typical scallop in the cumulative record.

The preceding results are consistent with the account of Ferster and Skinner (1957), of how typical FI performance develops (see also Skinner, 1938). This consistency, in particular for the first session, may be surprising at first because our birds were exposed to autoshaping training whereas Ferster and Skinner's were generally exposed to continuous reinforcement (CRF). On the other hand, given that during at least the last two sessions of autoshaping our birds pecked the key within 2 or 3 s from its illumination and received food immediately, any differences between autoshaping and CRF training must have been negligible.

The results of phase 1 also showed that the average rate curves rotated around a (roughly) fixed point (see Fig. 3), a moment in the interval during which response rate did not change significantly across sessions. This moment may identify a dynamic equilibrium point, that is, a moment when two opposite tendencies cancel each other, the tendency to reduce response rate, presumably an inhibitory effect related to the previous reinforcer (e.g. Kello, 1972; Staddon, 1972), and the tendency to increase it, presumably an excitatory effect related to the subsequent reinforcer. In the time axis, the first tendency is stronger to the left of the rotation point, whereas the second tendency is stronger to its right.

Finally, the data of phase 1 also showed that the initial rate during the FI converged to its steady state faster than the terminal rate (Fig. 4). In particular, the rate immediately after the reinforcer decreased substantially during the first two sessions, whereas the rates during the following moments decreased more slowly. These differences in the rates of convergence also suggest that different processes may underlie the changes in the initial and terminal rates.

Concerning the second question (i.e. behavior during extinction), our findings showed that during the first session of phase 2 all birds responded at a steady rate after the last reinforcer and for intervals much longer than the FI value. More specifically, five birds showed a sustained rate of pecking for intervals ten times longer than the inter-reinforcement interval, and one FI 80 s bird (10417) paused only twice during the first 750 s. Again, these findings agree with the observations of Ferster and Skinner (1957) (see also Kello, 1972; Staddon, 1972 on the reinforcement omission effect). In subsequent sessions, and despite the peck–pause oscillations that started to occur, the overall response rate declined only slightly during the first ten trials in extinction (see Fig. 5). However, with additional sessions of reconditioning and re extinction, the birds responded less or quit responding earlier in the absence food. In particular, the FI 40 s birds showed greater deceleration than the FI 80 s birds.

After the first sessions, however, overall response rate no longer described appropriately the major features of the birds' behavior, even during the first trials in extinction, because behavior changed from uninterrupted pecking to pause peck oscillations. These oscillations, which occurred in the absence of any external stimulus change, were related to the FI value. Thus, for the birds who experienced 40 s interfood intervals the period of the oscillations was close to 40 s, and for the birds who experienced 80 s intervals between food the period was close to 80 s. The preceding findings were revealed by cumulative records (Figs. 6–8), by power spectra obtained from a Fourier analysis (Figs. 9–11), and by histograms of pause-to pause cycle durations (Figs. 12–14). Although Ferster and Skinner (1957) [see also Innis and Staddon (1971) Fig. 5] had already reported that behavior during extinction resembles behavior during the FI training, our findings add a quantitative aspect to their report. For, not only do pigeons pause and then peck during extinction, but they seem to do so with a period related to the FI duration.

Given the preceding findings, two issues naturally come to mind: (1) to what extent did the pause peck oscillations stem from the FI training

rather than from the reconditioning and re extinction that took place across sessions? In other words, would reliable oscillations occur if the procedure in phase 2 were replaced by one long extinction session, for example? (2) could the pause–peck oscillations be due simply to spontaneous recovery? Concerning the first issue, the following findings suggest that the procedure could not by itself have engendered the oscillations. One bird (10417) showed reliable pause–peck oscillations during the first session of phase 2, that is, before any reconditioning and re-extinction (see Fig. 12), other birds showed reliable oscillations during the second or third sessions, that is, after minimal exposure to the procedure, occasionally the pause–peck oscillations started with a long period that then decreased to a value closer to the inter-reinforcement interval (e.g. session 10 in Fig. 6). These within-session changes suggest that the bird's own behavior, instead of, or in addition to, reconditioning and re-extinction, caused the change in the period of the oscillation. Concerning the second issue, spontaneous recovery is certainly consistent with the fact that the onset of the oscillatory process was preceded by a long period of pecking in extinction (see e.g. Fig. 12), but it is not clear how spontaneous recovery could account for the relationship between the average period of the oscillations and the duration of the FI.

The preceding results question common wisdom about what pigeons learn when they receive food periodically. For example, Gallistel (1990) states that in extinction 'the pigeon goes on pressing at a very high rate for an interval of time three or four times longer than the fixed interval, pauses for an unpredictable amount of time, resumes pecking with another burst much longer than the fixed interval, and so on' (p. 297). He further added that behavior in extinction indicates that during training pigeons formed aperiodic representations of the temporal interval since the last reinforcer. Our results, on the other hand, indicate that the duration of the pause peck cycles are neither unpredictable nor much longer, on average, than the FI. They are also difficult to interpret in terms of aperiodic representations, for they clearly contain periodicities related to the FI parameter.

Recently, Kirkpatrick-Steger et al. (1996) reported behavioral oscillations in a temporal task similar to an FI schedule. These authors trained pigeons in a peak procedure whereby regular FI trials (or food trials) alternated randomly with trials that were four times longer but ended without food (empty trials). They found that when the ratio of empty to food trials was 4 to 1, the response rate curves showed double peaks during the empty trials. That is, response rate increased from the beginning of the trial until the usual time of food, then it decreased symmetrically, increased again and peaked at three times the FI value, and then it decreased until the end of the trial. The double peaks occurred after 25 sessions of training, and the second peak emerged only after the first was well developed. In addition, no second peak was observed if the empty trials ended with response independent food. The authors considered their results as evidence for an internal oscillator.

There are some similarities between the study of Kirkpatrick-Steger et al. (1996) and the present one. Both used similar ranges of FI values (15–60 s in their study, 40–80 s here) and comparable amounts of training (30–60 and 55–60 sessions, respectively). Furthermore, the procedure used during phase 2 of the present experiment also resembles a peak procedure, for food trials occurred together with longer periods without food. But there were also important differences. First, if one considers the extinction portion of each session of phase 2 the equivalent of one empty trial, then in the present experiment the ratio between empty and food trials varied randomly from 10 to 1 (when extinction was preceded by 40 FI trials) to 40 to 1 (when extinction was preceded by 10 FI trials only). The fact that oscillations occurred with such high ratios (mean = 25) contrasts with the finding of Kirkpatrick-Steger et al. (1996) that oscillations did not occur when the ratio between empty and food trials was 8 to 1. Second, whereas Kirkpatrick-Steger et al. (1996) observed oscillations only after extensive alternation of food and empty trials, we observed them either during the first, or during the first five, sessions of phase 2 (i.e. after at most five alternations between the two types of trials). Third, Kirkpatrick-Steger et

al. (1996) observed oscillations for more than two cycles only when FI trials were mixed with two sets of empty trials, one 4 and the other 8 times longer than the FI. These sustained oscillations were visible after 13 sessions of training, but disappeared after the 25th session. In contrast, during the present experiment oscillations occurred reliably for several consecutive cycles. For example, in Fig. 6, the cumulative records of sessions 5, 7, 10, and 15 all show seven cycles in the span of ≈ 400 s (see also Figs. 7 and 8). Fourth, and related to the preceding point, whereas Kirkpatrick-Steger et al. (1996) observed oscillations with period $2T$, in the present study the period was generally between T and $2T$ (see Figs. 11, 13 and 14). We address the issue of a hypothetical internal oscillator after we discuss the implications of our findings to current models and theories of timing.

4.1. Implications for models and theories of timing

Two of the most influential theories of timing are the SET developed by Gibbon (1977, 1991) and the behavioral theory of timing (BeT) developed by Killeen and Fetterman (1988) (Killeen, 1991). In what follows, we contrast our findings with both theories. However, instead of focusing directly on BeT, we focus on the dynamic instantiation of BeT proposed by Machado (1997). Because this extension of BeT stresses the role of learning in timing, we refer to it as LeT. Three empirical findings will be contrasted with SET and LeT (a) the shape of the initial response rate curve and how it changed during training (b) the sustained rate of pecking during the first extinction session, and (c) the presence of reliable pause–peck oscillations.

4.1.1. Scalar expectancy theory

The simplest version of SET assumes an internal pacemaker that generates pulses at a high but variable rate, an accumulator that is reset to 0 at the beginning of each trial and then accumulates the pulses generated during the trial, and a memory store that retains the number of pulses present in the accumulator at the end of the trial, that is,

when food occurs. Because the rate of the pacemaker is assumed to vary from trial to trial, the number of pulses stored in memory will also vary¹. To time an interval, the animal extracts a number from its memory distribution and then, throughout the interval, it compares that number with the current number in the accumulator. When the ratio between the two numbers is less than a threshold the animal starts to respond. On occasion, SET has also assumed that the animal ceases to respond if the ratio between the two numbers exceeds the threshold. SET's account is illustrated in the top panel of Fig. 15.

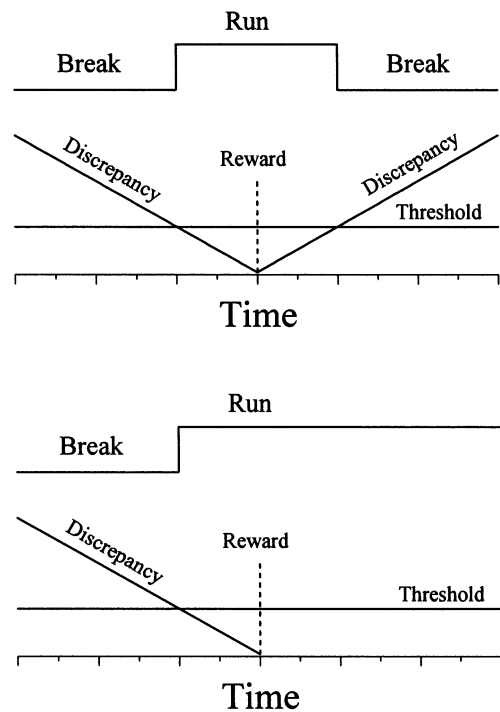


Fig. 15. Two alternative representations of SET's response rule. On the top, responding (the run state) starts when the discrepancy ratio between the number of pulses in the accumulator and the sample extracted from memory (represented by reward) falls below a threshold, and it stops when the discrepancy ratio exceeds the threshold. On the bottom, only the start rule applies.

¹ Alternatively, the rate of the pacemaker is constant but the number of pulses in the accumulator is multiplied by a random variable before the product is stored in memory. Because this alternative makes the same predictions, it will not be considered here.

Consider what this simple version of SET predicts concerning acquisition. If an animal is exposed to short interfood intervals during the preliminary training (e.g. after CRF or autoshaping), then its memory store is likely to contain only small numbers when the FI schedule is introduced. Hence, during the first session pecking will start early in the interval because a few pulses will suffice to bring the ratio below the threshold. But as the interval elapses, the number of pulses in the accumulator increases, and if the ratio between that number and the (small) memory sample exceeds the threshold before the end of the interval, then the animal stops pecking. The data from the first session agrees with the prediction of early pecking followed by no (or little) pecking.

As training proceeds, the number of memory samples related to the FI increases and eventually outnumbers the samples related to the previous training with shorter interfood intervals. Hence, with each new trial, the sample extracted at trial onset is more likely to come from the FI samples. When the sample does come from the FI samples, the bird responds as it will do during the steady state, when the sample does not, the bird responds as it did during the first session. Our findings did not support this prediction, instead, they indicated a more gradual change in response rate, a change with distinct intermediate stages (e.g. a relatively constant rate of responding during a major portion of the trial). Even when responding early in the trial was rare, performance was not at the steady state. Moreover, the different rates of change of the initial and terminal response rates shown in Fig. 4 are not derivable from a mixture of two distributions because, by definition, a mixture always yields the same normalized rate of convergence².

² Assume that response rate at time t is given by $F(t)$ in a proportion p of trials and by $G(t)$ in the remaining, i.e. $1 - p$, proportion of the trials. Across sessions, the value of p changes from 0 to 1. Let $p(i)$ be the p value during session i . Then the normalized rate at time t and during session i will equal:

$$R(i, t) = \frac{p(i) F(t) + (1 - p(i)) G(t) - G(t)}{F(t) - G(t)} = p(i)$$

which is independent of t .

Consider now the first moments in extinction. The assumption required to account for the decline of response rate during the first FI session, namely, that responding starts when the ratio falls below the threshold and stops when the ratio exceeds the threshold, predicts that responding will cease (or at least that the rate will decline) after a period of time without food no longer than two interval durations. No bird showed this behavior; instead, they all pecked for an interval at least seven times greater than the FI duration. However, an alternative reading of SET says that the use of the thresholds itself requires some experience. Hence, because the animals were never exposed to periods without food much longer than the FI duration, they would not use the threshold to stop responding. This hypothesis is illustrated in the bottom panel of Fig. 15. The problem is that the same argument would fail to predict the first stage of the FI training.

Finally, to account for the pause–peck oscillations, SET would need to invoke a spontaneous resetting of the internal clock or, more generally, a series of events functionally similar to the delivery of food. What these events might be remains unknown. In summary then until SET addresses the preceding issue and develops an explicit theory of its thresholds—how many they are, when is each one used, what determines its value, and the like—its ability to highlight the dynamics of timing will remain limited.

4.1.2. Learning to time

The top left panel of Fig. 16 illustrates the basic structure of LeT: a serial organization of behavioral states (top circles); a vector of associative links or couplings from the behavioral states to the operant response (middle connections); and the operant response itself (bottom circle). LeT assumes that after a time marker, typically a significant biological event or a reliable predictor of that event, a set of behavioral states is serially activated. The activation of state n at time t is represented by variable $X(t, n)$. During extinction the states lose their coupling with the operant response, whereas during reinforcement they increase that coupling. The strength of the coupling between state n and the operant response at time

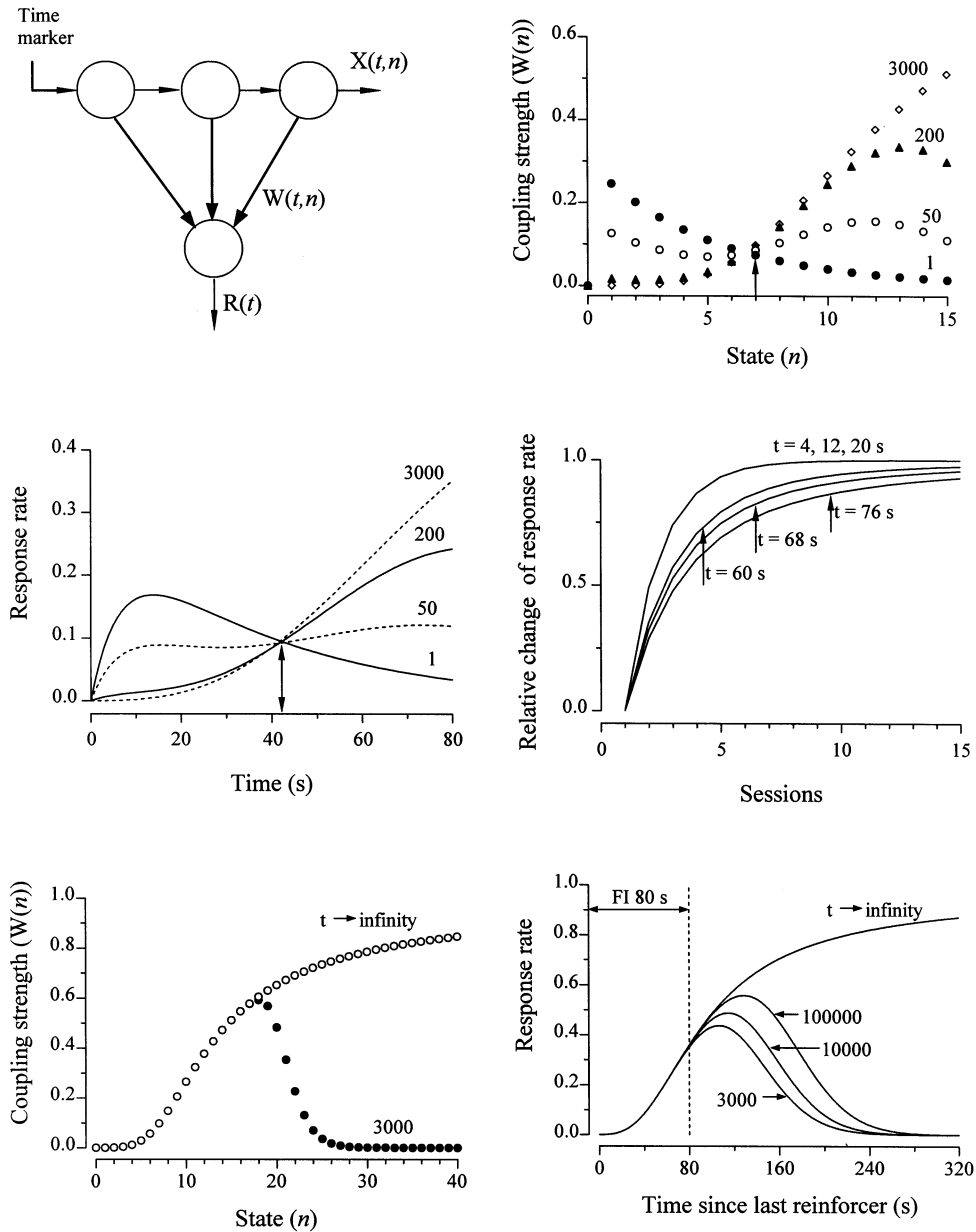


Fig. 16. The learning to time (LeT) model. The top left panel shows the serial organization of behavioral states (top circles), their couplings to the operant response (middle arrows) and the operant response itself (bottom circle). The remaining panels assume an FI 80 s schedule. The top right panel shows the strength of the couplings between the behavioral states and the operant response during different trials (see numbers), the coupling of state 7 does not change appreciably across trials (see arrow). The middle left panel shows the corresponding response rate curves; the double arrow identifies the moment in the interval during which response rate does not change. The middle right panel shows the rates of approach to the steady state of the initial and terminal response rates. The curves for the first three 8 s bins overlap. The bottom left panel shows the couplings between the behavioral states and the operant response after 3000 (filled circles) or an infinite number of trials (empty circles) and the bottom right panel shows the corresponding response rate curves in extinction. Throughout, $\lambda = 0.15$, $\alpha = 0.002$, and $\beta = 0.01$ (see Appendix A).

t is represented by variable $W(t, n)$. Paraphrasing Hodgson, as quoted in James (1961), the behavioral states are the measuring-tape, and reinforcement the dividing engine which stamps its length. (The length is stamped in the couplings). At any moment, response rate $R(t)$ is determined jointly by the profile of activation of the states and their couplings with the operant response. Appendix A shows the basic equations of LeT, and Machado (1997) may be consulted for a more extensive analysis of the model.

To understand LeT's predictions, we need to consider how the couplings between the behavioral states and the operant response (i.e. key pecking) change with training, and how such changes alter the response rate curve. Hence, in the analyses that follow, we always describe the couplings first and the corresponding response rate curve next. Assume then that an animal was exposed to short interfood intervals, as it happens during CRF or autoshaping training, before it experienced an FI 80 s schedule (the predictions remain the same for other FI values). Due to this preliminary training, only the early states will be strongly coupled with key pecking during the first FI trials. That is, the initial distribution of $W(n)$ is likely to be high for small n and low for large n . This prediction is illustrated by the filled circles in the top right panel of Fig. 16. It follows that during the first FI trials response rate should increase immediately after food and decrease thereafter. The middle left panel of Fig. 16 shows the predicted rate curve for the first FI trial (see curve labeled 1). The time of maximum responding, around 10 s, is independent of the FI duration. These qualitative predictions agree with our empirical findings.

As the FI training proceeds, the coupling between the early states and the operant response decreases because food is not available when these states are maximally active, but the coupling between the later states and the operant response increases because later states are the most active when food occurs. The different symbols in Fig. 16, top right panel, show the predicted changes in $W(n)$ after 50, 200, and 3000 trials (i.e. 1, 4, and 60 50-trial sessions). In terms of response rate profile (middle left panel) the rate during the

initial segment of the interval decreases across sessions while the rate during the last segment increases. In other words, the response rate curve rotates, and at the steady state response rate increases sigmoidally during the interval. The empirical findings agree with these predictions.

As the middle left panel illustrates, LeT predicts that whenever the flow rate of activation across the states (i.e. parameter λ in Appendix A, the equivalent of the pacemaker rate in BeT or SET) does not change appreciably across sessions, the point of rotation of the rate curves also will not change. To understand this prediction, look at the top right panel of Fig. 16. Initially, the early states are more coupled with the operant response than later states, but the reverse is true at the end of training (cf. filled circles and empty diamonds). Hence, by continuity, there is one intermediate state whose initial coupling value already equals, or is close to, its steady-state value. Because that state, $n = 7$ in the figure, is the most active at time $t^* = n/\lambda$, response rate at approximately that time will be mostly determined by the state's coupling value. If this value does not change, response rate at time t^* will not change. Again, this prediction agrees with the data.

The middle right panel in Fig. 16 shows that LeT also predicts that the initial response rate will converge to its steady state faster than the terminal response rate. The prediction requires only that reinforcement have larger effects on the couplings than extinction (i.e. in Appendix A, that $\beta > \alpha$). Although LeT failed to predict any appreciable differences between the convergencies of the response rates at 4, 12, or 20 s into the trial (i.e. the mid point of the first three bins for the FI 80 s birds), it did predict that response rate at $t = 60$ s (i.e. bin 8) would converge faster than at $t = 76$ s (bin 10; compare with Fig. 4, bottom).

LeT's predictions for the early moments of extinction depend on the amount of previous FI training. If we allow the number of trials to go to infinity, as Machado (1997) did, then LeT predicts that $W(n)$ will increase monotonically with n (see Machado (1997), for a proof). This prediction is illustrated by the unfilled circles in the bottom left panel of Fig. 16. Under these conditions, response rate in extinction will remain high for a period of

time much longer than the FI duration (see bottom right panel, curve ' $t \rightarrow \infty$ '). However, if the number of trials is of the same order of magnitude than that in the present experiment (i.e. around 3000 trials, or 60 sessions), then LeT's predictions are remarkably different. In this case, the distribution of $W(n)$ is bitonic (see bottom left panel, filled circles) and response rate during extinction will not remain high for more than twice the duration of the FI (see curve '3000' in the bottom right panel). More important, to predict uninterrupted pecking for intervals ten times greater than the FI, LeT would need to assume an enormous amount of training. As the two remaining curves in the panel show, not even after 100 000 trials (i.e. 2000 sessions) are the model's predictions correct.

The reasons for this scaling problem, which escaped Machado (1997), have to do with the type of dynamics assumed for the behavioral states, namely the Poisson distribution. To see this, consider the problem in the following light. First, for pecking to be sustained at such long intervals after the last reinforcer, the states that are ten or more times further down the chain of states than the state most active during reinforcement need to be strongly coupled with the operant response. In other words, if the most active state during reinforcement is state number x , then the state most active after ten intervals in extinction is state number $10x$. To sustain pecking, state $10x$ and its neighbors must be strongly coupled. Second, to be coupled with the operant response state $10x$ must have been active when reinforcement occurred during training because, according to LeT and BeT, that is the only moment during which coupling increases. Third, and here is the scaling problem, because in both LeT and BeT the state dynamics is Poisson with mean and variance equal to λT (where T is the FI duration), state $10x$ is $z = 9\sqrt{\lambda T}$ standard deviations away from the mean. Hence, for this state to be active at the time of food, λ would have to be so small (e.g. for $T = 80$ s and $z = 4$, $\lambda = 0.0025$) that LeT and BeT would then fail to predict the most elementary aspects of FI performance (e.g. the shape of the steady-state curves). Finally, the scaling problem is compounded when it is realized that some birds

continue to peck for periods longer than ten times the FI duration. We have found no satisfactory solution to the problem.

Finally, LeT also fails to account in a principled way for the development of reliable pause peck oscillations. True, if for some reason the behavioral states are 'reset' to their initial conditions, then LeT predicts a new scallop when responding resumes, but the model is silent on what could reset the states in the absence of any external stimulus change. Although we have explored several possibilities (e.g. the behavioral states are not fully reset after each reinforcer, the animal pauses when it emits a certain number of pecks), none has proved satisfactory. In summary then although LeT predicted the major qualitative findings obtained during phase 1, it failed to predict the main results from phase 2.

The foregoing discussion shows that some current models of temporal regulation have major difficulties in describing what pigeons learn under periodic reinforcement. In particular, no model provides a principled account of pause–peck oscillations in the absence of any stimulus change. To the preceding list containing SET, BeT, and LeT, one should add the diffusion model of Stadon and Higa (1991), for it also fails to predict oscillations during extinction. A radically different type of model would suggest that the behavioral oscillations that we and others have observed (Church and Broadbent, 1990; Broadbent, 1994; Kirkpatrick-Steger et al., 1996) are due to one or more internal oscillators, oscillators that were entrained by the periodic delivery of food. Although the full implications of these models concerning acquisition, for example, remain to be worked out, it is unclear how they would explain why in the study of Kirkpatrick-Steger et al. (1996), the two peaks were not generated simultaneously or, in the present study, why the behavioral oscillations did not occur immediately after the last reinforcer and why the cycle durations changed within a session (Fig. 12).

The shortcomings of the various models of timing may suggest that the phenomena we described concerning acquisition and extinction are not the proper domain of a timing model, for they may depend on variables and processes not di-

rectly related to the animal's timing processes. For example, the shape of the rate curve during the first session of phase 1 may depend mainly on the discriminative function of food (i.e. due to previous training, food cues more food), the pause–peck alternations may depend on the fact that when the bird stops pecking during extinction it re-establishes the stimulus conditions that prevailed during the early moments of the FI trial, stimulus conditions that may function as the initial link of a chain composed by the post-reinforcement pause and the ensuing pecks, and the period of the oscillations observed during extinction may depend on the total number of pecks emitted during the FI trials. To the extent that these variables and the processes they engage are extraneous to the animal's timing mechanisms, the failures of the models are not informative. But it seems to us that to strip the phenomenon of timing of all these 'extraneous' accoutrements is equivalent to dismantling a clock to investigate the essence of time.

Perhaps the most important contribution of the present findings is to invite researchers to reopen the long forgotten chapters on acquisition and extinction under periodic reinforcement. We do not know, for example, if the preceding findings extend to longer intervals or to different extinction procedures, how much training is required to observe pause–peck oscillations, or what sorts of behavioral processes may generate these oscillations. The chapters on acquisition and extinction may be pregnant with deep implications for our current understanding of temporal learning.

Acknowledgements

Parts of this paper were presented at the annual meeting of the Society for the Quantitative Analysis of Behavior, Chicago, 1997. We are grateful to Francisco Silva for his comments.

Appendix A

The following equations describe the dynamics of LeT (for a more detailed account see Machado, 1997)

Behavioral states: The activation of behavioral

state n ($n = 0, 1, \dots$) is a time-dependent variable with the following dynamics:

	$t = 0$	$t > 0$
$n = 0$	$X(0, 0) = 1$	$\frac{d}{dt} X(t, 0) = -\lambda X(t, 0)$
$n > 0$	$X(0, n) = 0$	$\frac{d}{dt} X(t, n) = \lambda X(t, n-1) - \lambda X(t, n)$

where $X(t, n)$ is the activation or strength of state n at time t , and λ is a rate parameter that controls how fast the activation spreads across the states. Time t is measured from the preceding time marker.

Learned associations: Each state is initially associated or coupled to some degree with the operant response; the degree of the association, represented by variable $W(t, n)$, changes in real time, decreasing during extinction and increasing during reinforcement. Specifically:

During extinction (no food)	$\frac{d}{dt} W(t, n) = -\alpha X(t, n) W(t, n)$
During reinforcement (food)	$\frac{d}{dt} W(t, n) = \beta X(t, n) [1 - W(t, n)]$

where α and β are rate parameters, and T is the interval from the preceding time marker to the onset of reinforcement.

Response rule: The strength of the operant response is obtained by adding the cueing function of all states, $W(t, n)$, each multiplied by the activation of the corresponding state, $X(t, n)$. The rate of the operant response (e.g. key pecking) equals:

$$R(t) = A \sum_n X(t, n) W(t, n)$$

where A , a scale parameter, maps the degree of temporal regulation given by the term $\sum_n X(t, n) W(t, n)$ onto measurable response

rate. Parameter A was set to 1 in Fig. 16 because only qualitative predictions were made.

References

- Baron, A., Leinenweber, A., 1994. Molecular and molar analyses of fixed-interval performance. *J. Exp. Anal. Behav.* 61, 11–18.
- Broadbent, H., 1994. Periodic behavior in a random environment. *J. Exp. Psychol.: Anim. Behav. Process.* 20, 156–175.
- Church, R., Broadbent, H., 1990. Alternative representations of time, number, and rate. *Cognition* 37, 55–81.
- Ferster, C., Skinner, B., 1957. *Schedules of Reinforcement*. Appleton-Century-Crofts, New York.
- Gallistel, C.R., 1990. *The Organization of Learning*. MIT, Cambridge.
- Gibbon, J., 1977. Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* 84, 279–325.
- Gibbon, J., 1991. Origins of scalar timing theory. *Learn. Motiv.* 22, 3–38.
- Higa, J.J., Wynne, C.D.L., Staddon, J.E.R., 1991. Dynamics of time discrimination. *J. Exp. Psychol.: Anim. Behav. Process.* 17, 281–291.
- Innis, N.K., Staddon, J.E.R., 1971. Temporal tracking on cyclic-interval reinforcement schedules. *J. Exp. Anal. Behav.* 16, 411–423.
- James, W., 1961. *Psychology: The Briefer Course*. University of Notre Dame Press, Notre Dame, IN. [Originally published in 1892.]
- Kello, J.E., 1972. The reinforcement-omission effect on fixed-interval schedules: frustration or inhibition? *Learn. Motiv.* 8, 138–147.
- Killeen, P.R., 1991. Behavior's time. In: Bower, G.H. (Ed.), *The Psychology of Learning and Motivation*. Academic Press, New York, pp. 295–334.
- Killeen, P., Fetterman, J.G., 1988. A behavioral theory of timing. *Psychol. Rev.* 95, 274–285.
- Kirkpatrick-Steger, K., Miller, S., Betti, C., Wasserman, E., 1996. Cyclic responding by pigeons on the peak timing procedure. *J. Exp. Psychol.: Anim. Behav. Process.* 22, 447–460.
- Machado, A., 1997. Learning the temporal dynamics of behavior. *Psychol. Rev.* 104, 241–265.
- Skinner, B.F., 1938. *The Behavior of Organisms*. Appleton-Century-Crofts, New York.
- Staddon, J.E.R., 1972. Reinforcement omission on temporal go–no-go schedules. *J. Exp. Anal. Behav.* 18, 223–229.
- Staddon, J.E.R., Higa, J.J., 1991. Temporal learning. In: Bower, G.H. (Ed.), *The Psychology of Learning and Motivation*. Academic Press, New York, pp. 265–294.