

*SHIFTS IN THE PSYCHOMETRIC FUNCTION AND  
THEIR IMPLICATIONS FOR  
MODELS OF TIMING*

ARMANDO MACHADO AND PAULO GUILHARDI

INDIANA UNIVERSITY AND UNIVERSITY OF SÃO CARLOS

This study examined how two models of timing, scalar expectancy theory (SET) and learning to time (LeT), conceptualize the learning process in temporal tasks, and then reports two experiments to test these conceptualizations. Pigeons responded on a two-alternative free-operant psychophysical procedure in which responses on the left key were reinforceable during the first two, but not the last two, quarters of a 60-s trial, and responses on the right key were reinforceable during the last two, but not the first two, quarters of the trial. In Experiment 1 three groups of birds experienced a difference in reinforcement rates between the two keys only at the end segments of the trial (i.e., between the first and fourth quarters), only around the middle segments of the trial (i.e., between the second and third quarters), or in both end and middle segments. In Condition 1 the difference in reinforcement rate favored the left key; in Condition 2 it favored the right key. When the reinforcement rates differed in the end segments of the trial, the psychometric function—the proportion of right responses across the trial—did not shift across conditions; when it occurred around the middle of the trial or in both end and middle segments, the psychometric function shifted across conditions. Experiment 2 showed that the psychometric function shifts even when the overall reinforcement rate for the two keys is equal, provided the rates differ around the middle of the trial. This pattern of shifts of the psychometric function is inconsistent with SET. In contrast, LeT provided a good quantitative fit to the data.

*Key words:* models of timing, learning to time, psychometric function, temporal bisection, key peck, pigeon

---

The scalar property has loomed large in recent studies, theories, and discussions of how animals time events. Indeed it has overshadowed the analysis of other equally critical aspects of timing. One of these aspects is how animals learn to regulate their behavior in time, or, equivalently, how current theories conceptualize the learning process in time-based tasks. The goal of the present study was to contrast the learning conceptualizations of two models of timing: scalar expectancy theory (SET; e.g., Gibbon, 1977, 1991) and a derivative of Killeen and Fetterman's (1988) behavioral theory of timing (BeT; see also Killeen, 1991) called learning to time (LeT; Machado, 1997; see also Machado & Cevik, 1998; Machado & Keen, 1999). We start with a description of the two models, paying particular attention to learning-related assump-

tions, and then use the models to try to understand a puzzling result obtained with a free-operant psychophysical procedure. This exercise will enable us to derive different implications from the models, implications that we subsequently tested in two experiments.

The simplest version of SET postulates an internal clock whose structure is represented in the left panel of Figure 1: A pacemaker generates pulses at a high rate; an accumulator adds the pulses emitted during the interval to be timed; and a long-term memory store saves the count obtained at the end of the interval. At the beginning of the interval to be timed, the animal samples a number from its long-term memory and then compares the sampled number with the number currently in the accumulator; the ratio between the two numbers controls, via one or more thresholds, the instrumental (operant) response. In some versions of SET the rate of the pacemaker is assumed to vary across trials; in other versions the rate of the pacemaker remains constant, but a random variable multiplies the number in the accumulator before that number is transferred to long-term memory; in yet other versions the random variable exerts its effect

---

This research was supported by a FIRST award grant from the National Institute of Mental Health to the first author. We thank students Mark Cannon, Erica Gerstman, and Wendy Tai for their help in conducting the experiments.

Address correspondence to Armando Machado, Department of Psychology, Indiana University, 1101 East 10th Street, Bloomington, Indiana 47405 (E-mail: amachado@indiana.edu).

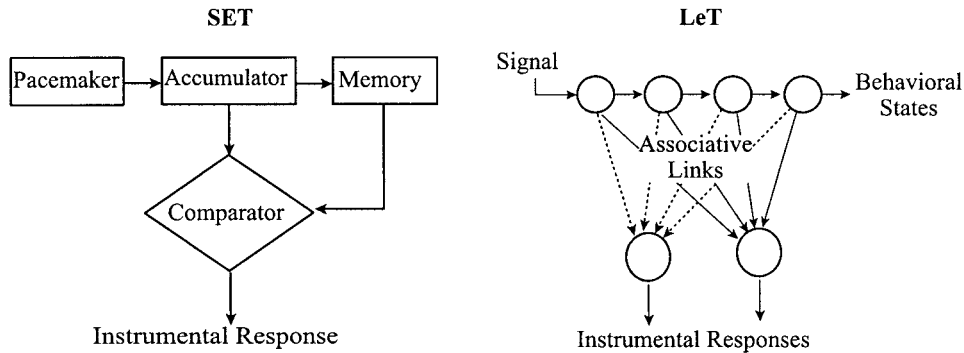


Fig. 1. Left: structure of the scalar expectancy theory (SET) model of timing. A pacemaker generates pulses that are accumulated in the accumulator and stored in long-term memory. A sample extracted from that memory is compared with the number currently in the accumulator, and the ratio between the two numbers determines behavior. Right: structure of the learning to time (LeT) model of timing. After a time marker, a series of behavioral states (top circles) is activated. The states may be coupled to various degrees (middle connections) with one or more operant responses (bottom circles). The strength of each response is determined by the dot product between the vectors of state activation and coupling.

only when the sample is retrieved from long-term memory. For our present purposes these differences are irrelevant because they all yield the same net result: An interval with duration  $t$  is represented by a distribution of values in memory with average and standard deviation proportional to  $t$ .

To deal with more complex situations, the basic version of SET is expanded by increasing the number of memory stores, sampling a number from each store at the beginning of the event to be timed, and comparing the current value in the accumulator with each sampled number. For example, in the bisection procedure (e.g., Catania, 1970; Church & Deluty, 1977; Stubbs, 1968) the choice of one of two alternatives is reinforced following a short stimulus, whereas the choice of the other alternative is reinforced following a long stimulus. SET (e.g., Gibbon, 1981) assumes that the number in the accumulator at the end of the short stimulus is stored in one memory, and the number at the end of the long stimulus is stored in another, distinct memory. The choice at the end of a stimulus is based on a comparison among three numbers: the number in the accumulator at the end of the stimulus, the number sampled from the “short” memory store, and the number sampled from the “long” memory store.

The right panel of Figure 1 shows the three key elements of LeT: a serial organization of behavioral states, a vector of associative links

coupling the behavioral states to one or more instrumental responses, and the instrumental responses themselves. At the onset of the interval to be timed, only the first state is active, but as time elapses the activation of each state flows to the next state in the series. Consequently the most active state changes with time. Each behavioral state is also coupled with one or more operant responses, and the degree of the coupling changes in real time, decreasing during extinction and increasing during reinforcement. States that are strongly active during extinction lose their coupling and eventually may not support the operant response, whereas states that are strongly active during reinforcement increase their coupling and may therefore sustain the response. Finally, the strength of each operant response is obtained by adding the activation strength of the states, each weighted by the coupling value between the corresponding state and the operant response.

From these descriptions we can identify some of the contrasting features of SET’s and LeT’s conceptualizations of learning.

1. In SET, learning consists of storing counts in distinct memory stores, whereas in LeT learning consists of strengthening and weakening connections between distinct behavioral states and the operant responses. Hence for SET the animal’s learning history is embodied in the contents of the memory stores, whereas for LeT it is embodied in the associative links.

2. In SET, the animal's memory stores are differentiated or indexed by structural features of the task (e.g., short and long choice alternatives in the bisection experiment), whereas in LeT, the equivalent of the memory stores is the associative links that are indexed by the behavioral states themselves.

3. Because in SET storage occurs when a reinforcer is delivered, and what is stored is a number that represents the duration of an interval, extinction plays no role in the theory. Therefore, according to SET the contents of the animal's memories will be the same across situations in which the periods of reinforcement availability remain the same but reinforcement rate changes. In contrast, in LeT reinforcement strengthens and extinction weakens the couplings between the active behavioral states and the operant response. Therefore, LeT predicts stronger connections between the states and the operant responses when the reinforcement rate increases, even though the periods of reinforcement remain unchanged.

4. Loosely speaking, in SET a well-trained animal has access to all the relevant information at any time, in particular at the onset of the time marker, because its memories are formed and they can be sampled continuously. In LeT relevant information becomes available only as successive states become active. This difference epitomizes the parallel and serial architectures of SET and LeT, respectively.

To illustrate some of the consequences of these differences (others were explored by Machado & Keen, 1999, and Machado & Cevik, 1998), we will consider a time-based choice situation known as the free-operant psychophysical procedure (FOPP; e.g., Bizo & White, 1994a, 1994b, 1995a, 1995b; Stubbs, 1980). With pigeons the FOPP proceeds as follows: A 50-s trial starts with the illumination of two keylights, left and right. For the first 25 s, only left choices are reinforceable; for the last 25 s, only right choices are reinforceable. During a baseline condition the reinforcers are scheduled according to two independent variable-interval (VI) 60-s schedules. The results of this experiment show that as time into the trial elapses, the proportion of right pecks increases sigmoidally from 0 to 1, with indifference around the middle of the trial at 25 s. This finding is

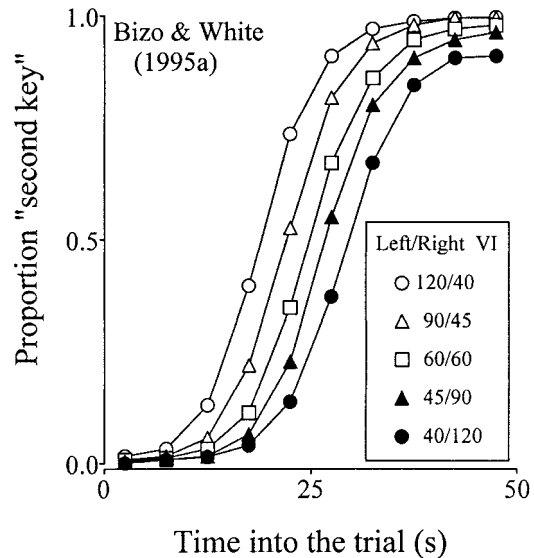


Fig. 2. Psychometric functions obtained by Bizo and White (1995a) with the FOPP. Across five conditions the authors varied the rate of reinforcement associated with the two response keys. When that rate favored the first key (first VI schedule), the psychometric function shifted to the right; when it favored the second key (second VI schedule), it shifted to the left.

illustrated by the open squares in Figure 2 (Bizo & White, 1995a). We refer to this curve as the psychometric function. If the experimenter then changes the reinforcement schedules associated with the two keys to VI 40 s for left and VI 120 s for right, for example, the bird switches to the right key later than during baseline, and the psychometric function shifts to the right. The filled circles in Figure 2 illustrate the result. Conversely, if the schedules change to VI 120 s for left and VI 40 s for right, the animal switches to the right key earlier than in baseline, and the psychometric function shifts to the left (open circles). Furthermore, the magnitude of the shift in the psychometric function seems to depend on the ratio of the two reinforcement rates (see remaining symbols in Figure 2).

How do SET and LeT account for this result? For SET the reinforcers received from the left key at time  $t$  will engender a distribution of counts in memory with a mean of  $c_1 t$  and a standard deviation of  $c_2 t$ . We seek the distribution that results when the effects of all reinforcers are taken into account, that is, when  $t$  varies from 0 s to 25 s. Figure 3 helps to derive this resulting distribution.

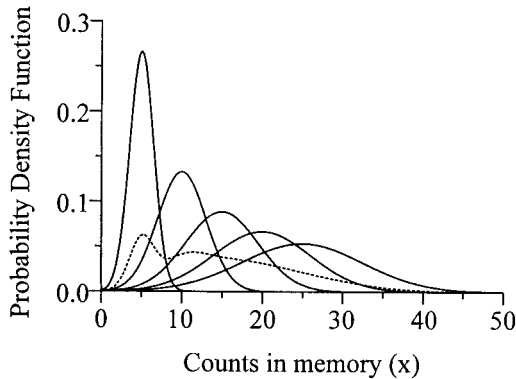


Fig. 3. Example to illustrate the memory contents predicted by SET. A reinforcer received at time  $t$  engenders a distribution of counts in memory with a mean of  $c_1 t$  and a standard deviation of  $c_2 t$  (in the example,  $c_1 = 1$  and  $c_2 = 0.3$ ). Assume that reinforcers were delivered only at 5 s, 10 s, 15 s, 20 s, and 25 s, and with the same probability at each time. The five solid curves show the distributions engendered by each set of reinforcers, and the dashed curve shows the distribution that results when the effects of all reinforcers are taken into account.

Take  $c_1 = 1$  and  $c_2 = 0.3$ , and consider for the moment only the set of reinforcers delivered at time  $t = 5$  s. The distribution of counts engendered by these reinforcers is represented by a normal density function with a mean of 5 and a standard deviation of 1.5. This density, which we denote by  $N(x, 5, 1.5)$ , corresponds to the leftmost solid curve in Figure 3. But reinforcers also occur at other times, and each of these sets of reinforcers will engender its own distribution. For example, the reinforcers delivered at times  $t = 10$  s, 15 s, 20 s, and 25 s will engender the distributions represented by the remaining solid lines in Figure 3. When *all* reinforcers are taken into account, a final or resulting distribution is obtained. The density function of this resulting distribution is a weighted average of the component densities, a weighted average in which the weights reflect the proportion of the reinforcers received at the various times. Thus if all reinforcers occurred at 5 s, 10 s, 15 s, 20 s, or 25 s, and if they were evenly distributed among these five times, the resulting density would be the dashed curve in Figure 3, which is simply the arithmetic mean of the five solid curves.

The FOPP is similar to the preceding example except that reinforcers from the left key can occur at all times between 0 s and 25 s. In fact, because the reinforcers are set up

according to a VI schedule, the probability density function of the reinforced times in the FOPP is also rectangular. Hence, the resulting density associated with the left key,  $f_L(x)$ , will equal the average of all contributing densities:

$$f_L(x) = \frac{1}{25} \int_0^{25} N(x, c_1 t, c_2 t) dt.$$

A similar expression, but with the integration taken from  $t = 25$  s to  $t = 50$  s, holds for the right key.

The two resulting densities correctly predict a preference for the left key at the beginning of the trial, indifference around the middle of the trial, and a preference for the right key at the end of the trial. Note, however, that the expression for each density does not include the schedule parameter, the reinforcement rate on the left or right key. Whether the schedule on the left key is a VI 10-s or a VI 1,000-s schedule, the resulting density is the same because the distribution of reinforcement times does not change. Hence, according to SET the psychometric functions should not shift. SET fails to predict the experimental findings shown in Figure 2 because its conception of learning as storing counts in memory when a reinforcer occurs makes the model insensitive to changes in reinforcement rate that are not accompanied by changes in the distribution of reinforcement times.

If SET cannot predict the shifts of the psychometric function on the basis of the contents of the animal's memory, perhaps it can predict them on the basis of some other feature, one that is directly related to reinforcement rate. One such feature may be the response thresholds typically included in the model's decision rules. Perhaps, so the argument could go, when Bizo and White (1995a) made the left key richer than the right key, the threshold for pecking the left key decreased with respect to the threshold for pecking the right key, and the animal chose the left key more often than before. This increased preference for the left key shifted the psychometric function to the right. Conversely, when they made the right key richer, the response thresholds changed in the opposite direction, the animal preferred the right key more than before, and the psychometric

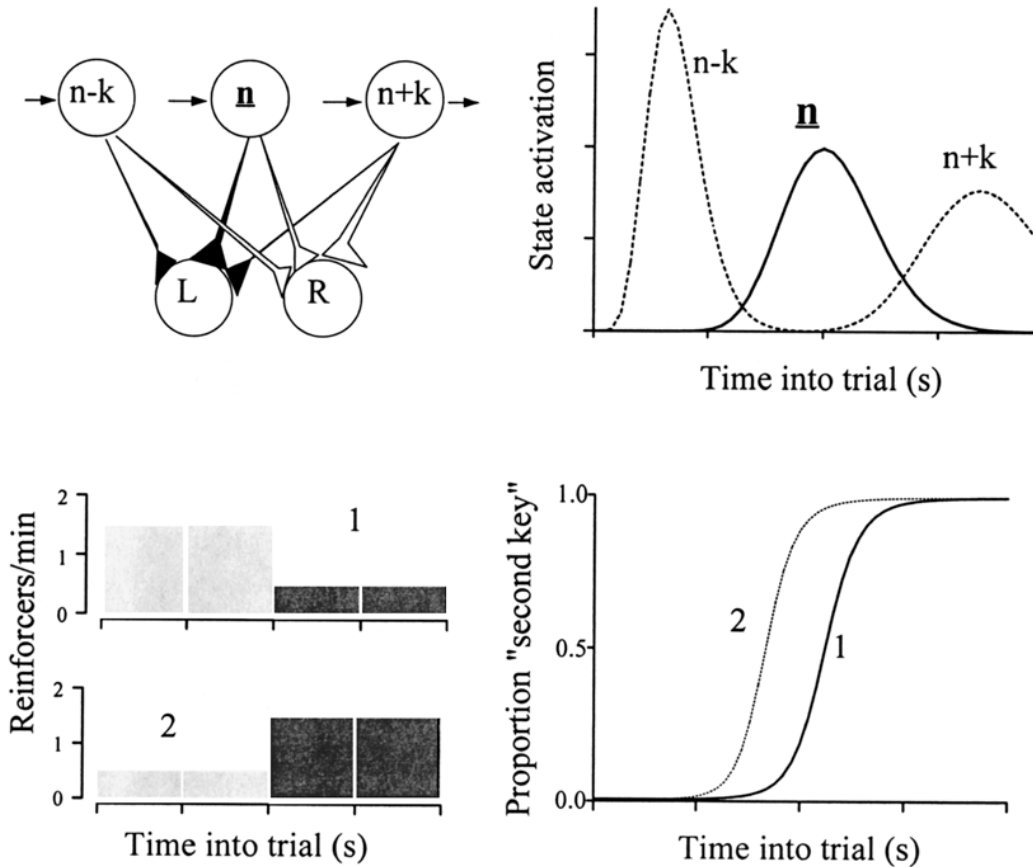


Fig. 4. LeT's account of the shifts of the psychometric function in the FOPP. Top left: the three states represent the states that are most active early in the trial ( $n - k$ ), around the middle of the trial ( $n$ ), and late in the trial ( $n + k$ ). Top right: activation of the states during the trial. Bottom left: two conditions (1 and 2) of an FOPP. The bars represent the reinforcement rate on the left (gray) and right (black) keys. Bottom right: LeT's predicted psychometric functions for the two conditions.

function shifted to the left. Thus, changing the relative reinforcement rates while preserving the distribution of reinforcement times induces threshold changes that in turn cause the shifts of the psychometric function. Henceforth, we refer to this hypothesis as the *threshold-based account* and will return to it after we have described how LeT explains the same findings.

For LeT the shifts of the psychometric function depend on the coupling asymmetries of the states that are most active during the middle of the trial. Figure 4 details the account at an intuitive level; we postpone until the General Discussion a more rigorous but less intuitive mathematical account. The top panels show three states (left) and their time course of activation (right). The first

state represents the states that are maximally active early in the trial; the second represents the states that are maximally active during the middle of the trial; and the third represents the states that are maximally active late in the trial. Because the maximally active state changes with time, each state has, as it were, a ruling epoch, an interval during which it affects responding more than the other states do. In Figure 4 we designate by  $n - k$  the states whose ruling epoch occurs at  $t \ll 25$  s; these states control (via their couplings) the operant responses during the beginning of the trial. We designate by  $n$  the states whose ruling epoch occurs around  $t = 25$  s; these states control the operant response around the middle of the trial. And we designate by  $n + k$  the states whose ruling epoch occurs at



$t \gg 25$  s and which control the operant responses during the end of the trial. These three states are singled out only to make LeT's account more intuitive and its description less cumbersome. The reader is asked to bear in mind that whenever a statement is made about state  $n$ , for example, it applies also to its neighboring states  $n - 1$ ,  $n + 1$ ,  $n - 2$ ,  $n + 2$ , and so on, but less and less so as their distance from state  $n$  increases.

During the baseline condition the two VIs are equal. Therefore, while state  $n$  is active (see the solid curve in the top right panel of Figure 4), reinforcers follow left and right pecks with approximately the same frequency. For this reason state  $n$  will become coupled equally with the two responses, and when this state is the most active the animal will be roughly indifferent between the two keys. The earlier states, represented by state  $n - k$ , will become coupled more with the left than the right response because while they are active reinforcers follow more left than right key pecks. When these states are the most active, more left pecks occur. For similar reasons the late states, represented by state  $n + k$ , will become coupled more with the right than the left response, and when they are the most active the animal pecks more at the right key. Hence, LeT correctly predicts a preference for the left key at the beginning of the trial because the early states are coupled mainly with the left response, indifference around 25 s because state  $n$  is coupled equally with the two responses, and a preference for the right key at the end of the trial because the late states are coupled mainly with the right responses.

When the VI schedules differ and, say, the schedule associated with the left key is richer (see the bottom left panel of Figure 4, Condition 1) state  $n$  becomes coupled more with the left than the right key because when this state is active more reinforcers follow left than right key pecks. At 25 s into the trial, when state  $n$  has its ruling epoch and exerts greater control over responding than the other states do, the animal continues to prefer the left key. Hence, LeT correctly predicts that the psychometric function will shift to the right. This prediction is illustrated by the solid curve in the bottom right panel of Figure 4. Conversely, and for similar reasons, when the VI for the right key is richer (left

panel, Condition 2), state  $n$  becomes coupled more with the right than the left key, and the psychometric function will shift to the left, as the dotted curve in the right panel shows. We will show later that LeT also predicts that the magnitude of the shift increases with the ratio between the two VI schedules, as Bizo and White's (1995a) findings suggest.

In summary, LeT provides a straightforward account of the shifts of the psychometric function: These shifts are caused by differences in reinforcement rate between the two keys, provided that these differences surround the middle of the trial (i.e., 25 s). This account is to be contrasted with the threshold-based account presented above. The difference is that whereas a threshold change implies a global, time-independent effect of the differences in reinforcement rate, LeT's account implies a local, time-dependent effect of these differences. For if the same differences in reinforcement rate were to occur far from the middle of the trial, then according to LeT either a much smaller shift or no shift at all should be observed, whereas according to the threshold-based account the same shift should happen. The two experiments reported below tested these predictions.

## EXPERIMENT 1

To introduce the logic of the experiment and to further clarify the difference between the two accounts, consider the following task. Each trial of an FOPP is divided into four equal periods. As before, during the first two only left pecks are reinforceable and during the last two only right pecks are reinforceable. The experimenter manipulates the reinforcement rates across the four trial periods. We distinguish three comparisons and illustrate them in the three left panels of Figure 5. Each comparison contains two experimental conditions. In Condition 1 the overall reinforcement rate on the left key is higher than on the right key (the area under the light bars is greater than the area under the dark bars); in Condition 2 the opposite is true. However, the overall reinforcement rate obtained from both keys is always equal (the total area under the light and dark bars always equals 4). The top panel reproduces the comparison analyzed above. Here, both LeT

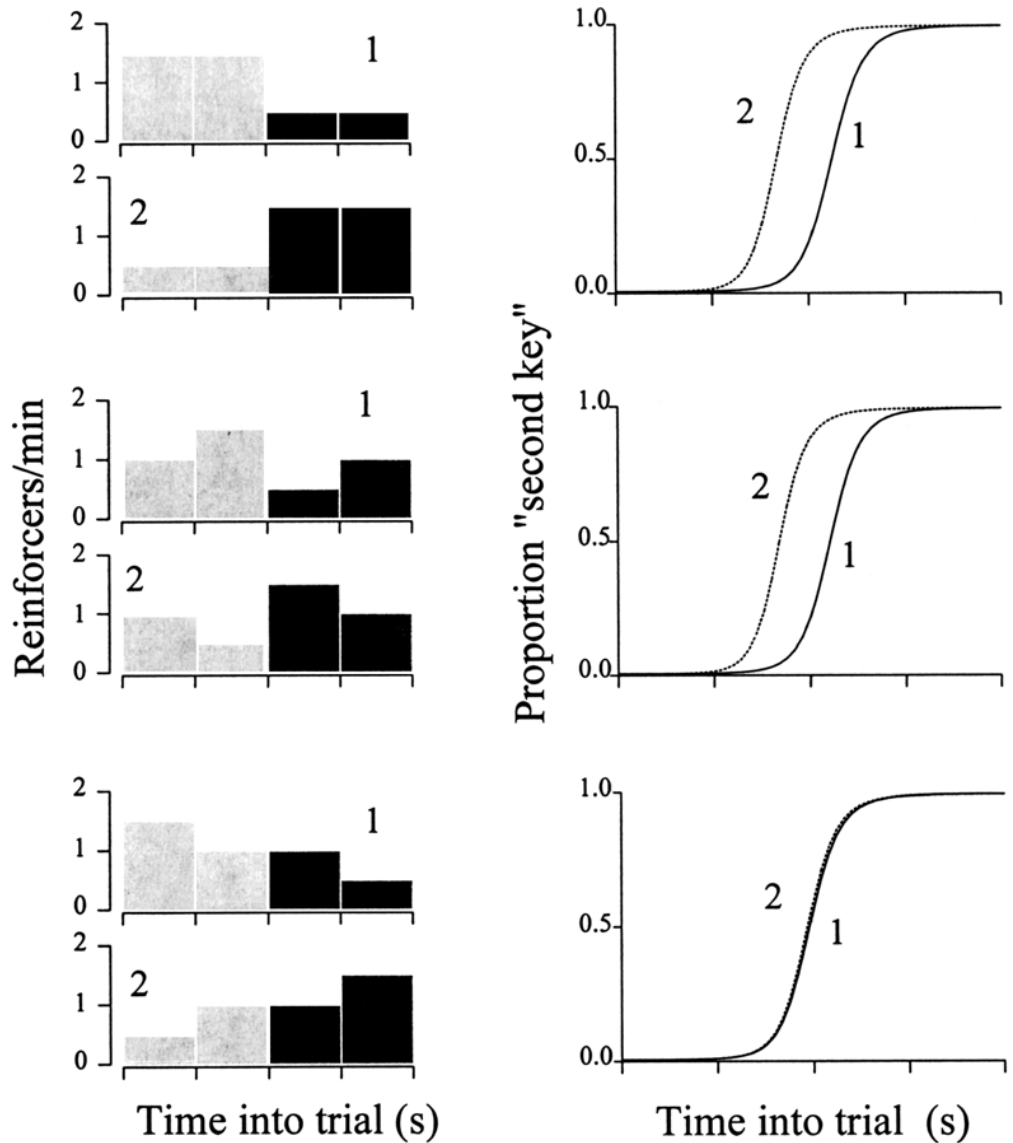


Fig. 5. Left panels: three pairs of FOPPs, each comparing two experimental conditions (1 and 2). In the top panel the reinforcement rates differ both around the middle and at the extremes of the trial. In the middle panel they differ around the middle but not at the extremes. In the bottom panel they differ at the extremes but not at the middle. Right panels: LeT's predicted psychometric functions. All curves used the same parameters.

and the threshold-based account predict a shift in the psychometric function with respect to baseline, a shift to the right in Condition 1 and to the left in Condition 2. The middle and bottom panels show two comparisons equal in all respects except that in one (middle panel) the differences in reinforcement rate between the two keys occur during the middle of the trial, whereas in the other

(bottom panel) the differences occur at the extremes of the trial. The threshold-based account predicts that the psychometric function will shift to the right in Condition 1 of both comparisons because the reinforcement rate on the left key is higher than the reinforcement rate on the right key; similarly, during Condition 2 the psychometric function will shift to the left because the reinforcement

rate on the right key is higher than that on the left key. In short, the threshold-based account predicts the same outcome for the middle and bottom comparisons.

In contrast, LeT predicts large shifts when the reinforcement rates differ around the middle of the trial (middle), but small or no shifts when they differ only at the extremes of the trial (bottom). These predictions are illustrated by the curves in the right panels. The shifts in the middle panel are predicted for the same reasons that shifts were predicted in the top panel: In both conditions, early and late states become coupled chiefly with the left and right responses, respectively, because most reinforcers will come from the left key while the early states are active and from the right key while the late states are active; state  $n$ , however, will be strongly coupled with the left response in Condition 1, but will be coupled with the right response in Condition 2; hence the shifts. The bottom comparison is different, though, not with respect to the early and late states but with respect to state  $n$ . When state  $n$  is active the reinforcers come equally often from the left and right keys, and therefore this state remains coupled equally with the two responses; hence the absence of shifts.

In summary, for LeT the comparisons represented in the top and middle panels of Figure 5 are essentially equal, but they differ significantly from the comparison represented in the bottom panel. The latter is in fact similar to baseline, notwithstanding the difference in relative reinforcement rates between the two keys. This feature distinguishes LeT from the threshold-based account, according to which all three comparisons represented in Figure 5 are essentially equal but are different from baseline.

Experiment 1 used three groups of pigeons to test these predictions. One group was exposed to the conditions depicted in the top panel of Figure 5. Because for this group the difference in reinforcement rates held throughout the trial, we call it Group TOT, for total. The other groups were exposed to the conditions depicted in the middle and bottom panels. We call them Group MID and Group EXT to emphasize that one experienced different reinforcement rates during the middle of the trial and the other experi-

enced different reinforcement rates during the extremes of the trial.

## METHOD

### *Subjects*

Eighteen experienced pigeons (*Columba livia*) participated in the experiment. The birds were housed individually and had free access to water and grit when not in the experimental chamber. A 12:12 hr light/dark cycle was in effect in the pigeon colony. Throughout the experiment, the birds were maintained at 80% of their free-feeding body weights.

### *Apparatus*

Three standard experimental chambers for pigeons (two from Lehigh Valley and one from Med Associates) were used. The front panel of each chamber contained three keys centered on the wall, 2.5 cm in diameter, 21 cm above the floor, and 8 cm apart center to center. The center key was not used during the experiment and always remained dark. The two side keys could be illuminated from behind with red light. In the Lehigh Valley chambers the hopper opening measured 5 cm by 6 cm and was centered on the wall, 7.5 cm above the floor grid; in the chamber from Med Associates the corresponding measurements were 7 cm by 6 cm and 10.5 cm. The bird had access to mixed grain when the hopper was raised and illuminated with a 7.5-W white light. On the back wall of the chamber another 7.5-W light provided general illumination. The experimental chamber from Med Associates and the intelligence panels from Lehigh Valley were enclosed by outer boxes equipped with ventilating fans. Three Dell 386 computers controlled the experimental events and recorded the data.

### *Procedure*

The 18 birds were assigned randomly to the three chambers, with 6 birds per chamber. The following procedural features remained constant throughout the experiment: Sessions were conducted daily at approximately the same time for each pigeon; each session was divided into trials and the trials were separated by 10-s blackout periods (intertrial interval); a peck to an illuminated key turned that key-light off for 50 ms; when a reinforcer was delivered all keylights and the houselight were



turned off, the hopper was raised, and its light turned on, for 2.5 s.

*Preliminary training.* During the first half of each 40-s trial, only one key was illuminated; during the second half of the trial, only the other key was illuminated. Three randomly selected birds in each chamber started the trial with one key; 3 started with the other. However, to facilitate the description of the procedure we will assume that for all birds the keys were illuminated in the order left first, right second. Pecks at the illuminated key were reinforced with a probability of .5, and a session ended when 40 reinforcers had been delivered. The trial duration was then increased to 60 s, the reinforcement probabilities were reduced to .2, and then were replaced by two VI 15-s schedules. The VI schedules were independent, and their timers ran only during the corresponding half of the trial. When a reinforcer was set up by a VI schedule, say the VI associated with the left key, its timer stopped until the bird pecked the left key during the first half of the trial. Hence, a reinforcer set up but not collected during one trial would be collected during subsequent trials.

In the final phase of this preliminary training, each 60-s trial began with the illumination of the houselight and both keylights. From then until the end of the experiment, the bird could always choose which key to peck. The VI schedules were increased gradually from 15 s to 50 s. A session ended when either 40 reinforcers had been delivered (and the current trial ended) or when 60 trials had been completed. This preliminary training lasted from 8 to 10 sessions.

*FOPP training.* All birds were first exposed to a baseline condition and then to Conditions 1 and 2. Table 1 shows the details. In the baseline condition, pecks at one key during the first 30 s of the trial were reinforced according to a VI 60-s schedule, and pecks at the other key during the last 30 s of the trial were reinforced according to another VI 60-s schedule. After about 55 sessions a 2-s changeover delay (COD) was introduced to reduce the high rate of switching between the keys observed in some birds. The COD remained in effect in all subsequent sessions.

After baseline training, the 18 birds were divided into three groups of 6 birds each. An attempt was made to balance the groups in

Table 1

Number of sessions during each condition of Experiment 1.

Group	Bird	Baseline	Condition 1	Condition 2
		60/60	40-40/120-120	120-120/40-40
TOT	4161	69	33	27
	4138	69	26	27
	1346	69	26	26
	5841	69	27	25
	1022	68	26	25
	2186	76	26	25
		60/60	60-40/120-60	60-120/40-60
MID	4171	69	25	26
	4170	76	25	25
	3856	68	33	25
	4913	69	25	26
	5259	69	25	26
	5320	69	25	26
		60/60	40-60/60-120	120-60/60-40
EXT	9854	69	26	28
	9825	69	34	25
	9805	69	25	26
	469	69	37	26
	2775	69	27	18
	3860	75	25	25

terms of baseline performance and yet retain (a) 2 birds per group in each chamber and (b) 3 birds per group starting the trials with the left key as the reinforceable key. For all groups the trial was divided into four periods of 15 s, and four independent VI timers controlled reinforcement rates during these periods. In Table 1 each condition is identified by the VI schedules; the first two (or only one during baseline) operated on the left key and the last two operated on the right key. For example, for Group MID, Condition 1, the VIs were 60-40/120-60. This means that pecks on the left key were reinforced according to a VI 60-s schedule during the first 15-s period of the trial and according to a VI 40-s schedule during the second 15-s period; pecks on the right key were reinforced according to a VI 120-s schedule during the third 15-s period and according to a VI 60-s schedule during the fourth 15-s period. Compared with Condition 1, Condition 2 simply reversed the order of reinforcement rates on the two keys. The remaining cells of Table 1 show the number of sessions per condition. Each condition remained in effect for a minimum of 25 sessions and until the psychometric function of the last 5 days did not change appreciably from the function of the preceding 5 days. Bird 2775 developed a tumor after 18 sessions

in Condition 2 and had to be removed from the experiment.

We restricted all data analyses to the last 10 sessions of Conditions 1 and 2 because the results from baseline do not bear on the major issues under examination and, in addition, because they were similar across groups: The average psychometric functions overlapped considerably, and the average bisection points (i.e., the times at which the psychometric functions crossed .5) did not differ significantly among the groups,  $F(2, 15) = 0.97$ ; these averages were 27.6 s, 25.5 s, and 25.0 s for Groups TOT, MID, and EXT, respectively. In statistical tests a  $p$  value less than .05 was considered a significant result.

#### RESULTS AND DISCUSSION

Figure 6 shows the average reinforcement rate obtained by each group during the four trial periods. In all cases, the obtained reinforcement rate was close to the scheduled reinforcement rate. In particular, for Group MID there were large differences in reinforcement rate around the middle of the trial (cf. second and third 15-s periods) but not at the extremes (cf. first and fourth 15-s periods), whereas for Group EXT the converse was true: There were large differences at the extremes but not around the middle of the trial. Furthermore, except for the fact that the assignment of the VIs to the two keys was reversed, the reinforcement rates during Conditions 1 and 2 were also similar.

A three-way analysis of variance (ANOVA) on the obtained reinforcement rates, with group as a between-subjects factor and condition and key as within-subject factors, yielded nonsignificant effects of group,  $F(2, 15) = 3.1$ ,  $p = .07$ ; condition,  $F(1, 5) = 1.24$ ,  $p = .28$ ; and key,  $F(1, 15) = 1.89$ ,  $p = .19$ . The interactions were also nonsignificant except in the two predicted cases: Condition  $\times$  Key,  $F(1, 15) = 2,876$ ,  $p < .001$ , because the left key was richer during Condition 1 but poorer during Condition 2, and Group  $\times$  Condition  $\times$  Key,  $F(2, 15) = 155$ ,  $p < .001$ , because the interaction effect just stated differed between Group TOT and the other two groups (i.e., there were two VI 40-s and two VI 120-s schedules in Group TOT but only one VI 40-s and one VI 120-s schedule in Groups MID and EXT). We conclude that the birds experienced the intended reinforcement rates.

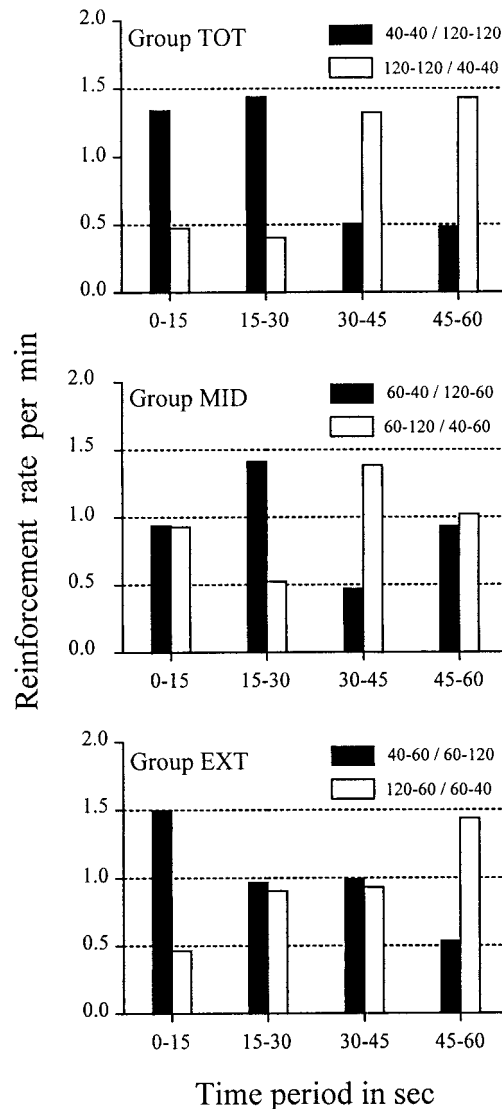


Fig. 6. Average reinforcement rate obtained across the four trial periods during Conditions 1 (dark bars) and 2 (light bars) of Experiment 1. The dotted lines correspond to the maximum rates predicted by a VI 40-s schedule (1.5 reinforcers per minute), a VI 60-s schedule (1 reinforcer per minute), and a VI 120-s schedule (0.5 reinforcer per minute).

The next three figures show the psychometric functions for each bird. To obtain these functions we divided the trials into 12 5-s periods, determined the total number of left and right responses during each period of each session, added the results from the last 10 sessions, and then computed the proportion of right responses in each period. For

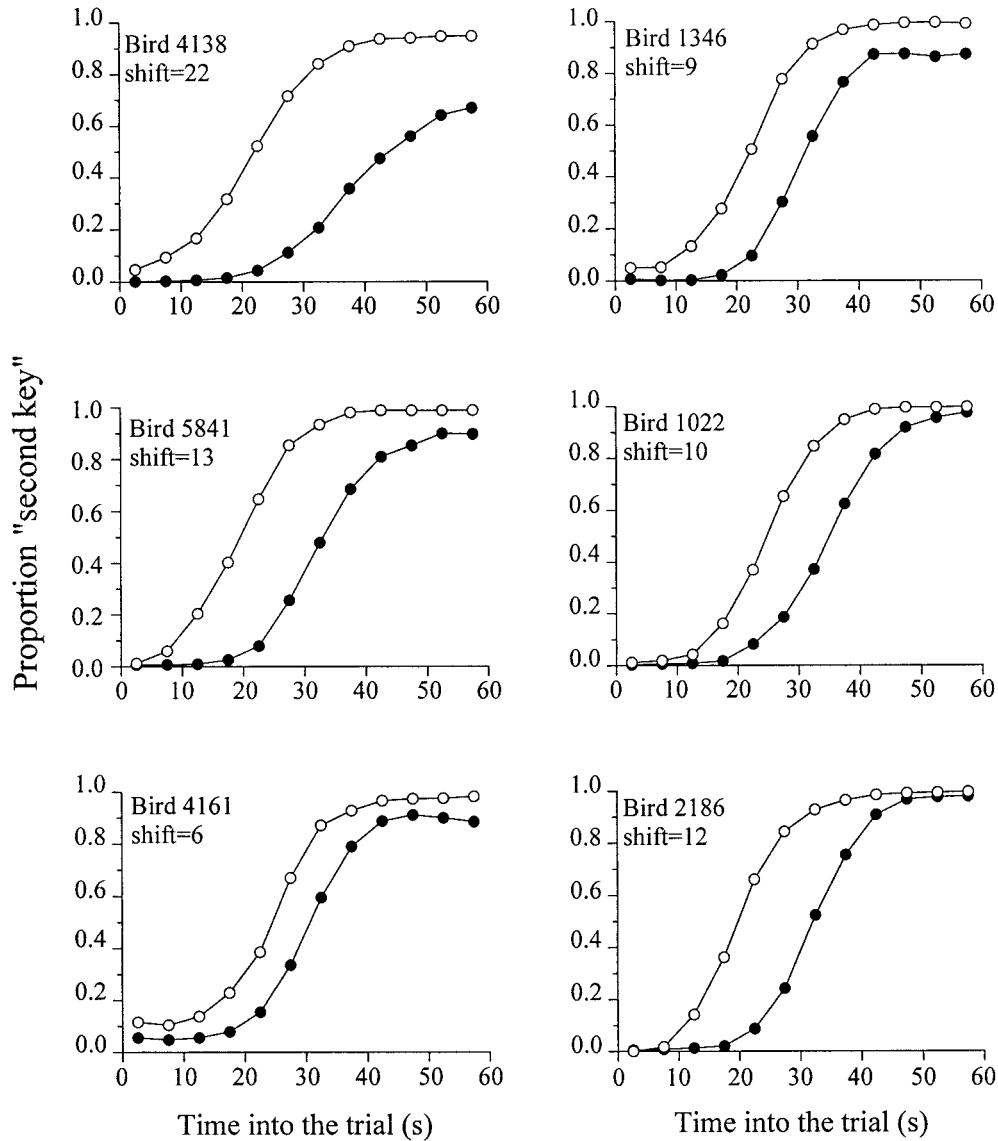


Fig. 7. Individual psychometric functions during Conditions 1 (filled circles) and 2 (open circles) from Group TOT of Experiment 1. The shift is measured from the bisection points of the curves.

each function we also determined by linear interpolation the time at which the two response proportions were equal—the bisection point. The difference between the bisection points of the functions for Conditions 1 and 2 defined the magnitude of the shift.

Figure 7 presents the results for Group TOT. All birds showed the typical ogive, with proportion of right increasing with time into the trial. More important, they also showed a large shift between conditions, in that the

curve for Condition 2 was always to the left of the curve for Condition 1. The magnitude of the shift ranged from 6 s to 22 s (mean = 12 s, median = 11 s). These results replicate Bizo and White's (1995a) findings displayed in Figure 2. The proportions at the beginning and end of the trial showed the following patterns: In 4 pigeons the proportion at the end of the trial was below .95 (Birds 4138, 1346, 5841, and 4161, Condition 1); in 3 pigeons the proportions at the beginning of the trial

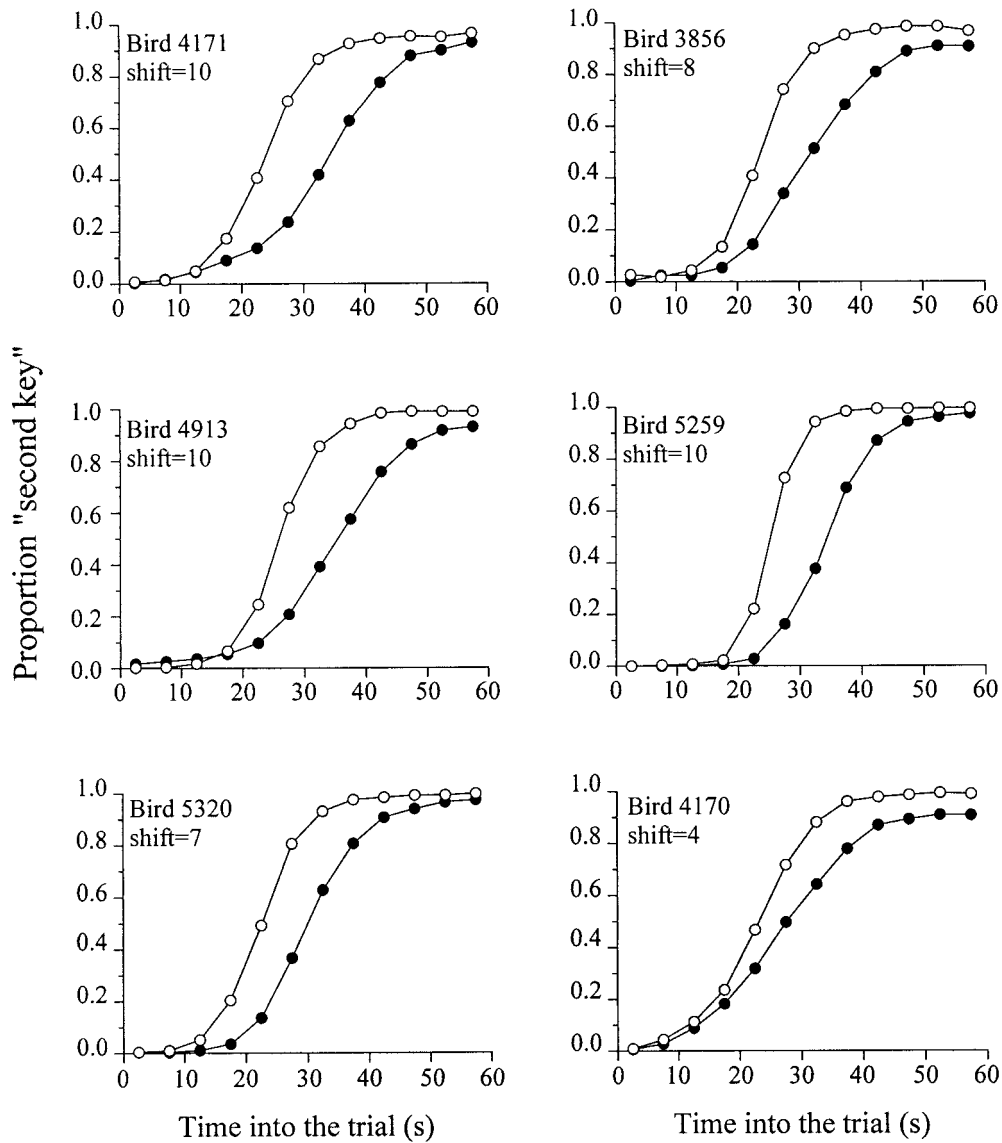


Fig. 8. Individual psychometric functions during Conditions 1 (filled circles) and 2 (open circles) from Group MID of Experiment 1. The shift is measured from the bisection points of the curves.

were above .05 (Bird 4161, Conditions 1 and 2, and Bird 1346, Condition 2); in all other pigeons the proportions at the beginning and end of the trial were very close to 0 and 1, respectively.

Figure 8 presents the results for Group MID. All birds showed the typical ogive and, as for Group TOT, the curve for Condition 2 was also always to the left of the curve for Condition 1. The magnitude of the shifts of the psychometric function ranged from 4 s to

10 s (mean = 8 s, median = 9 s), values that were smaller than those for Group TOT. Most curves started at 0 and ended at 1. The three exceptions occurred during Condition 1: The curves of Birds 4171, 3856, 4913, and 4170 ended below .95.

Figure 9 presents the results for Group EXT. The birds showed the typical ogive, but the two curves either overlapped considerably or remained in close proximity. The magnitude of the shifts ranged from -2 s to

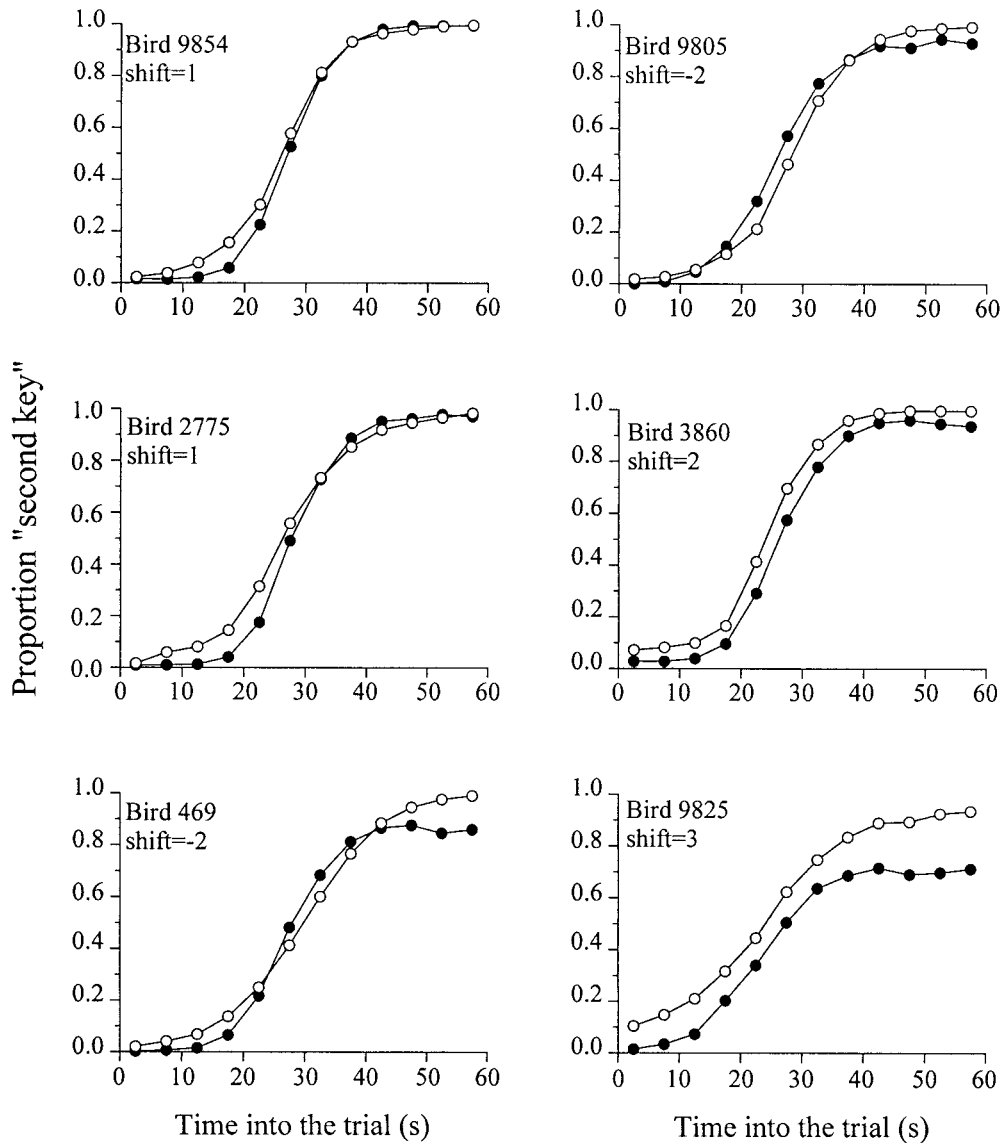


Fig. 9. Individual psychometric functions during Conditions 1 (filled circles) and 2 (open circles) from Group EXT of Experiment 1. The shift is measured from the bisection points of the curves.

3 s (mean = 0.5 s, median = 1 s), values that were considerably smaller than those for Groups TOT and MID. During Condition 1, 4 birds (9805, 3860, 469, and 9825) showed proportions at the end of the trial below .95; during Condition 2, Birds 3860 and 9825 showed a proportion at the beginning of the trial above .05.

In summary, the psychometric functions shifted significantly in Groups TOT and MID but not in Group EXT. An ANOVA on the

magnitude of the shifts yielded a significant effect,  $F(2, 15) = 15.4, p = .0002$ , and Tukey's post hoc test revealed that Group EXT was significantly different from the other two groups, but these did not differ from each other. This pattern of results is consistent with LeT because the two groups for which LeT predicted a substantial shift showed the shift, and the group for which LeT predicted no shift did not show it. In contrast, this pattern of results is inconsistent with a threshold-



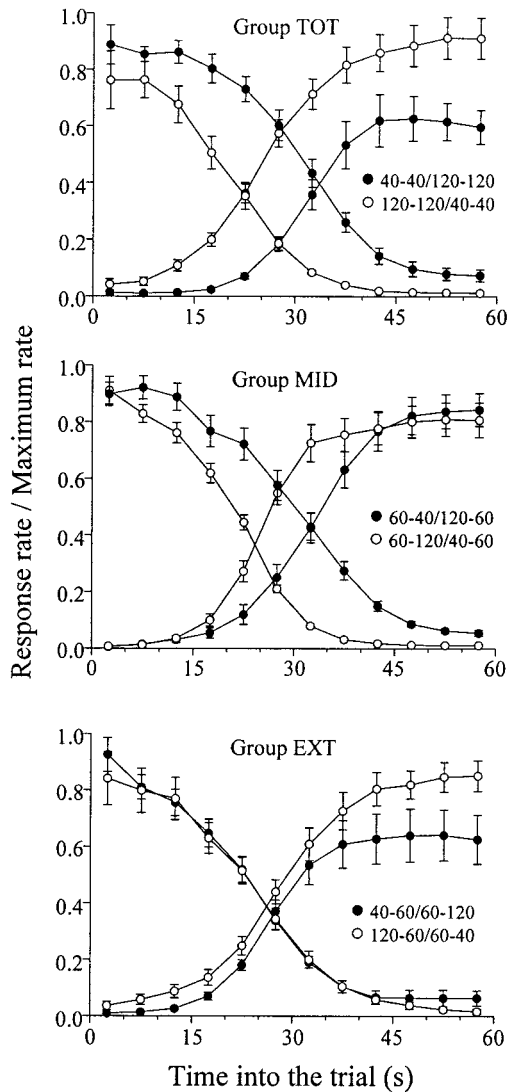


Fig. 10. Average normalized response rate across the trial for Conditions 1 (filled circles) and 2 (open circles) of Experiment 1. The descending curves are for the left key; the ascending curves for the right key. The bars represent the standard error of the mean.

based account because, according to this account, all groups should have shown similar shifts.

Figure 10 shows the average of the normalized response-rate curves for the three groups. Each curve was obtained as follows. First, we normalized the left and right rate curves of each session by dividing the 12 values along each curve by the highest value of the two curves (i.e., the highest of the 24 val-

ues). This transformation corrects for variation in absolute response rates across sessions and birds, and it reveals more clearly the pattern of change in response rate across the trial. Next, we averaged the normalized rate curves of the last 10 sessions. Finally, we averaged the individual curves to obtain the group curve. As might have been expected, the results support the conclusions reached above. For Group TOT the curves for Condition 1 intersected at approximately 34 s into trial, whereas the curves for Condition 2 intersected at 22.5 s. This difference is consistent with the average 12-s shift observed in the psychometric functions. For Group MID the intersections occurred at approximately 32.5 s (Condition 1) and 24 s (Condition 2), also consistent with the 8-s average shift of the psychometric functions. For Group EXT the intersections occurred at about 27 s (Condition 1) and 26 s (Condition 2), consistent with the 0.5-s average shift of the psychometric functions.

The rate curves shed further light on the reasons for the different degrees of shift of the psychometric functions. For Condition 1, response rate on the left key decreased and response rate on the right key increased slightly more rapidly in Group EXT than in Groups TOT and MID. Condition 2 showed the opposite finding: Response rate on the left key decreased and response rate on the right key increased more slowly in Group EXT than in the other groups. That is to say, when the reinforcement rates around the middle of the trial differed, responding on the richer key either decreased at a lower rate (Condition 1, left key) or increased at a higher rate (Condition 2, right key), whereas responding on the poorer key either increased at a lower rate (Condition 1, right key) or decreased at a faster rate (Condition 2, left key). The combination of these effects yielded the different amounts of shift in the psychometric functions.

## EXPERIMENT 2

The threshold-based account cannot explain why the psychometric function shifted in Group MID but not in Group EXT of Experiment 1. For both groups the overall reinforcement rate (L + R) was one reinforcer per minute, and the difference between the

reinforcement rates delivered by the two keys ( $|L - R|$ ) was 0.25 reinforcers per minute. If the response thresholds depended only on relative reinforcement rate, then the two groups should have displayed similar shifts, but they did not. This conclusion would be much stronger if the shifts of the psychometric function had been obtained with equal reinforcement rates on the two keys. If response thresholds are a function of the difference in reinforcement rates between the keys, then eliminating this difference should eliminate the difference in thresholds. However, for LeT the key issue is not the difference in reinforcement rates per se but when the difference occurs. If the reinforcement rates differ around the middle of the trial, then LeT predicts a shift in the psychometric function, even if the keys deliver the same overall number of reinforcers per minute.

Consider then the following situation. One group of pigeons is exposed to a schedule similar to that of Group MID in Experiment 1, but with a slight difference: The four VI schedules are now 120-40/120-40 during Condition 1 and 40-120/40-120 during Condition 2. That is, during Condition 1 left choices are reinforced according to a VI 120-s schedule during the first period and a VI 40-s schedule during the second period; right choices are reinforced according to a VI 120-s schedule during the third period and a VI 40-s schedule during the fourth period. In Condition 2 the order of the VIs for each key is reversed. Both keys deliver 0.5 reinforcers per minute during Conditions 1 and 2. Because there is no difference in overall reinforcement rates, the response thresholds should not change. Hence, according to the threshold-based account the psychometric function should not shift. For LeT, however, during Condition 1 state  $n$  will become coupled more strongly with the left than the right choice, but during Condition 2 the opposite will happen. Hence the psychometric function should shift.

Consider now a group exposed to a slight variation of the contingencies used for Group EXT in Experiment 1. The VI schedules are 120-40/40-120 during Condition 1 and 40-120/120-40 during Condition 2. The reinforcement rates remain the same, 0.5 per minute, but LeT predicts no asymmetry in

the couplings of state  $n$  because the reinforcement rates are equal around the middle of the trial. The psychometric function should not shift. In summary, whereas LeT predicts a shift in one but not in the other group, the threshold-based account predicts no shift in either group. Experiment 2 tested these predictions.

## METHOD

### *Subjects and Apparatus*

Fourteen of the 18 pigeons used in Experiment 1 served as subjects. The experimental chambers were also the same.

### *Procedure*

Experiment 2 followed Experiment 1 immediately. The birds were divided into two groups, one with 9 birds and the other with 5 (initially this group had 6 birds, but one became sick and was removed from the experiment). To balance the groups with respect to the birds' experimental histories, each group included birds from all three groups of Experiment 1. Except for the changes in the reinforcement schedules, all experimental details remained as in Experiment 1.

The larger group was exposed to the VI schedules 120-40/120-40 during Condition 1 and 40-120/40-120 during Condition 2. We designated this Group DIF because the scheduled reinforcement rates differed around the middle of the trial. The smaller group was exposed to the schedules 40-120/120-40 during Condition 1 and 120-40/40-120 during Condition 2. The group was designated EQU because the scheduled reinforcement rates were equal during the middle of the trial. Each condition lasted for a minimum of 25 sessions and until the psychometric function of the last 5 days did not differ appreciably from the function of the preceding 5 days (the actual number of sessions ranged from 25 to 40). All data analyses were based on the last 10 sessions of each condition.

## RESULTS AND DISCUSSION

Figure 11 shows the average reinforcement rates for each group during the four trial periods. The obtained values were close to the scheduled ones, but there was some indication that the VI 40-s schedules delivered slightly more reinforcers when they operated




Fig. 11. Average reinforcement rate obtained across the four trial periods during Conditions 1 (dark bars) and 2 (light bars) of Experiment 2. The dotted lines correspond to the rates predicted by a VI 40-s schedule (1.5 reinforcers per minute) and a VI 120-s schedule (0.5 reinforcer per minute).

on the left key (cf. the heights of the longer bars during the first two and the last two 15-s periods). This difference was picked up by two repeated measures ANOVAs with condition and key as factors. The main effect of key was significant both for Group DIF,  $F(1, 8) = 10.74, p = .01$ , and for Group EQU,  $F(1, 4) = 15.9, p = .02$ ; the remaining effects were not significant. Although statistically significant, these differences were probably of no practical significance because they were small (0.12 reinforcers per minute for Group DIF and 0.13 for Group EQU), always favored the left key, and affected the two groups in the same way.

Figure 12 shows the individual psychometric functions for Group DIF. For all pigeons the function for Condition 2 was to the left of the function for Condition 1. The magnitude of the shifts ranged from 4 s to 11 s (mean = 7.2 s, median = 8 s). These values

were close to those observed in Group MID during Experiment 1 (range, 4 s to 10 s; mean = 8 s; median = 9 s). The 2 birds that came from Group EXT of Experiment 1 showed larger shifts during Experiment 2 (Bird 9854: 1 s in Experiment 1 vs. 7 s in Experiment 2; Bird 3860: 2 s in Experiment 1 vs. 6 s in Experiment 2). Of the 4 birds that came from Group MID of Experiment 1, 2 showed a smaller shift during Experiment 2 (Bird 4171: 10 s vs. 4 s; Bird 5259: 10 s vs. 8 s) and 2 showed a larger shift (Bird 4170: 4 s vs. 9 s; Bird 5320: 7 s vs. 8 s). Finally, the 3 birds that came from Group TOT of Experiment 1 showed a smaller shift during Experiment 2 (Bird 4138: 22 s vs. 11 s; Bird 1346: 9 s vs. 4 s; Bird 5841: 13 s vs. 8 s). No psychometric curve in Experiment 2 started above .05, but four ended below .95 (Bird 4171, Condition 2; Bird 4138, Condition 1; Bird 4170, Condition 2; and Bird 5259, Condition 2).

Figure 13 shows the results for Group EQU. For 3 birds the two psychometric functions overlapped considerably; for the other 2 birds there was a slight left shift of the function for Condition 2. The magnitude of the shifts ranged from 0 s to 3 s (mean = 1.2 s, median = 1 s). These values were also close to those of Group EXT in Experiment 1 (range, -2 s to 3 s; mean = 0.5 s; median = 1 s). The 4 birds that came from Group TOT or Group MID of Experiment 1 showed smaller shifts in Experiment 2 (Bird 4161: 6 s vs. 2 s; Bird 2816: 12 s vs. 0 s; Bird 3856: 8 s vs. 3 s; Bird 4913: 10 s vs. 1 s). Bird 9805, which in Group EXT of Experiment 1 had shown a shift of -2 s, showed no shift during Experiment 2. Most psychometric functions started close to 0 and ended close to 1. The exceptions were the functions for Condition 1 of Birds 3856 and 2816, which started above .05, and that of Bird 4161, which ended below .95.

In conclusion, the birds from Group DIF showed larger shifts of the psychometric function than did the birds from Group EQU. The difference between the two groups in the magnitude of the shifts was statistically significant,  $t(12) = 6.12, p < .001$ . These results are consistent with LeT but not with a threshold-based account of the shifts of the psychometric function.

Figure 14 shows the average of the nor-

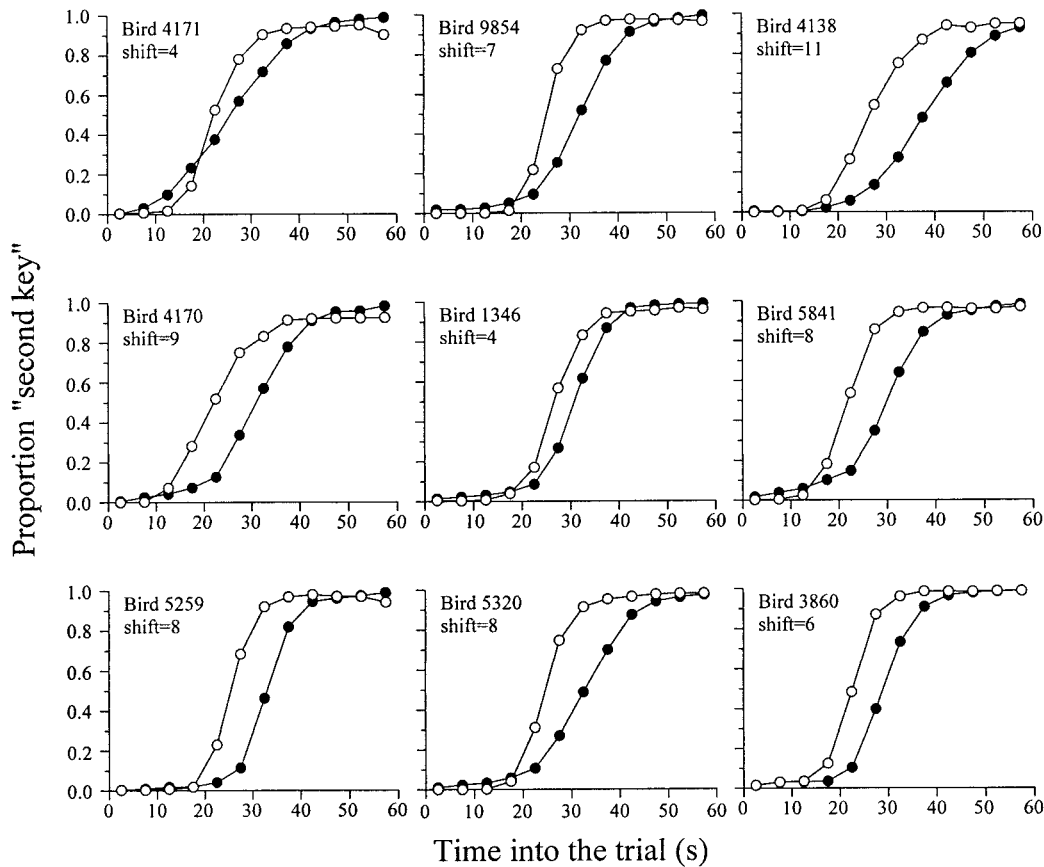


Fig. 12. Individual psychometric functions during Conditions 1 (filled circles) and 2 (open circles) from Group DIF of Experiment 2. The shift is measured from the bisection points of the curves.

malized response-rate curves. The curves for Group DIF intersected at approximately 32 s in Condition 1 and 24 s in Condition 2, yielding an 8-s average shift. For Group EQU, the two pairs of curves intersected at approximately 27 s, yielding a 0-s average shift. As in Experiment 1, the rate curves can shed light on the reasons for the different degrees of shift in the psychometric functions. Take the descending curve with filled circles (left key): The curve for Group DIF (top panel) descends more slowly than the curve for Group EQU (bottom panel). This result cannot be due to the reinforcement conditions operating on the left key because they were equal. Therefore, it must stem from the reinforcement conditions on the right key, in particular the conditions holding during the third 15-s trial period. For Group DIF the right key delivered 0.5 reinforcers per minute (VI 120 s), whereas for Group EQU it delivered 1.5

reinforcers per minute (VI 40 s). The higher reinforcement rate on the right key during the third trial period may therefore explain why the left curve for Group EQU decreased more rapidly than the curve for Group DIF. The same reason may also explain why the ascending curve with filled circles (right key) increased more rapidly in Group EQU (bottom) than in Group DIF (top). The remaining curves (open circles) support this interpretation.

The results from both experiments support the following generalization: When the reinforcement rate associated with one response varies across time, the strength of that response will also vary across time—the bird discriminates the different reinforcement rates during different periods. Consequently, the choice ratios expressed by the psychometric function show a large shift when the reinforcement differential holds around the

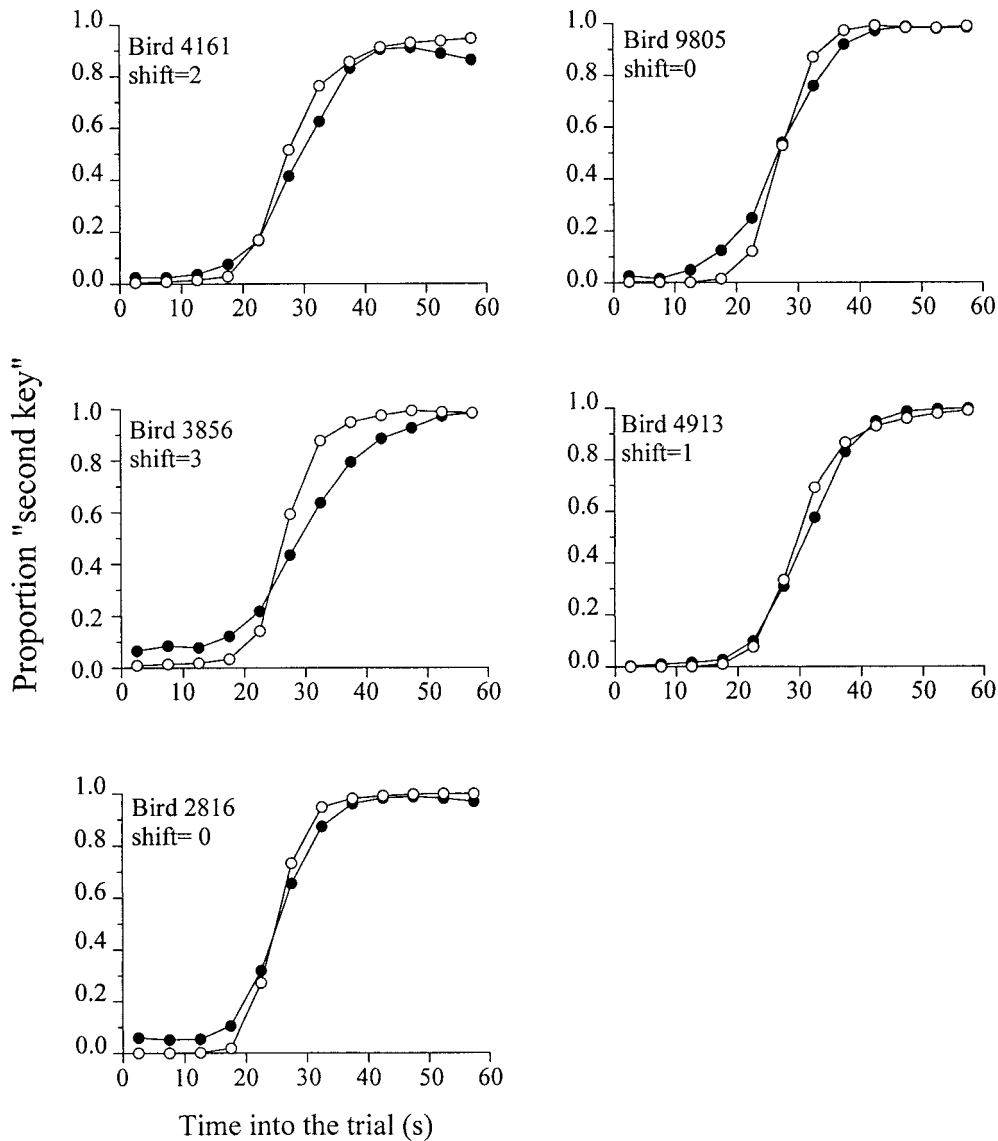


Fig. 13. Individual psychometric functions during Conditions 1 (filled circles) and 2 (open circles) from Group EQU of Experiment 2. The shift is measured from the bisection points of the curves.

middle of the interval but show a small or zero shift when it holds at other temporal locations.

#### GENERAL DISCUSSION

In Experiment 1, Group TOT experienced a threefold difference in reinforcement rate between the left and right keys, a difference that held both around the middle and at the extremes of the trial. Groups MID and EXT

were the critical groups because (a) the difference in reinforcement rate between the keys was similar in the two groups, and (b) Group MID experienced different reinforcement rates around the middle of the trial and similar rates at the extremes of the trial, whereas Group EXT experienced a difference at the extremes but not at the middle of the trial. According to the threshold-based account, all groups should show a shift; according to LeT, only Groups TOT and MID



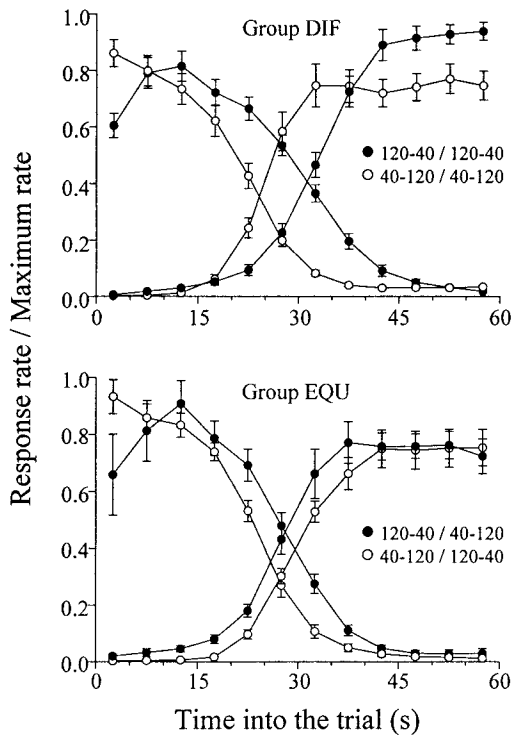


Fig. 14. Average normalized response rate across the trial for Conditions 1 (filled circles) and 2 (open circles) of Experiment 2. The descending curves are for the left key; the ascending curves are for the right key. The bars represent the standard error of the mean.

should show a shift. Moreover, the magnitude of the shifts for Groups TOT and MID should be roughly similar because for LeT the difference in reinforcement rates at the extremes of the trial contributes much less to the shift than the difference around the middle of the trial; the latter was the same for these two groups. The results were consistent with LeT, in that Groups TOT and MID showed a large shift and Group EXT did not. Furthermore, although the 11-s average shift for Group TOT was greater than the 8-s average shift for Group MID, this difference was not statistically significant.

In Experiment 2, the left and right keys always delivered the same reinforcement rate. However, for Group DIF the reinforcement rates around the middle of the trial differed, whereas for Group EQU they were equal. Without its putative cause for response threshold changes, the threshold-based account predicted no shift in either group, whereas LeT predicted a shift in Group DIF

but not in Group EQU. The results favored LeT because Group DIF showed a 7-s average shift, whereas Group EQU showed a 1-s average shift. Also in agreement with LeT, the magnitude of the shift for Group DIF was similar to that shown by Groups TOT and MID in Experiment 1 (in these three groups the difference in reinforcement rates during the middle of the trial was the same).

Although SET cannot account for the shifts of the psychometric function in terms of either memorial contents or threshold changes, a third possibility remains open: a change in the speed of the pacemaker. In fact some studies have found that shifts of the psychometric function similar to those displayed in Figure 2 may be obtained by using drugs (e.g., Meck, 1983) or by changing reinforcement rate (Bizo & White, 1994a, 1994b; Fetterman & Killeen, 1991; Killeen & Fetterman, 1988; Morgan, Killeen, & Fetterman, 1993), two manipulations that are presumed to change the speed of the pacemaker. Speeding the pacemaker causes a left shift in the psychometric function because the same interval seems longer than before, whereas slowing the pacemaker causes a right shift because the same interval seems shorter. Although changes in pacemaker speed have not figured prominently in SET, they are certainly not excluded by it. However, there are two problems with this account. First, to predict the direction of the shifts in Bizo and White's (1995a) experiment, in which the overall reinforcement rate remained constant across conditions, one would need to assume that the speed of the pacemaker varied directly with the reinforcement rate on the right key: Increasing that rate increased the speed of the pacemaker, which in turn caused the left shift of the psychometric function; decreasing it slowed the pacemaker and the psychometric function shifted to the right. But this account seems implausible because it implies that the reinforcement rate on the left key had no effect. The authors who first advanced this account were not convinced: "It does seem unlikely that the conditions in the left-key component would be totally ineffective in influencing timing" (Bizo & White, 1995a, p. 230). Second, even if the speed of the pacemaker changed, after sufficient training with the new schedules the animal should relearn the temporal discrimination and the shift

should disappear. In other words, because accurate timing may be obtained with different pacemaker speeds, any shift should be temporary. In Bizo and White's (1995a) experiment, the shifts lasted more than 20 sessions.

*Fitting LeT to the data.* Although LeT provides a good account of the main qualitative aspects of the data, the quantitative fit of its predictions to the data remains to be tested. The Appendix details how LeT's predictions for the FOPP are derived. We reproduce here the final equations. Each behavioral state is characterized by its degree of activation, a time-dependent variable represented by  $X_n(t)$ , where  $t$  is the time elapsed since the onset of the trial, and  $n \geq 0$  is the state number; the asymptotic couplings between state  $n$  and the left and right operant responses are represented by the variables  $WL_n$  and  $WR_n$ , respectively; the strengths of the left and right responses  $t$ s into the trial are represented by  $RL(t)$  and  $RR(t)$ , respectively; and the probability of choosing the right key at time  $t$  is represented by  $p(t)$ . Concerning the task variables, we let a generalized FOPP trial be divided into  $k$  periods, not necessarily equal,  $(T_0, T_1)$ ,  $(T_1, T_2)$ ,  $\dots$ ,  $(T_{k-1}, T_k)$ . In Experiments 1 and 2, for example,  $k = 4$ ,  $T_0 = 0$ ,  $T_1 = 15$  s,  $T_2 = 30$  s,  $T_3 = 45$  s, and  $T_4 = 60$  s. The reinforcement probabilities per second for the left and right responses during the  $i$ th time period ( $1 \leq i \leq k$ ) are represented by  $p_i$  and  $q_i$ , respectively; in the derivations we have approximated these probabilities by the reciprocal of the VI parameter.

The activation of the behavioral states is the familiar Poisson distribution, the cornerstone of BeT and LeT:

$$X_n(t) = \frac{\exp(-\lambda t)(\lambda t)^n}{n!}, \quad (1)$$

where  $\lambda$  is a pacemaker-like parameter that controls how fast the activation flows across the states. For the first state,  $n = 0$ , the activation follows the exponential function  $\exp(-\lambda t)$ . For the remaining states, the activation increases until  $t = n/\lambda$  and then decreases (see examples in the top right panel of Figure 4).

The equations for the asymptotic values of the couplings are

$$WL_n = \sum_{i=1}^k p_i \times \frac{f_n(T_i, T_k) - f_n(T_{i-1}, T_k)}{1 - f_n(0, T_k)} \quad (2)$$

and

$$WR_n = \sum_{i=1}^k q_i \times \frac{f_n(T_i, T_k) - f_n(T_{i-1}, T_k)}{1 - f_n(0, T_k)}, \quad (3)$$

where the function  $f_n$ , used to simplify the notation, is defined by  $f_n(t_i, t_j) = \exp[-\beta \int_{t_i}^{t_j} X_n(\tau) d\tau]$ , and  $\beta$  is a learning parameter that determines how much the couplings change with reinforcement and extinction (the Appendix provides various interpretations of these coupling equations). Response strength is obtained from the dot product of the vectors of state activation and the couplings

$$RL(t) = \sum_n X_n(t) WL_n(t) \quad (4)$$

and

$$RR(t) = \sum_n X_n(t) WR_n(t). \quad (5)$$

Finally, the probability of choosing the right key at time  $t$  is given by the relative response strength,

$$P(t) = \frac{RR(t)}{RR(t) + RL(t)}. \quad (6)$$

The full model has two free parameters,  $\lambda$ , which is included in  $X_n(t)$ , and  $\beta$ , which is included in  $WL_n$  and  $WR_n$ . However, in the Appendix we show that when  $\beta$  is small the coupling equations can be simplified and  $\beta$  can be eliminated. In this case LeT becomes a one-parameter model.

We fitted the model to the average results of five experiments, Experiments 1 and 2 of the present study, Bizo and White's (1995a) study already mentioned, and two experiments reported by Stubbs (1980) and described below. It would be misleading to use different sets of parameters to fit different curves from the same experiment, because our claim has been that the shifts of the psychometric function stem directly from the pattern of reinforcement contingencies within the trial. For this reason we forced the parameters to remain the same within each experiment.

Figure 15 shows the fits to the average data from Experiment 1. The theoretical curves fit the data well in the middle segments of the

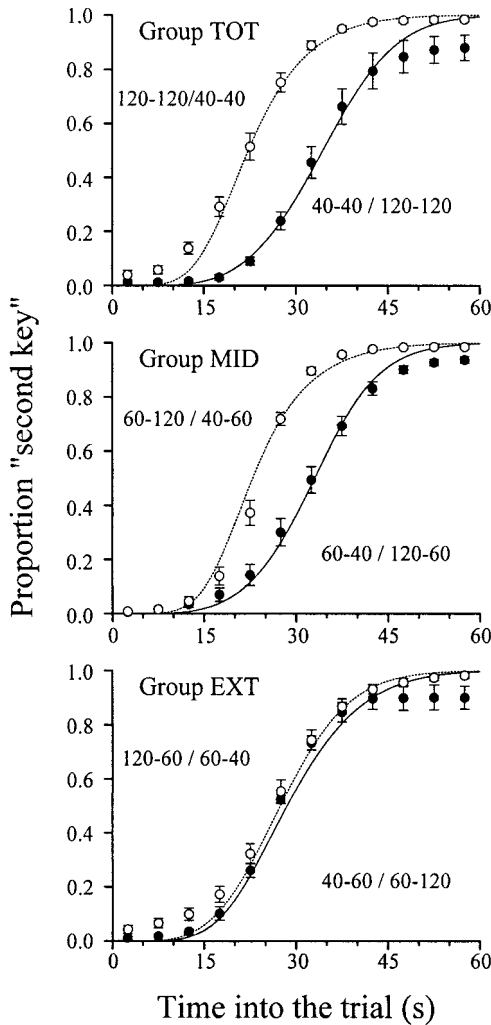


Fig. 15. Average psychometric functions obtained in Experiment 1 during Conditions 1 (filled circles) and 2 (open circles). The bars represent the standard error of the mean. The curves show LeT's predicted functions.

trial, but less well at the extremes. The parameters for the six curves were  $\lambda = 0.7$  and  $\beta = 0.5$ , and the variance accounted for ranged from .982 to .993 ( $\overline{\omega^2} = .988$ ). The approximation that uses only one parameter yields an average goodness of fit of .985 ( $\lambda = 0.6$ ). Furthermore, if instead of the scheduled reinforcement rates one uses the average of the obtained reinforcement rates, the fit remains equally good ( $\overline{\omega^2} = .985$ ). These findings confirm LeT's basic predictions.

Figure 16 shows the fits of the model to the data from Experiment 2. The parameters for the four curves were  $\lambda = \beta = 1$ , and the

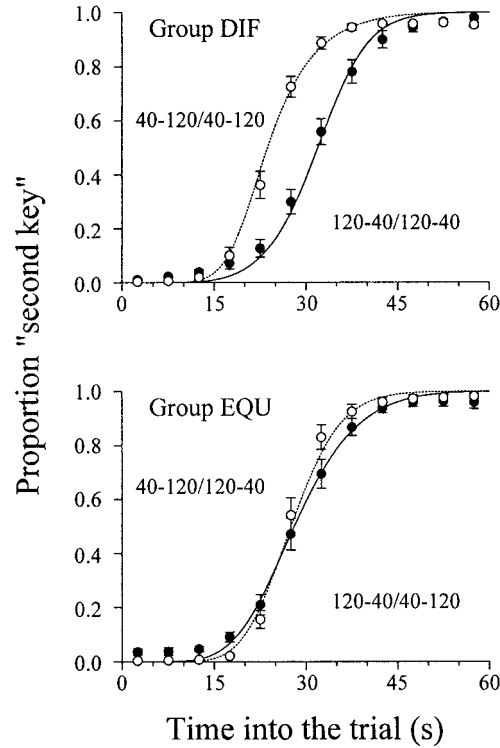


Fig. 16. Average psychometric functions obtained in Experiment 2 during Conditions 1 (filled circles) and 2 (open circles). The bars represent the standard error of the mean. The curves show LeT's predicted functions.

goodness of fit ranged from .992 to .996 ( $\overline{\omega^2} = .995$ ); with the parameter values used for Experiment 1 the average goodness of fit was .993. We see also that the fit is better at the middle than at the extremes of the trial.

Figure 17 shows the fits of the model to three data sets reported by other researchers. The top panel replots Bizo and White's (1995a) findings (already displayed in Figure 2). The parameters were  $\lambda = 1$  and  $\beta = 0.1$ , and goodness of fit ranged from .987 to .995 ( $\overline{\omega^2} = .993$ ). The single-parameter approximation with  $\lambda = 1$  yielded an equally good fit ( $\overline{\omega^2} = .994$ ). Again, the theory deviates from the data at the end of the trial, particularly for Conditions 40/120 (filled circles) and 45/90 (filled triangles).

The middle and bottom panels show the fits of the model to data from two experiments by Stubbs (1980). The first experiment used the same rationale as Bizo and White's (1995a) study, and the second used the same rationale as our Experiment 2. However, the

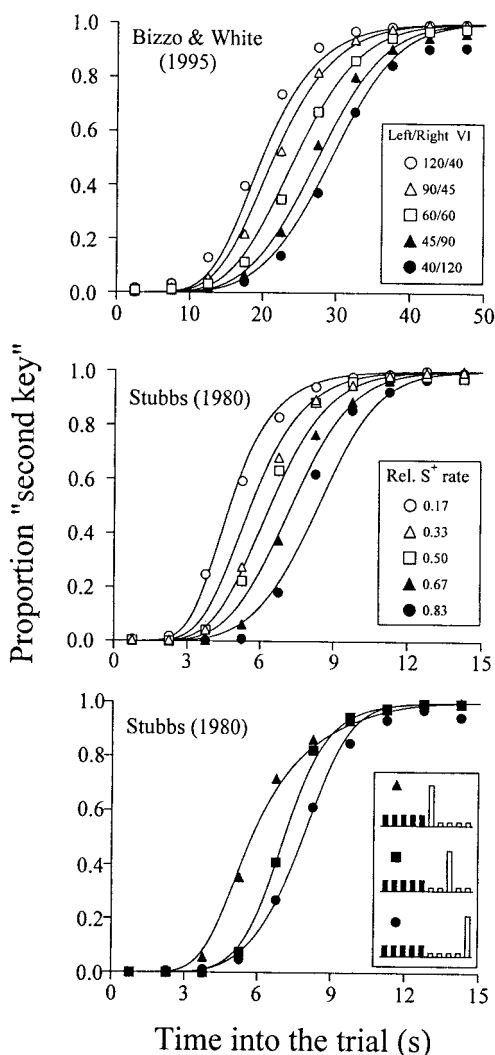


Fig. 17. Top: psychometric functions obtained by Bizo and White (1995a) when the reinforcement rate between the two keys varied across conditions. The first and second VI schedules were associated with the left and right keys, respectively (same data as Figure 2). Middle: psychometric functions obtained by Stubbs (1980, Experiment 2) when the relative reinforcement rate for the left key varied across conditions. Bottom: psychometric functions obtained by Stubbs (1980, Experiment 3) when the distribution of the reinforcers during the second half of the trial varied across conditions. The distributions during the first and second halves are shown by the black and light bars, respectively. The curves are LeT's predicted psychometric functions.

procedure was substantially different, because instead of the two-key left-right choice procedure, Stubbs used Findley's (1958) alternative with reinforcement key and change-

over key. The reinforcement key could be either red or green. Pecks at the changeover key changed the color of the reinforcement key and with it the operative reinforcement schedule. The trials were at most 15 s long and started with the red light on the reinforcement key. During the first half of the trial, only pecks in the presence of the red light were reinforceable; pecks in the presence of the green light were reinforceable only during the second half of the trial. One reinforcer was always scheduled for each trial. One additional feature of Stubbs' procedure complicates the analysis of his experiments, namely, ending a trial as soon as the reinforcer was delivered. This meant that well-trained birds had much more exposure to the red lights than the green lights, that the two choices differed in opportunities for extinction, that the obtained overall reinforcement rate was not constant across the experimental conditions (described below), and that scheduled and experienced relative reinforcement rates also differed. To fit LeT to Stubbs' data, we assumed that his procedure was functionally equivalent to the procedure used in the present study—the red and green lights were equivalent to the left and right choice keys; ending the trial when the reinforcer was delivered was inconsequential; and the scheduled and obtained reinforcement rates were equal.

In one experiment with 3 pigeons, Stubbs (1980) varied the relative reinforcement rate on the left key from .17 to .83 and obtained the results displayed in the middle panel of Figure 17. These results show that, as in Bizo and White's (1995a) study, the psychometric function shifts to the right when the relative reinforcement rate for the left key increases. The parameters used to fit LeT were  $\lambda = 3.3$  and  $\beta = 8$ , and goodness of fit ranged from .969 to .997 ( $\omega^2 = .986$ ). Note that the fit is not particularly good for the filled circles because the data approached a step function: In order to fit that function, LeT would require a higher value for  $\lambda$ , but then the magnitude of the shifts predicted for the remaining conditions would be smaller than the observed values.

In another experiment with 2 pigeons, Stubbs (1980) scheduled the same reinforcement rate on the two keys by having 50% of the reinforcers allocated to the first half and

50% to the second half of the trial. However, the temporal distribution of the reinforcers differed between the halves of the trial. A reinforcer set up for the first 7.5 s was delivered with equal probability in each of the five 1.5-s bins that comprised that part of the trial, but a reinforcer set up for the second 7.5 s was not delivered with equal probability in each of the five bins of that part of the trial. Instead, its temporal distribution varied across three conditions. The inset graphs in the bottom panel of Figure 17 illustrate the details. In one condition (triangles) most second-half reinforcers fell into the 7.5- to 9.0-s bin, the beginning of the second half of the trial; in another (squares) they fell into the 10.5- to 12.0-s bin, the middle of the second half; and in a third one (circles) they fell into the 13.5- to 15-s bin, the end of the second half. The results of this ingenious experiment showed that the psychometric function shifted increasingly more toward the left (i.e., to an earlier moment) as the majority of the reinforcers delivered during the second half of the trial approached the midpoint of the trial. Although the procedural details of this experiment differed markedly from our Experiment 2, its results support the same conclusion: A difference in relative reinforcement rates around the middle of the trial shifts the psychometric function by a greater amount than the same difference at other locations. Furthermore, because in Stubbs' experiment each trial ended when the reinforcer was delivered, the shifts of the psychometric function cannot be due to some momentary response-perseveration effect of the reinforcer. Finally, the theoretical curves shown in the bottom panel were obtained with  $\lambda = 3.9$  and  $\beta = 6$ , values roughly comparable to those used in the preceding fit. LeT's goodness of fit was .990 (circles), .993 (triangles), and .996 (squares).

In conclusion, LeT provided a reasonably good fit to the five data sets. The values of its two parameters were also reasonably close when the fitted data came from similar experiments; in some sets the model could fit the data well with one single parameter; and the fits of the model were robust in the sense that they were not strongly affected by parameter values. Perhaps the most surprising finding revealed by these analyses is that one set of parameters generated curves for different

experimental conditions that shifted in the same direction and by the same amount as the empirical curves. This finding is significant because it shows that the model is as sensitive as the birds are to the variations in the independent variable.

But LeT also showed weaknesses. These were of two types. First, its predictions for the initial and final values of the psychometric function were on occasion too extreme when compared with the data. Both group data (e.g., in Figure 15: top panel, empty circles at the beginning and filled circles at the end; bottom panel, filled circles at the end) and individual data (e.g., Figure 7) revealed that the pigeons' temporal discriminations at the extremes of the trial were sometimes worse than LeT predicts. Second, Stubbs' (1980) data in particular showed that a few psychometric functions are so steep that LeT cannot fit them properly. The reasons for these misfits are unclear. On the one hand they could stem from processes not included in the model. For example, several data sets show that the discrimination tends to be worse at the end than at the beginning of the trial, even when the reinforcement schedules are the same during these periods. Perhaps the event that signals the beginning of the trial—typically the illumination of the keys and the houselight—exerts some form of situational control (Staddon, 1974) that improves the discrimination. This event is absent from the end of the trial. On the other hand, the misfits could stem from one or more incorrect model assumptions, the Poisson distribution assumed to govern the dynamics of the behavioral states, the linear learning rules assumed to govern the dynamics of the couplings between the states and the operant responses, or the ratio rule assumed to map response strengths onto response probabilities.

The matter awaits further investigation, but the following results from our theoretical analyses may pave the way to a better model. It is well known that the Poisson distribution does not describe correctly the dynamics of the states without ancillary assumptions (e.g.,  $\lambda$  is proportional to overall reinforcement rate). However, when we replaced the Poisson by a normal distribution with mean and standard deviation proportional to the interval to be timed, that is, when we assumed the scalar



property, we found essentially the same goodness of fit—and the same problems. Concerning the learning rules, which incidentally were dictated more by parsimony and mathematical tractability than by any other reason, we have found that their linear form leads to coupling values for the early and late states roughly proportional to the reinforcement rates at the beginning and end of the trial, respectively. It follows that the predicted response strengths at the extremes of the trial tend to be also proportional to the reinforcement rates during those periods (see the Appendix). However the observed response-rate curves displayed in Figures 10 and 14 show that this is not the case. For example, response rates were not three times higher under the VI 40-s schedules than under the VI 120-s schedules. Some form of nonlinearity seems to be required.

*Broader implications for SET and LeT.* Much of contemporary research on timing has not been concerned with issues related to learning, but with issues related to the scalar property. Yet even to deal only with the scalar property, researchers must assume some form of learning process or, to put it differently, what they assume in terms of internal architectures, processing devices, rules, and the like, has implications for the learning process. In all likelihood experimental tests will show that these assumptions and implications are wrong—that all current theories of timing were born posthumously, one is tempted to say. However, one hopes that the tests will also identify the nature of the problems with the theories and thus contribute to a better understanding of the animal's learning processes. It is with these ideas in mind that we now discuss some of the shortcomings of SET and LeT (see also Kirkpatrick & Church, 1998; Staddon, 1999, and ensuing commentary; Zeiler, 1998).

Consider the memorial structure postulated by SET. In most if not all descriptions of the theory, the animal is assumed to form different memory stores, each containing the counts that were in the accumulator when a reinforcer was obtained. But how are these memories indexed? How does the animal know, as it were, in which memory store to save the current number in the accumulator? Although not always stated as explicitly as one would like, the reply is “If the reinforcer

came from the left key it is saved in one memory store; if it came from the right key it is saved in another.” The logic of the reply has some important and hitherto unforeseen consequences. To expose them, let us examine a simpler experimental task. A pigeon receives food for pecking a key after either 30 s or 240 s have elapsed since the onset of the trial. No cue signals whether the current trial will be short or long, and the two trial types are equally probable. The results of the experiment show that during the long trials average response rate increases from the beginning of the trial until approximately 30 s have elapsed, then it decreases, and then it increases again until the end of the trial (e.g., Catania & Reynolds, 1968; Leak & Gibbon, 1995). This performance is derived from SET by assuming that the animal stores the counts obtained at 30 s and 240 s into distinct memory stores. As Leak and Gibbon (1995, p. 6) put it, “In SET, there is assumed to be a single clock but an independent memory distribution for each criterion time interval.” Then at the beginning of the trial the bird samples a number from the short store, compares that number with the current number in the accumulator, pecks the key when the two numbers are sufficiently close, stops pecking when they become sufficiently different again, at which time it samples a number from the long store, and then executes the same routine. The account predicts the two peaks in the response-rate curve and the fact that their respective widths show the scalar property. Alas, the account also begs the question because that which was supposed to be explained was assumed in the explanation. (Logicians call this type of error *petitio principii*.) In fact, how does the animal know in which memory store to save a particular count? The reinforcers have the same source, and no distinct signal cues the two trials. To reply that when the count is small it is saved in the short store and when large it is saved in the long store is to explain nothing at all, because the meanings of *small* and *large* have yet to be defined in the theory and their discrimination accounted for. Once we reject this spurious explanation, we find ourselves without an answer to the question. (That Leak & Gibbon, 1995, did not raise the question of how the temporal discrimination might be learned according to SET; simply

presupposed some unstated answer to this question; and then moved on to examine the widths of the peaks of the response-rate function reveals how much the study of the scalar property has overshadowed the study of other timing issues.)

To be consistent and avoid the *petitio principii*, SET must assume that the animal's memories are indexed (formed, accessed, etc.) by structural features of the situation, by the signal being timed or by the source of the reinforcers, for example, and not by time itself. A coherent account would proceed by stating that when the reinforcers come from a single source and are not correlated with distinct exteroceptive stimuli, the counts in the accumulator are all lumped into one and the same memory store. Therefore, when the reinforcers are obtained at two distinct moments, as in the task analyzed above, the distribution of the counts in memory will be a mixture of two distributions, the one induced by the reinforcers delivered at short intervals (30 s) and the other by the reinforcers delivered at long intervals (240 s). The predicted pattern of behavior will also be a mixture of two patterns, the pattern predicted for a simple fixed-interval 30-s schedule and the pattern predicted for a fixed-interval 240-s schedule. This prediction is incorrect.

In addition to the problem of how the animal's memories are formed and accessed, SET must also reckon with some undesirable consequences of the one-dimensional contents it imputes to those memories. Recall that in SET each memory store contains only the distribution of reinforcement times associated with one or the other key. Therefore some other feature of the model must mediate the effects of reinforcement rate illustrated in Figure 2. We have seen that changes in the response thresholds or in the pacemaker speed cannot account for the entire pattern of shifts of the psychometric function observed in Experiments 1 and 2 (Figures 15 and 16), in Bizo and White's (1995a) study, and in Stubbs' (1980) earlier studies (Figure 17). Now we can state the general problem: Any evidence that the animal has learned that different reinforcement rates associated with one key hold during different periods within the trial cannot be explained by current versions of SET. For although the memory contents represent the reinforcement times and

the response thresholds or pacemaker speed may represent overall reinforcement rate, nothing in the model represents the variation of reinforcement rate with time or enables the derivation of that aspect from the model's architecture. To rephrase this problem in the language of information processing, one would say that the memorial representation must be two-dimensional because the FOPP experiments suggest (if our reading is correct) that the animal stores both the reinforcement time and the reinforcement frequency *at that time*.

In summary, we have identified two problems with SET that are related, not to the scalar property or to specific mathematical formulations, but to the model's overall conceptualization of the learning process, the memorial architecture that it postulates, and the rules for constituting and indexing those memories. At the root of these problems lies the assumption that animals lump into one memory store all the counts, and nothing but the counts, obtained when reinforcers are delivered from the same source. The challenge is clear: To solve the first problem SET must explain without assuming a temporal discrimination—for that is the very purpose of the theory—how distinct memories are formed; to solve the second problem SET must explain how the effects of reinforcement rate may be modulated by time into the trial. The first challenge is more logical than psychological and requires conceptual clarification rather than more experiments. The second is more psychological than logical and requires empirical tests such as those conducted in this study to determine whether the contents in memory are not at least two-dimensional. We anticipate one counterargument, namely, that SET is concerned exclusively with steady-state performance; that insofar as the issues we have raised address learning, they fall outside the boundaries of the model; that in their 1995 study Leak and Gibbon were not trying to explain how the pigeons learn the discrimination, but were examining only whether the birds' final performance in a simultaneous timing task reveals the scalar property. However, if this is the case then SET's status as a theory of timing becomes unclear, because the counterargument seems to reduce SET to a statement of a property of data, the scalar property. What then would

be the value of a complex architecture—switches, pacemaker, accumulator, memory stores, and response thresholds—that refuses to go beyond the scalar property, important and undisputed as it may be, and invokes for its refusal vague psychological distinctions between learning and performance or acquisition and terminal performance?

The problems with LeT seem to be related less to the model's overall architecture than to its specific mathematical formulations. In addition to the well-known problem of deriving the scalar property from the assumption that  $\lambda$  is strictly proportional to overall reinforcement rate—an assumption that is not supported by the data (e.g., Bizo & White, 1994a, 1994b; Fetterman & Killeen, 1991)—it seems that LeT's learning rule must also be revised to account for the nonlinear relation between reinforcement and response rates. (Additional quantitative difficulties with the model are discussed by Kirkpatrick & Church, 1998; Machado & Cevik, 1998; Rodríguez-Gironés & Kacelnik, 1999.) But one should remember that LeT's ability to predict the conditions in which the psychometric function should and should not shift depends more on the model's structural features than on the specific forms taken by their mathematical instantiation. These features include the serial process represented by the cascade of behavioral states and the reinforcement-based learning process represented by the couplings. In contrast with SET, LeT assumes that when reinforcers are received from one single key at distinct times, the couplings of the states are strengthened differentially and separately. Memory is two-dimensional in LeT because the couplings reflect both the distribution of reinforcement times associated with one key (via which couplings are strong) and the reinforcement rate associated with that key (via how strong the couplings are). In other words, to characterize each coupling two numbers are required, the number of its corresponding behavioral state and the coupling's strength. Therefore, the effects of reinforcement rate do not *have* to be mediated by other time-independent features of the model, although it seems likely that reinforcement rate also affects the speed of the pacemaker and response thresholds. These features enable LeT to capture something fundamental about temporal learning, some-

thing strongly evinced by the data from the FOPP, and something not captured by current versions of SET: the animal's sensitivity to differences in reinforcement rates at specific times during the trial.

Perhaps a hybrid between SET and LeT may solve some of the preceding problems and yet retain the strengths of each model (see Church, 1997, and Kirkpatrick & Church, 1998, for the benefits of hybridization in theoretical research). In the spirit of SET, the dynamics of the states would follow the scalar property (e.g., by having  $\lambda$  vary across trials; see Gibbon, 1992), and in the spirit of BeT and LeT, these states would be coupled with the operant response according to some nonlinear learning rule. The hybrid model would account for the findings described in this study as well as the findings reported by Bizo and White (1995b) and Stubbs (1980, Experiment 1) concerning the scalar property in the FOPP. It remains to be seen whether the new model will confirm the well-known vigor of the intraspecies hybrid or the equally well-known fact that most interspecies hybrids are sterile.

## REFERENCES

- Bizo, L. A., & White, K. G. (1994a). The behavioral theory of timing: Reinforcer rate determines pacemaker rate. *Journal of the Experimental Analysis of Behavior*, *60*, 19–33.
- Bizo, L. A., & White, K. G. (1994b). Pacemaker rate in the behavioral theory of timing. *Journal of Experimental Psychology: Animal Behavior Processes*, *20*, 308–321.
- Bizo, L. A., & White, K. G. (1995a). Biasing the pacemaker in the behavioral theory of timing. *Journal of the Experimental Analysis of Behavior*, *64*, 225–235.
- Bizo, L. A., & White, K. G. (1995b). Reinforcement context and pacemaker rate in the behavioral theory of timing. *Animal Learning & Behavior*, *23*, 376–382.
- Catania, A. C. (1970). Reinforcement schedules and the psychophysical judgments: A study of some temporal properties of behavior. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 1–42). New York: Appleton-Century-Crofts.
- Catania, A. C., & Reynolds, G. S. (1968). A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, *11*, 327–383.
- Church, R. (1997). Quantitative models of animal learning and cognition. *Journal of Experimental Psychology: Animal Behavior Processes*, *23*, 379–389.
- Church, R. M., & Deluty, M. Z. (1977). Bisection of temporal intervals. *Journal of Experimental Psychology: Animal Behavior Processes*, *3*, 216–228.
- Fetterman, J. G., & Killeen, P. R. (1991). Adjusting the pacemaker. *Learning and Motivation*, *22*, 226–252.

- Findley, J. D. (1958). Preference and switching under concurrent scheduling. *Journal of the Experimental Analysis of Behavior*, *1*, 123–144.
- Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, *84*, 279–325.
- Gibbon, J. (1981). On the form and location of the psychometric bisection function for time. *Journal of Mathematical Psychology*, *24*, 58–87.
- Gibbon, J. (1991). Origins of scalar timing theory. *Learning and Motivation*, *22*, 3–38.
- Gibbon, J. (1992). Ubiquity of scalar timing with a Poisson clock. *Journal of Mathematical Psychology*, *36*, 283–293.
- Killeen, P. R. (1991). Behavior's time. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 295–334). New York: Academic Press.
- Killeen, P., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review*, *95*, 274–285.
- Kirkpatrick, K., & Church, R. (1998). Are separate theories of conditioning and timing necessary? *Behavioral Processes*, *44*, 163–182.
- Leak, T. M., & Gibbon, J. (1995). Simultaneous timing of multiple intervals: Implications for the scalar property. *Journal of Experimental Psychology: Animal Behavior Processes*, *21*, 3–19.
- Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review*, *104*, 241–265.
- Machado, A., & Cevik, M. (1998). Acquisition and extinction under periodic reinforcement. *Behavioral Processes*, *44*, 237–262.
- Machado, A., & Keen, R. (1999). Learning to time (LET) or scalar expectancy theory (SET)? A critical test of two models of timing. *Psychological Science*, *10*, 285–290.
- Meck, W. (1983). Selective adjustment of the speed of internal clock and memory processes. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*, 171–201.
- Morgan, L., Killeen, P. R., & Fetterman, J. G. (1993). Changing rates of reinforcement perturbs the flow of time. *Behavioral Processes*, *30*, 259–272.
- Rodríguez-Gironés, M. A., & Kacelnik, A. (1999). Behavioral adjustment to modifications in the temporal parameters of the environment. *Behavioral Processes*, *44*, 173–191.
- Staddon, J. E. R. (1974). Temporal control, attention and memory. *Psychological Review*, *81*, 375–391.
- Staddon, J. E. R. (1999). Time and memory: Towards a pacemaker-free theory of interval timing. *Journal of the Experimental Analysis of Behavior*, *71*, 215–251.
- Stubbs, D. A. (1968). The discrimination of stimulus duration by pigeons. *Journal of the Experimental Analysis of Behavior*, *11*, 223–238.
- Stubbs, D. A. (1980). Temporal discrimination and a free-operant psychophysical procedure. *Journal of the Experimental Analysis of Behavior*, *33*, 167–185.
- Zeiler, M. D. (1998). On sundials, springs, and atoms. *Behavioral Processes*, *44*, 89–99.

Received September 13, 1999  
Final acceptance April 25, 2000

## APPENDIX

LeT has three components: the behavioral states, their couplings with the operant responses, and the strength of the operant responses (see Figure 1). To obtain LeT's predictions one needs the expression for  $X_n(t)$ , the activation of state  $n$  at  $t$  s into the trial, and the expressions for the asymptotic values of  $WL_n$  and  $WR_n$ , the couplings between state  $n$  and the left and right responses, respectively. From these expressions one gets the strengths of the left and right responses at time  $t$ ,  $RL(t) = \sum_n X_n(t) WL_n(t)$  and  $RR(t) = \sum_n X_n(t) WR_n(t)$ , respectively, and the probability of choosing the right key at time  $t$ ,  $p(t) = RR(t) / [RR(t) + RL(t)]$ .

In LeT, as in BeT,  $X_n(t)$  follows the Poisson distribution  $X_n(t) = \exp(-\lambda t) (\lambda t)^n / n!$ , where  $\lambda > 0$  is a parameter that controls how fast the activation spreads across the states. The derivation of this equation is given in Machado (1997).

The asymptotic values of  $WL_n$  and  $WR_n$  are obtained in similar ways. Hence to lighten the notation, we let  $W_n$  refer generically to

the coupling between state  $n$  and one operant response. With the help of Figure 18 we proceed in three steps. First, we derive an equation that describes how  $W_n$  changes during a short interval. Next we derive an equation relating the value of  $W_n$  at the beginning of one trial to the value of  $W_n$  at the beginning of the next trial. And finally we derive an expression for  $W_n$  after a large number of trials.

*Short interval.* Consider a short interval of, say, 1 s, from  $t$  to  $t + 1$  (see top of Figure 18). The values of  $W_n$  at the extremes of this interval are  $W_n(t)$  and  $W_n(t + 1)$ . We seek an equation relating these two values. We assume that if a reinforcer occurs during the interval, then  $W_n(t)$  changes by the amount  $\Delta W_n(t) = \beta \times X_n(t) \times [1 - W_n(t)]$ ; if a reinforcer does not occur during the interval, then  $W_n(t)$  changes by the negative amount  $\Delta W_n(t) = -\alpha \times X_n(t) \times W_n(t)$  (Machado, 1977). Therefore the expected value of  $W_n(t + 1)$  conditional on  $W_n(t)$  equals

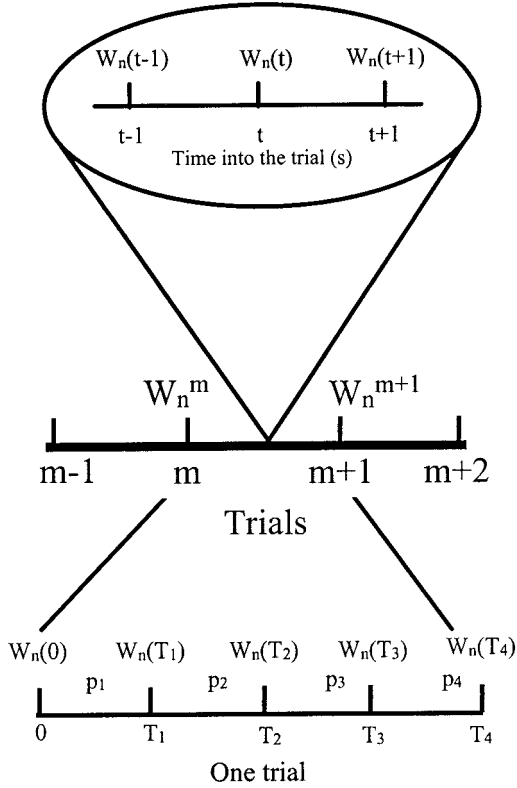


Fig. 18. Middle line: The FOPP session is divided into trials indexed by  $m$ . The value of coupling  $W_n$  at the end of each trial is represented by  $W_n^m$ . Bottom: Each trial is subdivided into  $k$  periods indexed by  $T_i$ ,  $1 \leq i \leq k$ . Top: Each period within the trial is subdivided into intervals of 1 s, indexed by  $t$ .

$$\begin{aligned}
 E[W_n(t+1) | W_n(t)] &= W_n(t) + \overbrace{\beta X_n(t) [1 - W_n(t)] \times I}^{\text{reinforcement}} \\
 &\quad - \underbrace{\alpha X_n(t) W_n(t) \times (1 - I)}_{\text{extinction}},
 \end{aligned}$$

where  $I = 1$  if a reinforcer occurred during the 1-s interval under consideration, and  $I = 0$  otherwise. To simplify the preceding equation and retain only two parameters in the model, we let  $\alpha = \beta$ . Then, after some simplification we get

$$\begin{aligned}
 E[W_n(t+1) | W_n(t)] &= W_n(t) + X_n(t) \times \beta \\
 &\quad \times [I - W_n(t)].
 \end{aligned}$$

The preceding expression is a conditional expectation—conditional on the value of  $W_n$  at

the beginning of the time period. To obtain the unconditional expectation of  $W_n$  at time  $t+1$ , one must take expectations again. This yields, after rearranging,

$$\begin{aligned}
 E[W_n(t+1)] - E[W_n(t)] &= X_n(t) \times \beta \times \{p - E[W_n(t)]\},
 \end{aligned}$$

where  $p$  is the probability of reinforcement during the 1-s interval. The solution of this difference equation may be approximated by the solution of the corresponding differential equation

$$\frac{dW_n(t)}{dt} = X_n(t) \times \beta \times [p - W_n(t)], \quad (\text{A1})$$

where we have dropped the expectation symbol to simplify the notation. The solution of Equation A1 is

$$\begin{aligned}
 W_n(t) &= p - [p - W_n(0)] \\
 &\quad \times \exp\left(-\beta \int_0^t X_n(\tau) d\tau\right),
 \end{aligned}$$

where  $W_n(0)$  is the value of  $W_n$  at the beginning of the integration period.

To ease the notation we introduce the following definition:

$$f_n(t_i, t_j) = \exp\left(-\beta \int_{t_i}^{t_j} X_n(\tau) d\tau\right)$$

and note three properties of this function, the last one of which is valid for small  $\beta$ :

$$f_n(t, t) = 1$$

$$f_n(t_i, t_j) \times f_n(t_j, t_k) = f_n(t_i, t_k)$$

$$f_n(t_i, t_j) \approx 1 - \beta \int_{t_i}^{t_j} X_n(\tau) d\tau.$$

Using the definition of function  $f_n$ , the expression for  $W_n(t)$  can be written more compactly,

$$\begin{aligned}
 W_n(t) &= p[f_n(t, t) - f_n(0, t)] \\
 &\quad + f_n(0, t) W_n(0). \quad (\text{A2})
 \end{aligned}$$

*One trial.* We divide an FOPP trial into  $k$  periods (see bottom part of Figure 18). The operant response is reinforced with probability  $p_1$  per second during the period  $(0, T_1)$ , with probability  $p_2$  during the period  $(T_1,$



$T_2$ ), and so on; during the last period,  $(T_{k-1}, T_k)$ , the probability of reinforcement is  $p_k$ . To relate the values of  $W_n$  at the beginning and at the end of the trial, we use Equation A2 for each period with the appropriate value of  $p$  and the appropriate integration limits. Thus at time  $T_1$ ,  $W_n$  will equal

$$W_n(T_1) = p_1[f_n(T_1, T_1) - f_n(0, T_1)] + f_n(0, T_1)W_n(0).$$

At time  $T_2$ ,  $W_n$  will equal

$$W_n(T_2) = p_2[f_n(T_2, T_2) - f_n(T_1, T_2)] + f_n(T_1, T_2)W_n(T_1).$$

Replacing  $W_n(T_1)$  by its expression yields

$$W_n(T_2) = p_2[f_n(T_2, T_2) - f_n(T_1, T_2)] + p_1[f_n(T_1, T_2) - f_n(0, T_2)] + f_n(0, T_2)W_n(0).$$

The general pattern is

$$W_n(T_k) = \sum_{i=1}^k p_i[f_n(T_i, T_k) - f_n(T_{i-1}, T_k)] + f_n(0, T_k)W_n(0). \quad (\text{A3})$$

*Steady state.* To obtain the expression for  $W_n$  after a large number of trials, we note that the value of  $W_n$  at the end of one trial equals the value of  $W_n$  at the beginning of the next trial. In other words, if superscript  $m$  denotes the trial number, then

$$W_n^m(T_k) = W_n^{m+1}(0).$$

The left side of the preceding equality is the strength of the coupling between state  $n$  and the operant response at the *end* of trial  $m$  (i.e., at time  $t = T_k$ ); the right side is the strength of that coupling at the *beginning* of trial  $m + 1$  (i.e., at time  $t = T_0$ ). From Equation A3 we get

$$W_n^{m+1}(0) = \sum_{i=1}^k p_i[f_n(T_i, T_k) - f_n(T_{i-1}, T_k)] + f_n(0, T_k)W_n^m(0).$$

The solution of this difference equation yields the long-term value of  $W_n(0)$ ,

$$W_n^\infty(0) = \sum_{i=1}^k p_i \frac{f_n(T_i, T_k) - f_n(T_{i-1}, T_k)}{1 - f_n(0, T_k)}. \quad (\text{A4})$$

This is equivalent to Equations 2 and 3 in the text. If the learning parameter  $\beta$  is small, we may approximate the steady-state value of  $W_n$  by using the third property of function  $f_n$  defined above. The result is

$$W_n^\infty(0) \approx \sum_{i=1}^k p_i \frac{\int_{T_{i-1}}^{T_i} X_n(\tau) d\tau}{\int_0^{T_k} X_n(\tau) d\tau}. \quad (\text{A5})$$

Equation A5 shows two things. First, when  $\beta$  is small, the steady-state distribution of  $W_n$  depends only on  $\lambda$  (included in  $X_n$ ). Second, the steady-state value of  $W_n$  is approximately a weighted sum of the reinforcement probabilities. Each weight is a ratio, with the denominator equal to the total activation of state  $n$  at the end of one trial and the numerator equal to the part of the total activation gained during the corresponding reinforcement period.

To understand intuitively the form of the steady-state distribution of  $W_n$ , consider Equation A1 again:

$$\frac{dW_n(t)}{dt} = X_n(t) \times \beta \times [p - W_n(t)].$$

At time  $t$ ,  $W_n(t)$  is attracted to the current reinforcement probability,  $p$ : When  $W_n(t) > p$ , the time derivative of  $W_n(t)$  is negative and therefore  $W_n(t)$  decreases towards  $p$ ; when  $W_n(t) < p$ , the time derivative is positive and therefore  $W_n(t)$  increases towards  $p$ . Moreover, the rate at which  $W_n(t)$  approaches  $p$  depends on the activation of state  $n$ ,  $X_n(t)$ . But  $p$  and  $X_n(t)$  change across the trial, which means that  $W_n(t)$  is attracted to different values of  $p$  and at different rates. Hence its final value is a weighted sum of the various  $p$  values with the normalized integrated rates of approach as the weights.

Another way to understand the steady-state distribution of  $W_n$  is to revert to the stochastic interpretation favored by BeT. Each ratio in Equation A5 is the probability that the animal will have reached state  $n$  while the reinforcement probability equals  $p_i$  (numerator), provided that the animal has reached state  $n$  before the trial ended (denominator). Because  $W_n$  approaches  $p_i$  only when the animal is in state  $n$  and the reinforcement probability

equals  $p_i$ , the final value of  $W_n$  is a weighted average of the  $p_i$ s.

*A specific example.* To illustrate how the model is applied to a specific case, consider one of the conditions of Bizo and White's (1995a) study. During the initial 25 s of the 50-s trial, left pecks were reinforced according to a VI 40-s schedule and during the last 25 s of the trial right pecks were reinforced according to a VI 120-s schedule. In this case there were  $k = 2$  periods, the first from  $T_0 = 0$  s to  $T_1 = 25$  s and the second from  $T_1 = 25$  s to  $T_2 = 50$  s. The reinforcement probabilities for left pecks were  $p_1 = 1/40$  during  $(0, T_1)$  and  $p_2 = 0$  during  $(0, T_2)$ ; for right pecks they were  $q_1 = 0$  and  $q_2 = 1/120$  during the same periods. Recall that  $WL_n$  and  $WR_n$  stand for the asymptotic couplings of state  $n$  with the left and right responses, respectively. If we use the exact Equation A4 we get

$$WL_n = \frac{1}{40} \frac{f_n(25, 50) - f_n(0, 50)}{1 - f_n(0, 50)},$$

and

$$WR_n = \frac{1}{120} \frac{1 - f_n(25, 50)}{1 - f_n(0, 50)}.$$

If we use the approximation, Equation A5, we get

$$WL_n = \frac{1}{40} \frac{\int_0^{25} X_n(\tau) d\tau}{\int_{25}^{50} X_n(\tau) d\tau},$$

and

$$WR_n = \frac{1}{120} \frac{\int_{25}^{50} X_n(\tau) d\tau}{\int_{25}^{50} X_n(\tau) d\tau}.$$

Replacing either set of values and the expression for  $X_n(t)$  in Equations 2, 3, and 6

gives the predicted response probability at time  $t$ .

*Shifts of the psychometric function.* The bisection point of the psychometric function may be approximated by determining first the value of  $n$  for which  $WL_n$  is approximately equal to  $WR_n$ . We call this value  $n^*$ . When  $n^*$  is the most active state, responding is affected mainly by its couplings with the two operant responses; if these couplings are equal, then choice probability is close to .5. In the foregoing example,  $n^*$  is the solution to the following equation:

$$3 \int_0^{25} X_{n^*}(\tau) d\tau = \int_{25}^{50} X_{n^*}(\tau) d\tau.$$

This equation shows that  $n^*$  is that state whose total activation during the first 25 s of the trial is three times smaller than the total activation during the last 25 s. Or in BeT's stochastic interpretation, state  $n^*$  is the state for which the following statement is true: The probability that state  $n^*$  has been reached by 25 s into the trial is three times less than the probability that state  $n^*$  is reached only during the last 25 s of the trial.

Consider now what happens when the two VIs are reversed, that is, when left and right keypecks are reinforced according to VI 120-s and VI 40-s schedules, respectively. In this case indifference occurs approximately at the time when the most active state is  $n^{**}$ , the solution to the equation

$$\int_0^{25} X_{n^{**}}(\tau) d\tau = 3 \int_{25}^{50} X_{n^{**}}(\tau) d\tau.$$

That is,  $n^{**}$  is the state whose cumulative activation during the first 25 s is three times larger than its cumulative activation during the last 25 s. Obviously,  $n^{**}$  must be less than  $n^*$ , which means that reversing the VIs brings about an earlier switch from the left to the right key. The foregoing analysis also shows that the difference between the bisection points when the VIs are reversed increases with the ratio of the two VI schedules.