Available online at www.sciencedirect.com

SCIENCE DIRECT®

Behavioural Processes xxx (2004) xxx–xxx

BEHAVIOURAL PROCESSES

www.elsevier.com/locate/behavproc

Reinforcer concentration effects on a fixed-interval schedule☆

Frances J. Blomeley[a], C.F. Lowe[a], J.H. Wearden[b,*]

[a] *University of Wales, Bangor, UK*

[b] *Department of Psychology, Manchester University, Manchester M13 9PL, UK*

Received 30 January 2004; received in revised form 10 February 2004; accepted 11 February 2004

Abstract

Four rats received training on a mixed FI 30-s FI 150-s schedule, where the different FI values were associated with different levers. During baseline, the reinforcer was a 30% concentration of condensed milk. During subsequent testing sessions, the reinforcer concentration was varied within sessions over values of 10, 30, 50, and 70%. Measures of behaviour were taken from the FI 30-s lever during trials where the reinforcer was delivered for responses on the other lever. Increasing the reinforcer concentration which began the interval (a) increased the time to start responding in the interval, and (b) increased the location of the response peak on the FI 30-s lever (often to values well above 30 s). Response rate at the peak, and spread of the response rate versus time function, changed much less with reinforcer concentration. The data are discussed relative to predictions derived from Scalar Expectancy Theory, the Behavioural Theory of Timing, and the Tuned-trace model.

Keywords: Fixed-interval (FI) schedule; Mixed FI; Reinforcer concentration; Scalar expectancy theory; Behavioural theory of timing; Tuned-trace theory

1. Introduction

Few behavioural processes in operant conditioning can be more fundamental than learning that a reinforcer is available for responses at one particular time, but not others. The fixed-interval (FI) schedule, introduced in its classic form by Skinner (1938), arranges a reinforcer for the first response occurring at some time, *t*, after a time marker, most often the previous reinforcer delivery. As is well known, the behaviour of many animal species adapts well to the FI contingency, with averaged response rate changing as a function of elapsed time since reinforcer delivery, from near zero responding just after the reinforcer to a maximum response rate at or near the time when the next reinforcer is available (Lejeune and Wearden, 1991).

Behaviour occurring on FI schedules, and variants like the peak procedure (Church et al., 1984; Roberts, 1981) has received extensive theoretical analysis. We might group models which treat FI performance into three categories.

One is a clock-based approach, typified most clearly by scalar expectancy theory (SET: Gibbon et al., 1984). An internal clock provides the raw material for judgements of time that are expressed in behaviour, and the elapsed time in the interval is assumed to be represented internally by an increasing accumulation of pulses from a pacemaker. During learning, some

☆ The experiment reported here was conducted in partial fulfilment of the requirements for the degree of Ph.D. awarded to the first author by the then-titled University College of North Wales in 1984.

* Corresponding author. Tel.: +44-1612752592; fax: +44-1612742588.

E-mail address: wearden@psy.man.ac.uk (J.H. Wearden).

0376-6357/$ – see front matter © 2004 Published by Elsevier B.V.

doi:10.1016/j.beproc.2004.02.005

critical number of pulses has been stored as representing the time of reinforcement, *t*, and as time elapses in the interval the organism is assumed to be comparing the number of pulses currently accumulated with the number representing *t*, and a response is generated when the two are sufficiently close. The time marker (usually the previous reinforcer) is supposed to reset the timing process back to zero so that the current interval can be timed.

Another approach is behaviourally based, and two closely related accounts of this sort are the Behavioural Theory of Timing (BeT: Killeen and Fetterman, 1988) and the Learning to Time model (LeT: Machado, 1997). Both assume that the timing of behaviour on FI is accomplished not by reading on some internal clock, but by behavioural sequences. These "adjunctive" sequences, usually unmeasured behaviours (but see Lejeune et al., 1998), are initiated by the time marker which starts the interval, then adjunctive behaviours are emitted in succession until one of these behaviours cues the measured response. In BeT, the rate of transition between adjunctive sequences is controlled by the "arousal" level of the organism, which is related to its expectancy of reinforcement.

Finally, there is the recently developed "Tuned-trace" theory of Staddon et al. (2002). The essential feature of this complex model is that timing depends, not on a clock process or a behavioural sequence, but on memory of the "trace" of the time marker which started the interval. So, for example, the trace of the reinforcer delivery which starts the interval decays as the interval elapses, and the animal learns to respond when some particular trace value is reached.

The clock and trace models are similar in that responding is cued by the value of some internal variable (pulse accumulation or trace value), whereas the behavioural theories appear to differ from them in that the cue for the measured response is another behavioural response, rather than an internal trace. However, since the adjunctive behaviours are usually unspecified by BeT, and seem reduced to the status of unmeasured "states" by LeT (e.g. Machado, 1997), it may be that the two types of theories are more similar than they at first appear. Hopson (2003) discusses some of the theoretical relations between BeT and LeT, and in particular the question of whether the "states" of LeT are behavioural in the same sense as those of BeT.

A problem in distinguishing the theories comes from the fact that they all make similar predictions in many cases. All, for example, predict that responding will on average increase with elapsed time in the FI. The clock and trace theories predict this because as time elapses the clock-reading or trace value that is currently present becomes more and more like that obtaining at the end of the interval (i.e. more and more like the value associated with reinforcement), so some threshold for responding is crossed with growing probability as time in the interval passes. The behavioural theories assume that the behavioural sequences become organised by reinforcement experiences so that behaviours which trigger the operant response become more probable towards the end of the interval on FI.

Although all current theories predict the basic properties of responding on simple FI schedules, more complicated experimental situations may distinguish them. These include experiments where animals can receive reinforcers at two potential times after the start of the interval (Leak and Gibbon, 1995; Whitaker et al., 2003), or where rate or probability of reinforcement are manipulated (Fetterman and Killeen, 1991, 1995; Morgan et al., 1993). The present article reports data from a situation where the concentration of a constant quantity of condensed milk reinforcer was varied within experimental sessions, after training with a constant concentration in baseline. As will be seen later, this manipulation produced some very clear behavioural effects, and current theories appear to differ in their predictions about what should occur in this situation. Variation in the concentration of sweet liquid reinforcers, when rats serve as subjects, is regarded as a manipulation of reinforcer magnitude, with higher concentrations representing greater magnitudes (e.g. Heyman and Monaghan, 1994). However, unlike increases in reinforcer magnitude occasioned by increasing the number of pellets delivered, or duration of access to grain, the time to consume the liquid reinforcers which differ in concentration but not volume presumably remains roughly constant.

Like those in another recent article (Whitaker et al., 2003), the data reported here have the unusual characteristic that they were collected 20 or more years ago (and reported in Blomeley, 1984, but otherwise unpublished). This means that the data predate most of the theories to which they now appear relevant. As in

the Whitaker et al.'s study, the age of the data mean that some behavioural measures which would be standard in modern work were not collected, and cannot now be derived from the data set which exists but, as will be seen later, the data that exist provide a very clear picture of the effects of reinforcer concentration on responding on an FI-like schedule.

2. Experiment

The basic experimental arrangement used in our study was a mixed fixed-interval (mixed FI) schedule, with two equally likely components, FI 30-s and FI 150-s. On each interval, reinforcer delivery was equally likely for the first response occurring either 30 s after the start of the interval, or 150 s, and nothing signalled to the animal which reinforcement time was in force. However, our experiment used a variant of mixed FI where the two different schedules (30 and 150 s) were consistently associated with different levers, so if the FI 30-s schedule was in force, the reinforcer was delivered for responses on one lever, and if the FI 150-s schedule was in force the reinforcer was delivered for responses on the other lever. Almost all the results that we have come from responses on the FI 30-s lever but, as will be seen later, these data show consistent effects of reinforcer concentration.

3. Method

3.1. Subjects and apparatus

Four male hooded rats, approximately 14 months old at the start of the experiment, served. All had previous experience of two-lever mixed-FI schedules like that used in the present study, except that food pellet reinforcers were employed. The operant chamber used had a floor area of 20 cm × 24 cm, and was 18.5 cm high. Three walls were made of bakelite, lined with sanded aluminium. The fourth wall was a metal panel on which lever levers were mounted either side of a 4 cm × 5 cm recess. The reinforcer (0.05 ml Nestlé condensed milk, diluted with water) was delivered up to the floor of the recess by a motor-operated dipper mechanism. The dipper was normally in the "up" position: at reinforcer delivery it was depressed into a container placed below, then returned, full, to the resting position. Six containers holding the condensed milk were mounted on a circular aluminium plate, the centre of which was attached to the spindle of a 24 V ac motor/gearbox. The motor was operated by external programming equipment, and rotation of the aluminium plate could thus change the reinforcer used from one interval to another. The operant chamber was housed in a sound-attenuating cubicle, with an exhaust fan, producing an ambient noise level of 50 ± 2 dB. The experiment was controlled by a Data General NOVA 1200 computer.

3.2. Procedure

As animals had received previous training with the mixed-FI schedule, no shaping was needed. In both baseline and test sessions, the tray holding the containers was rotated following reinforcer delivery across a pre-programmed number of components. To do this, a number from 1 to 6 was selected at random, and the tray rotated over this number of containers after reinforcer delivery. During baseline sessions the milk concentration was 30%, although the tray was rotated during baseline sessions to accustom the animals to the noise of its rotation. During the testing phase, compartments 1, 3 and 5 contained the 30% milk concentration used in the baseline sessions, compartment 2 was 10%, compartment 4 was 50% and compartment 6 was 70%. The FI associated with the short lever was 30 s and that associated with the long lever 150 s. Rats received between 17 and 19 baseline sessions, and which levers were associated with the short and long FI values was counterbalanced. The test phase lasted from six to eight sessions. Animal G3 died unexpectedly after the sixth test session, but provided sufficient data to be included.

4. Results

As mentioned above, almost all the data available come from responding on the short (FI 30-s) lever, but were taken from intervals in which the reinforcer was actually delivered for responding on the other (FI 150-s) lever.

The upper panel of Fig. 1 shows mean times to the first response on the FI 30-s lever (the post-

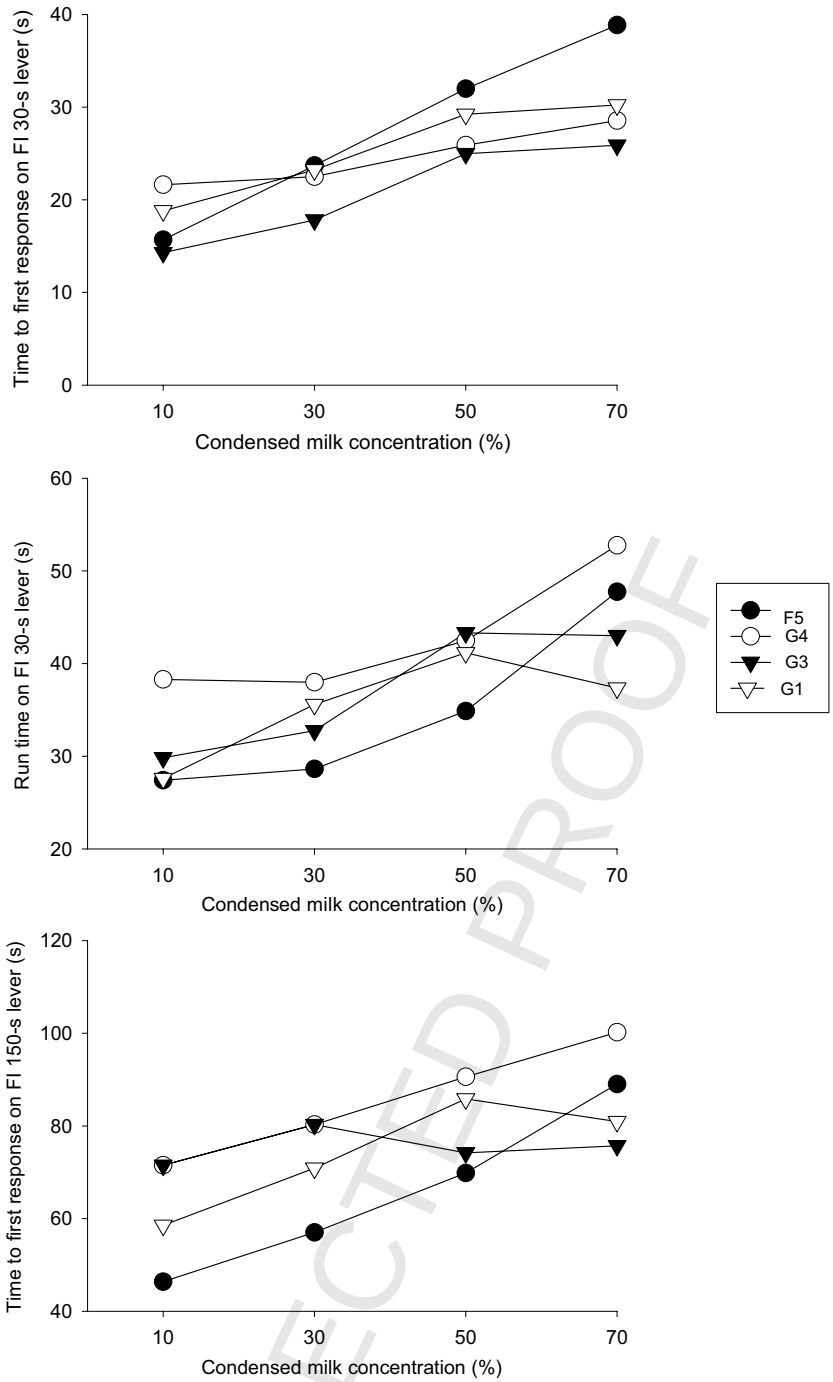

Fig. 1. Upper panel: time to the first response (start time, or porstreinforcement pause, in s) on the FI 30-s lever as a function of reinforcer concentration that started the interval. Data are shown separately for individual rats (see key for identifiers). Centre panel: time spent responding on the FI 30-s lever (time between the last response on the lever in each interval and the first one), plotted against reinforcer concentration. Lowest panel: time to the first response on the FI 150-s lever, plotted against reinforcer concentration.

reinforcement pause, start time, or wait time) as a function of the reinforcer concentration that began the interval. For all animals, mean wait times increased with increasing reinforcer concentration, and for all animals the increase in wait time was monotonic with increases in reinforcer concentration. Mean wait times were 17.6 s (10%), 21.8 s (30%), 28.01 s (50%) and 30.9 s (70%), thus not only did the mean wait times nearly double on average from a 10% concentration to a 70% one, but mean wait times at the two highest concentrations were close to or exceeded the reinforcement time, 30 s.

The centre panel of Fig. 1 shows the average run time (the time from the first to the last response on the FI 30-s lever in each trial) plotted against the reinforcer concentration. All rats showed increases in run time with increasing reinforcer concentration, but these increases were not always monotonic for individual rats. The bottom panel of Fig. 1 shows the time of the first response on the FI 150-s lever (the only data we have from this lever). Three of the four rats showed clear (although not always completely monotonic) increases in the time of the first response, but G3 showed no effect.

Overall, therefore, analysis of the wait times on both levers, and the run time on the FI 30-s lever indicates that increases in reinforcer concentration (a) made the rats start later on the FI 30-s lever, (b) made them stay longer on this lever (although this effect was less strong) and (c) also delayed the start of responding on the FI 150-s lever, probably as a consequence of effects (a) and (b).

As well as data on when responding started and stopped on the levers, we also collected response rate data from the FI 30-s lever, plotted against elapsed time in the interval, with intervals separated according to the reinforcer concentration that began the interval. The response rates were collected in 20 successive 3-s bins, and the non-linear regression subprogram of SPSS 10 was used to fit Gaussian curves to the response rate versus elapsed time in the interval functions. The response rate in each bin was regressed against the bin mid-point (i.e. 1.5 s for the first bin, 4.5 s for the second one, and so on).

The raw data from the four rats were averaged together, and Gaussian curves fitted to the resulting response function. The top panel of Fig. 2 shows the resulting plots. The Gaussian curves were associated with r^2 values of 0.89, 0.92, 0.95 and 0.97 for the reinforcer concentrations from 10 to 70%. Response rates increased from near zero levels early in the interval to a peak later, then declined. The position of the peak was progressively displaced later in the interval as reinforcer concentration increased.

The Gaussian curves can be analysed to yield a number of parameters. One is peak location (or peak time: the elapsed time in the interval when the peak of responding identified by the Gaussian curve is found). Another is peak height, the response rate at the peak, and another is the standard deviation of the Gaussian curve, a measure of its spread. Dividing the standard deviation by the peak location yields the coefficient of variation, a measure of the relative spread of the response distribution, as peak location varies.

The lower four panels of Fig. 2 show the four measures (peak location, peak height, standard deviation, and coefficient of variation) derived from the Gaussian curves fitted to averaged data. Peak location increased markedly with reinforcer concentration, whereas peak height was more stable, although was lowest with the 70% concentration. Standard deviation of the curve remained at values between 10 and 12 over the four concentrations, but coefficient of variation declined with increasing concentration, probably as the result of dividing a little-changed standard deviation by larger and larger peak location values as reinforcer concentration increased.

Fig. 3 shows the same four measures, but this time derived from Gaussian curve fits to data from individuals rats. The r^2 values for the fits to individuals were low in two cases, both with the 10% concentration (0.68 and 0.76 for F5 and G1, respectively), but were above 0.85, and mostly above 0.90, for the other 14 cases. Inspection of the data show that (i) peak locations increased with increasing reinforcer concentration for all rats (usually monotonically), (ii) peak heights showed little change, (iii) standard deviations tended to increase slightly at the highest concentrations, and (iv) coefficients of variation declined with increasing concentration.

Data from the last five baseline sessions were available for the four measures, and Fig. 4 shows peak location, peak height, standard deviation, and coefficient of variation, for individual rats from the test sessions, as well as their average (line in each panel of Fig. 4), this time expressed as a fraction of the baseline val-

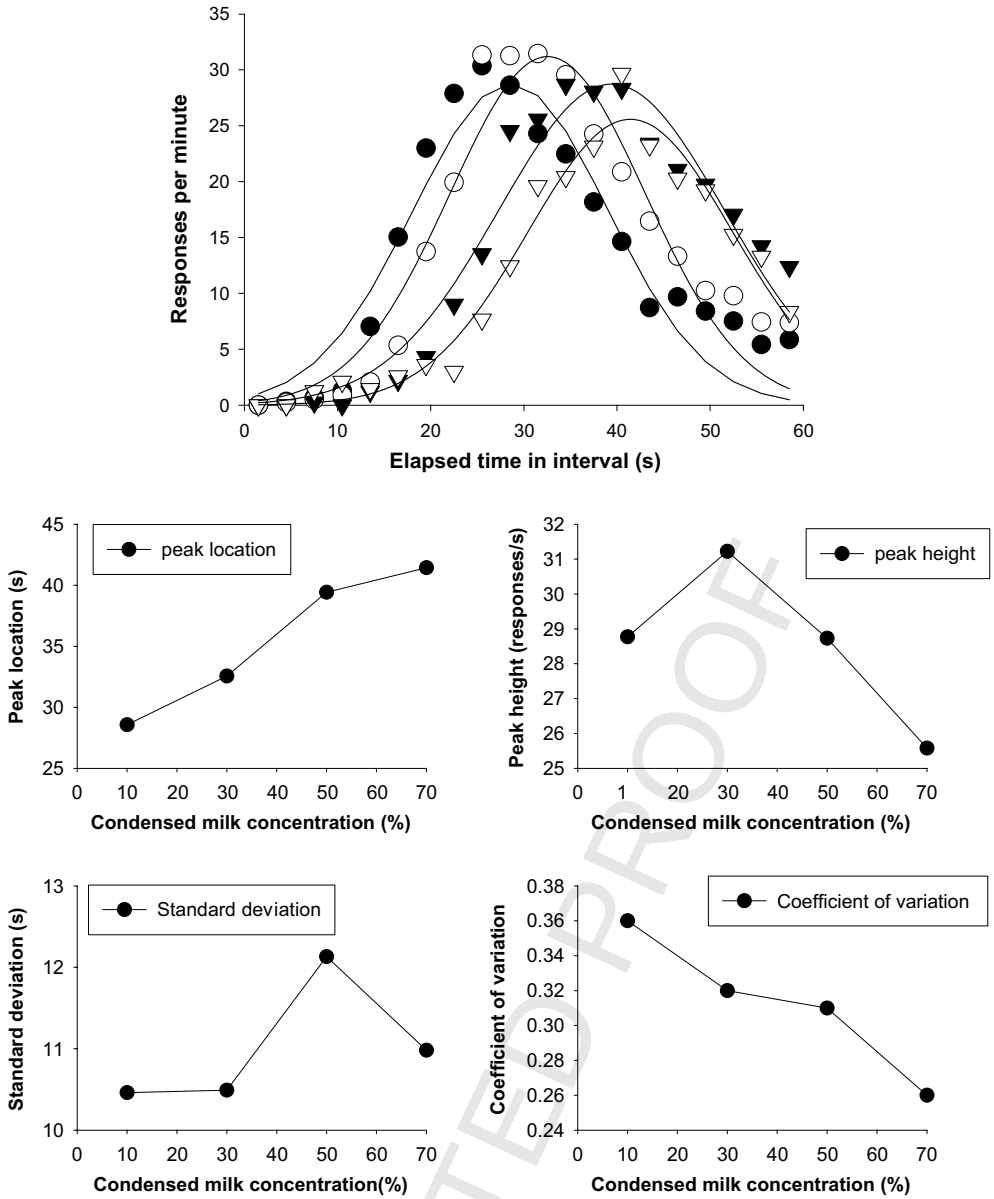

Fig. 2. Upper panel: response rate vs. elapsed time in interval functions derived from data averaged over the four rats. Data are shown separately for the different reinforcer concentrations that started the interval: filled circles, 10%; open circles, 30%; filled triangles, 50%; open triangles, 70%. Also shown are best fitting Gaussian curves for each response function (lines). Lower four panels: measures derived from the Gaussian curves shown in the upper part. Peak location: middle left; peak height: middle right; standard deviation: lower left; coefficient of variation, lower right.

ues. Consider first peak location. Although values increased with increasing reinforcer concentration, all rats showed peak location values lower than baseline levels when the reinforcer concentration was 10% (i.e. also lower than baseline values). Peak heights, on the other hand, remained close to baseline levels at concentrations up to 70%, where they were then lower in 3/4 rats. Standard deviations were close to baseline

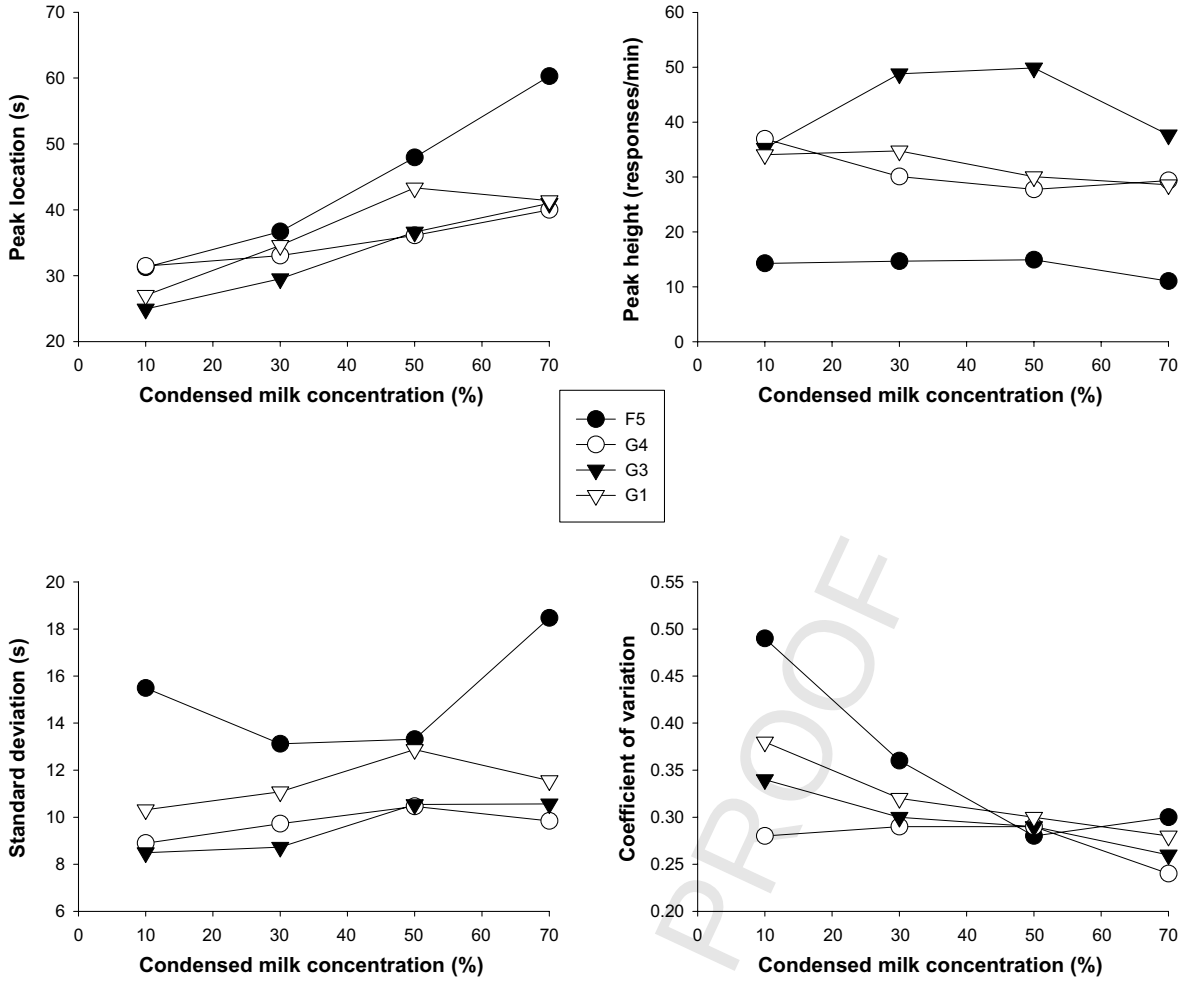

Fig. 3. Measures derived from fits of Gaussian curves to data from individual rats. Measures shown are peak location (upper left), peak height (upper right), standard deviation (lower left), and coefficient of variation (lower right). Data are shown separately for individual rats (see key).

levels at reinforcer concentrations of 10 and 30% but slightly higher on average (10 or 20% higher) at reinforcer concentrations of 30 and 50%. Coefficients of variation were above baseline levels at the 10 and 30% concentration, but closer to baseline at 50 and 70%.

5. Discussion

Our results showed very marked effects of reinforcer concentration on some measures of operant responding but not others. The post-reinforcement pause, start time or wait time (Fig. 1) and peak location (Figs. 2–4) showed the most marked effect, whereas peak response rate (peak height) showed only a small change, and then only at the 70% concentration. The absolute measure of variability of the Gaussian curve, standard deviation, showed a slight increase with increasing concentration for some rats, whereas coefficient of variation showed a clear decrease.

Some of the effects noted in our data are probably secondary to others, and there are two sorts of these. The first sort are secondary because of calculation. For example, the decline in coefficient of variation

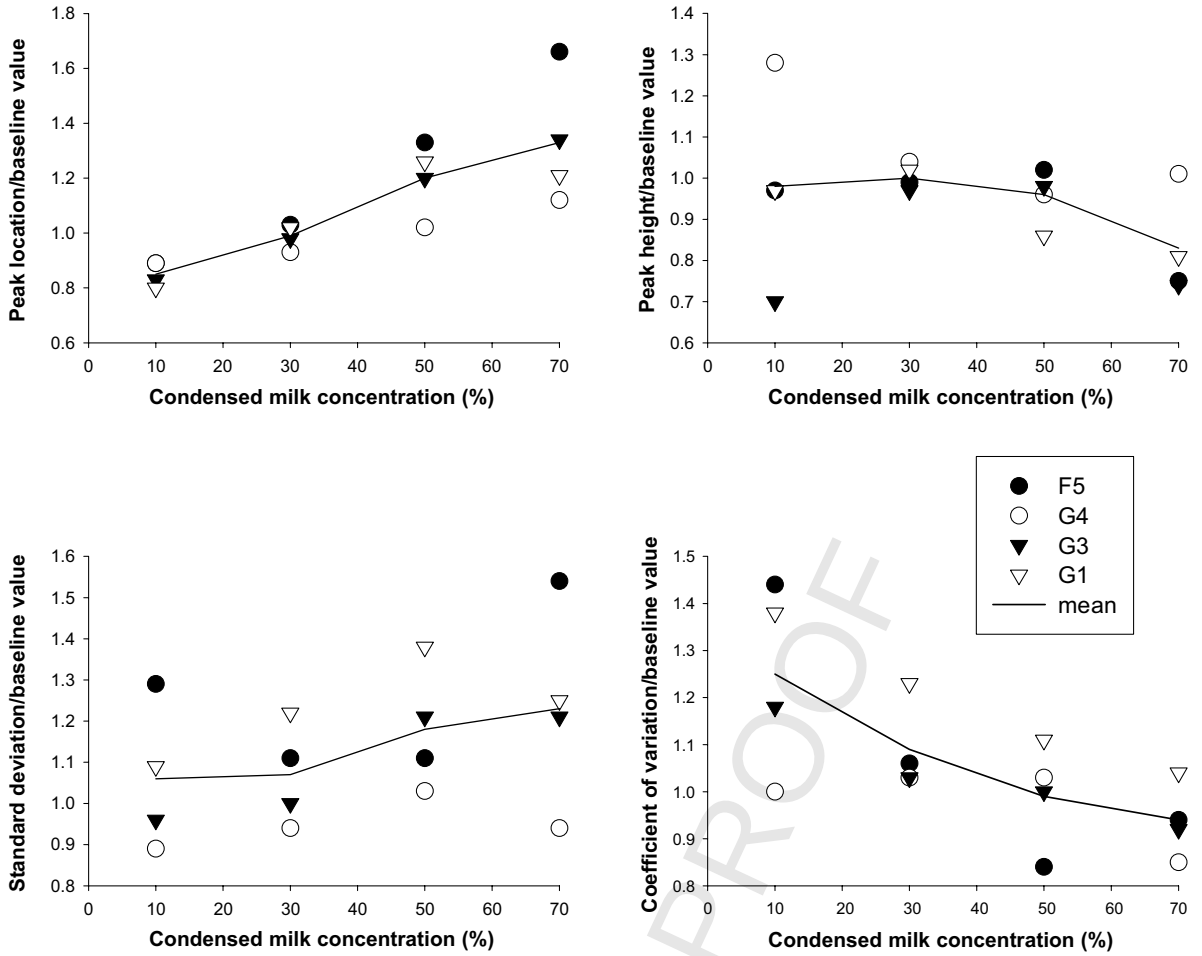

Fig. 4. Measures derived from fits of Gaussian curves to data from individual rats, expressed as a fraction of that measure during baseline. Measures shown are peak location (upper left), peak height (upper right), standard deviation (lower left), and coefficient of variation (lower right). Data are shown separately for individual rats (see key). The line within each panel shows the average of the four rats.

with increasing concentration, implying *more* precise timing, is probably a secondary effect of dividing a nearly constant or very slightly increasing standard deviation (implying *less* precise timing with increasing concentration) by a markedly increasing peak location. Similarly, if the time to start responding on the FI 30-s lever (upper panel of Fig. 1), and the time spent responding on it (centre panel of Fig. 1), both increase with increasing reinforcer concentration, then the time to start responding on the FI 150-s lever (lowest panel of Fig. 1) seems likely to increase as a consequence.

Other changes might be secondary for a different reason. Increasing reinforcer concentration increased peak location and the time to start responding. The start times at the highest concentrations were in fact on average close to the actual time of reinforcer delivery on the FI 30-s lever, so with variability would sometimes have been higher, possibly considerably higher. Although our data come only from intervals in which the reinforcer was delivered for responses on the FI 150-s lever, there is no reason to suppose that the behaviour before reinforcer delivery would have been different on the intervals where the reinforcer was assigned to the FI 30-s lever, so on some of these intervals, the long start times resulting from the high reinforcer concentration would mean that the actual time

of reinforcement experienced was (a) greater than 30 s and (b) variable from trial to trial around values above 30 s. This has two implications. One is that some of the changes we observed in peak location and start time may have been a secondary consequence of experiencing reinforcer delivery at times substantially greater than 30 s; the other is that possible variability in times of reinforcer delivery may have caused the animals to persist longer on this lever than normal, as the reinforcer was available at several potential times above 30 s, so producing both the run time increase, and the slight increase in standard deviation of the Gaussian curves observed for some rats.

It may be, in fact, that the main effect of reinforcer concentration is just to change the start of the timing process, which is otherwise unaffected. So, for example, the timing process starts earlier than baseline times at 10%, but increasingly later at 50 and 70%.

How would some contemporary models of timing account for the present results? A problem is that only one addresses the effect of reinforcer concentration directly, so there is a problem of deriving predictions from others, but some previously published material offers suggestions as to what these predictions would be.

SET (Gibbon et al., 1984) derives timed behaviour from an interaction of clock, memory and decision processes. SET proposes a timing process where times of reinforcer delivery are stored with average accuracy, and a variance process that produces standard deviations which are proportional to the time of reinforcer delivery. On an FI-like schedule, response rates increase with elapsing time in the interval on average because the time in the interval currently measured by the internal clock becomes increasingly close to the stored time of reinforcement as the interval elapses. When the time indicated by the clock and the stored time are "close enough", the animal responds.

The function of the reinforcer is to reset the clock back to zero for the next interval, so "incomplete reset", where the clock is not reset back to zero but still contains some residual time value, is a possibility. Mellon et al. (1995) discuss the well-known reinforcement omission effect on FI in these terms. In a reinforcement omission study, the reinforcer that normally terminates one interval of FI and starts the next one is either omitted altogether, or replaced by a brief signal. Complete omission of the reinforcer leaves the clock running, so responding persists past the normal time of reinforcer delivery on FI (see Mellon et al., Fig. 6, p. 295). When the reinforcer is replaced by a brief signal, the usual finding is that the animal pauses after the brief signal, but pauses for a shorter time (i.e. resumes responding earlier in the interval) than on normal FI intervals. Mellon et al. (1995) interpret this in terms of "incomplete reset" of the clock, so when the brief stimulus is delivered the "accumulator begins with some nonzero value, and reaches its criterion time [. . . for initiating responding . . .] earlier" (p. 296 with our addition in parenthesis).

The incomplete reset account may deal with the leftward shift of the peak, compared to baseline, when the reinforcer concentration was less than baseline (e.g. Fig. 4). This result closely resembles that obtained when a brief stimulus is substituted for reinforcer delivery and, as in our data, responding begins earlier on average, and peaks earlier (cf. Mellon et al., 1995, Fig. 7, p. 295). It is perhaps suggestive that for two of our rats, the Gaussian fits at the 10% concentration were notably poorer than at higher concentrations, perhaps indicating a disturbance of temporal control of responding as a result of incomplete reset.

The problematical result for SET would seem to be the systematic increase in peak location when the reinforcer concentration is higher than during baseline. Presumably, the clock cannot be reset "below zero", so some additional mechanism is needed to explain why the peaks are shifted rightward in the high concentration conditions.

However, there are some possible similarities between our reinforcer-concentration effects, and results reported from Experiment 3 of Roberts (1981) and in an experiment by Meck and Church (1987). Roberts (1981) used a peak-interval procedure with a reinforcement time of 40 s, and in Experiment 3 pre-fed his rats half their daily food ration about 40 min before the experimental session. This pre-feeding operation decreased peak rate very markedly (roughly halving it, see Roberts' Fig. 9, p. 251), but also shifted the peak time slightly to the right (from 42.6 to 44.7 s), at least over the first few sessions with pre-feeding. Meck and Church (1987) used a slightly different pre-feeding operation, where rats were given "snacks", which were intended to load them with specific nutrients rather than partially satiate them, then tested the rats on a peak-interval procedure with a 20-s re-

inforcement time. Lecithin (choline), casein (protein), and sucrose (carbohydrate) snacks were used in different conditions, and for our purposes the effect of the carbohydrate snack is the most relevant. This manipulation shifted the peak location to the right (from 18.1 to 21.7 s), with little change in peak rate (see their Fig. 3, p. 471), and the effect of the "snack" wore off gradually with continued training.

Both Roberts (1981) and Meck and Church (1987) interpreted their operations as producing a slight slowing in pacemaker speed. This interpretation is consistent with both (a) the initial shift to the right of the timing function and (b) the fact that the effect gradually wore off. While the effects reported in these experiments are reminiscent of the ones we found, some caution is needed before direct parallels can be drawn. For one thing, our peak location shifts are much greater than those reported by Roberts and Meck and Church (around 50% between the 10 and 70% concentrations, rather than the 5–20% shifts noted by the other authors). For another, if our higher reinforcer concentrations act the like the "snacks" in Meck and Church's study, then their effects must be due to some sort of learning mechanism, as the reinforcer delivered at the start of an interval is not immediately absorbed by the body, so cannot produce any direct change in bodily state. A speculative proposal is that our concentrated reinforcer conditions evoked memories of carbohydrate consumption which acted like "associative snacks", and thus slowed the pacemaker of the internal clock down, producing effects similar to those found by Meck and Church (1987). Even if this possibility is entertained, however, the very strong effects of intermixing reinforcer concentrations within a single session, as in our study, compared with the pre-feeding manipulations, remains to be explained.

From the point of view of variance changes, SET assumes that the spread of response functions, like our Gaussian curves, should be proportional to their means, thus giving rise to a constant coefficient of variation. Our declining coefficients of variation seem to clearly contradict SET but, as noted above, the decline may be a secondary consequence of the shift of the peaks. One characterisation of our data is that the timing process is essentially little changed by reinforcer concentration: it just starts later as concentration increases. Standard deviations increase for some rats because some times of reinforcement greater than 30 s are experienced in the higher concentration conditions, a result which is generally consistent with SET predictions.

Killeen and Fetterman's (1988) BeT derives measured operant behaviour from a sequence of (usually unobserved) adjunctive states. So, each adjunctive behaviour is succeeded by the next one, until one cues the measured operant response. Transition from one adjunctive behaviour to another is assumed to be governed by a pulse from a pacemaker, and this pacemaker pulses more rapidly in conditions with higher arousal (i.e. those associated with higher rates or probabilities of reinforcement). Such an incentive-based account might be expected to make predictions about reinforcer magnitude effects, as Fetterman and Killeen (1991) themselves state "If magnitude and probability of reinforcement act like rate of reinforcement, the more favorable conditions should increase the rate of the pacemaker" (p. 231).

At first sight, BeT's predictions seem to be the opposite of the results found in our experiment: if reinforcer concentration is a reinforcer magnitude manipulation, then higher magnitudes will be expected by the animal after receiving a higher concentration than a lower one, so the pacemaker should run faster after a reinforcer with higher concentration relative to a lower one. Thus, the sequence of adjunctive states is traversed more rapidly with higher concentrations, so responding occurs earlier with these, and later with lower ones. In consequence wait times should be shorter, and peak locations earlier, with higher concentrations than with lower ones, the opposite of the results we obtained. A further consequence of changing pacemaker rates is changing variability: standard deviations of response functions will be smaller with higher pacemaker rates. In our case, therefore, we would expect standard deviations to decline with increasing reinforcer concentration.

Fetterman and Killeen (1991) conducted a manipulation where hopper duration was varied across conditions with pigeons. Their main measure was standard deviation, and they found small effects in their predicted direction (their Fig. 3, p. 236). Fetterman and Killeen (1995) also used a categorical timing procedure which resembles our two-lever mixed FI schedule, except they used three response keys, each with a different associated FI (e.g. FIs of 8, 16 and 32 s). If we consider responding on the key associated with

the "middle" FI (16 s in our example), this increased with elapsed time in the interval to a peak near 16 s, then declined thereafter, a pattern like that shown in our data (upper panel of Fig. 2). Fetterman and Killeen (1995) varied probability of reinforcement in the experimental situation, and found shifts in the peak on the 16-s key that were in accord with their predictions: the peak shifted to the right when reinforcement probability was decreased between conditions, and to the left when it was increased. These results are consistent with higher pacemaker rates in conditions with higher reinforcement rates.

Although there is some evidence that reinforcement probability or rate will change behaviour in the way predicted by BeT (e.g. Fetterman and Killeen, 1991, 1995), and some slight evidence that reinforcer amount operates in the same way (Fetterman and Killeen, 1991), our data show effects which are the opposite of those BeT predicts. Start times and peak locations are later when reinforcer magnitude is increased, and standard deviation, if changed at all by reinforcer concentration, is slightly increased rather than decreased.

Machado's (1997) version of BeT (the LeT model) does not explicitly link transition between its "states" to motivational factors in the way that BeT does, although the rate of transition is "proportional to the overall density of food", as in BeT (Machado, 1997, p. 258). This suggests that LeT either has no specific predictions about our reinforcer concentration effects to contradict, or would encounter the same apparent difficulties as BeT.

One way that BeT (and possibly LeT) could accommodate our findings would be to propose that, although the rate of the pacemaker was increased by increasing reinforcement magnitude, the *number* of adjunctive states was also increased, and this increase more than offset the changes in pacemaker rate. So for example, increasing numbers of adjunctive states would occur as reinforcer concentration increased, thus the terminal state that cued the operant response would occur later and later, even though the pacemaker was running faster and faster. Although this is a possible account it is unclear why reinforcer concentration should have this effect, particularly when, as in our study, it is varied within experimental sessions. Perhaps a more promising approach would be to argue that reinforcer concentration increases prolonged the duration of the immediately post-reinforcement adjunctive state or states. For example, suppose that the adjunctive state after reinforcer delivery is related to the consumption of the reinforcer, and that this is prolonged if the reinforcer concentration is increased, then such a change would tend to shift response functions to the right with increasing concentration, as we observed. Although there is some evidence that increasing reinforcer magnitude can increase the duration of post-reinforcement activities (e.g., Pitts and Malagodi, 1996), this argument may not apply well to our data for two reasons. Firstly, the quantity of the reinforcer was kept constant, so presumably the consumption time remains fairly constant as well. If reinforcer magnitude had been varied by giving, say, 10 food pellets instead of 1, then the time taken the animal to consume the pellets (and consequent prolongation of early adjunctive activities) would presumably increase more. Secondly, the magnitude of our effects was large, with average peak locations being 50% higher with the 70% reinforcer concentration than with 30%, and start times of some individual rats more than doubling. Such a marked effect seems inconsistent with the prolongation of some presumably brief post-reinforcement consummatory state.

One account which fares better with our results is the "Tuned-trace" model of Staddon et al. (2002). This model, like its predecessor the multiple time scales (MTS: Staddon and Higa, 1999) model, effectively assumes that timing is accomplished by decaying memory of the reinforcer (the "trace"). More potent reinforcers produce larger initial traces so if the animal initiates responding when the trace strength reaches some threshold level, it will take longer to reach this level after larger reinforcers than smaller ones, so timing functions should be systematically displaced, to the right if the reinforcer is larger than "normal" and to the left if the reinforcer is smaller. Staddon et al. (2002) in fact discuss unpublished data which suggest that intermixing reinforcer durations within an experimental session should have this effect (see their Fig. 7, p. 113, and their discussion, pp. 113–114). In addition, two earlier experiments (Lowe et al., 1974; Staddon, 1970) varied reinforcer magnitude within sessions on simple FI, and found that the postreinforcement pause (start time) increased systematically with increasing reinforcer magnitude. These results closely resemble our own (upper panel of Fig. 1), as does the result that

effects of reinforcer magnitude on response rates were much less marked.

Staddon et al.'s model has the advantage compared with others that it naturally predicts the pattern of results we obtained: both increased wait times and, presumably, later response peaks will be occasioned by reinforcers of greater magnitude than experienced in training. It is unclear what prediction this model makes about response variability, so our finding of near constant, or only very slightly increasing, standard deviation with increasing reinforcer concentration, may or may not be compatible with Staddon et al.'s (2002) position.

The effects of motivational manipulations, like changes in reinforcer value, were once standard benchmarks against which theories of animal behaviour were tested (see Wearden, 1989, for discussion). With some exceptions, notably those deriving from tests of BeT (Fetterman and Killeen, 1991, 1995), and tests of the Matching Law (Heyman and Monaghan, 1994), motivational effects have been less frequently studied in recent years, but the present data show that reinforcer concentration, to name but one, has consistent and strong effects on timed behaviour on an FI-like schedule. Attempts to theoretically account for such effects might well provide ways of distinguishing the predictions of different current theories of timing in animals, predictions which can be difficult to test under normal conditions.

References

Blomeley, F.J., 1984. Determinants of temporal control in operant performance. Unpublished Ph.D. thesis, University of Wales, Bangor.

Church, R.M., Meck, W.H., Gibbon, J., 1984. The application of scalar timing theory to individual trials. J. Exp. Psychol. Anim. Behav. Process. 20, 135–155.

Fetterman, J.G., Killeen, P.R., 1991. Adjusting the pacemaker. Learn. Motiv. 22, 226–252.

Fetterman, J.G., Killeen, P.R., 1995. Categorical scaling of time: implications for clock-counter models. J. Exp. Psychol. Anim. Behav. Process. 21, 43–63.

Gibbon, J., Church, R.M., Meck, W., 1984. Scalar timing in memory. In: Gibbon, J., Allan, L. (Eds.), Annals of the New York Academy of Sciences, Vol. 423: Timing and Time Perception (pp. 52–77). New York Academy of Sciences, New York.

Heyman, G.M., Monaghan, M.M., 1994. Reinforcer magnitude (sucrose concentration) and the matching law theory of response strength. J. Exp. Anal. Behav. 61, 505–516.

Hopson, J.W., 2003. General learning models: timing without a clock. In: Meck, W.H. (Ed.), Functional and Neural Mechanisms of Interval Timing (pp. 24–60). CRC Press, Boca Raton, FL.

Killeen, P.R., Fetterman, G.J., 1988. A behavioral theory of timing. Psychol. Rev. 95, 274–295.

Leak, T.M., Gibbon, J., 1995. Simultaneous timing of multiple intervals: implications of the scalar property. J. Exp. Psychol. Anim. Behav. Process. 21, 3–19.

Lejeune, H., Wearden, J.H., 1991. The comparative psychology of fixed-interval responding: some quantitative analyses. Learn. Motiv. 22, 84–111.

Lejeune, H., Cornet, S., Ferreira, J., Wearden, J.H., 1998. How do mongolian gerbils (*Meriones unguiculatus*) pass the time? Adjunctive behavior during temporal differentiation in gerbils. J. Exp. Psychol. Anim. Behav. Process. 24, 325–334.

Lowe, C.F., Davey, G.C., Harzem, P., 1974. Effects of reinforcement magnitude on interval and ratio schedules. J. Exp. Anal. Behav. 22, 553–560.

Machado, A., 1997. Learning the temporal dynamics of behavior. Psychol. Rev. 104, 241–265.

Meck, W.H., Church, R.M., 1987. Nutrients that modify the speed of internal clock and memory processes. Behav. Neurosci. 101, 465–475.

Mellon, R.C., Leak, T.M., Fairhurst, S., Gibbon, J., 1995. Timing processes and the reinforcement omission effect. Anim. Learn. Behav. 23, 286–296.

Morgan, L., Killeen, P.R., Fetterman, J.G., 1993. Changing rates of reinforcement perturbs the flow of time. Behav. Process. 30, 259–272.

Pitts, R.C., Malagodi, E.F., 1996. Effects of reinforcement amount on attack induced under a fixed-interval schedule in pigeons. J. Exp. Anal. Behav. 65, 93–110.

Roberts, S., 1981. Isolation of an internal clock. J. Exp. Psychol. Anim. Behav. Process. 7, 242–268.

Skinner, B.F., 1938. The Behavior of Organisms. Appleton-Century-Crofts, New York.

Staddon, J.E.R., 1970. Effect of reinforcement duration on fixed-interval responding. J. Exp. Anal. Behav. 13, 9–11.

Staddon, J.E.R., Higa, J., 1999. Time and memory: towards a pacemaker-free theory of interval timing. J. Exp. Anal. Behav. 71, 215–251.

Staddon, J.E.R., Chelaru, I.M., Higa, J.J., 2002. A tuned-trace theory of interval-timing dynamics. J. Exp. Anal. Behav. 77, 105–124.

Wearden, J.H., 1989. Mysteries of the organism: Clark L. Hull's principles of behavior and some problems in contemporary schedule theory. J. Exp. Anal. Behav. 51, 277–286.

Whitaker, J.S., Lowe, C.F., Wearden, J.H., 2003. Multiple-interval timing in rats: performance on two-valued mixed fixed-interval schedules. J. Exp. Psychol. Anim. Behav. Process. 29, 277–291.