

This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>

When to respond? And how much? Temporal control and response output on mixed-fixed-interval schedules with unequally probable components[☆]

S. Whitaker¹, C.F. Lowe², J.H. Wearden^{*}

School of Psychology, Keele University, Keele, Staffordshire ST5 5BG, United Kingdom

Received 21 December 2006; received in revised form 21 April 2007; accepted 6 June 2007

Abstract

Rats were trained on mixed-fixed-interval (FI) schedules, with component FIs of 30 and 60 s. The probability of reinforcement according to FI 30 s varied between conditions, across values of 0.1, 0.3, 0.5, 0.7 and 0.9. When response rate in the 60 s intervals was measured, separate response peaks, one close to 30 s, the other at 60 s, could be identified when the probability of reinforcement at 30 s was 0.3 or greater. Nonlinear regression found that the location of the earlier peak was always close to 30 s, that the coefficient of variation of the response functions at 30 and 60 s were unaffected by reinforcement probability, but that the 30 s component appeared to be timed slightly more precisely than the 60 s one. Response rate at around 30 s increased with increasing probability of reinforcement according to FI 30 s, but responding at 60 s was unaffected by reinforcement probability. The data are discussed with respect to a number of contemporary models of animal timing (scalar expectancy theory, the Behavioural Theory of Timing and the Learning to Time model), and a recent account of response output on FI-like schedules.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Mixed-FI schedules; Probability of reinforcement; Temporal control

1. Introduction

When to respond? And how much? These two questions address two of the most fundamental issues in the analysis of operant behaviour. The first poses the question of how the temporal distribution of responses should vary as a function of temporal periodicities or temporal constraints imposed by schedules. The second asks how much response output should occur, usually as a function of the rate or probability of reinforcement for that response. With a few exceptions, discussed in detail later, the two issues have been kept separate in theoretical terms, with the first addressed by theories of temporal control, such as scalar expectancy theory (SET: Gibbon et al., 1984), and its more recent competitors, discussed later. The second issue, that of how much

to respond, has generally been related to principles such as the Matching Law (Herrnstein, 1961), or Herrnstein's hyperbolic equation (Herrnstein, 1970).

Studies of *when* to respond have usually employed schedules exposing the animal to temporal periodicities or constraints. For example, on a fixed-interval (FI) schedule, the first response occurring *ts* after some previous time marker (usually the previous reinforcer) is itself reinforced. This temporal periodicity produces temporally regulated behaviour in a range of animal species (e.g. Lejeune and Wearden, 1991), with responding being at low or zero levels after the time marker, and on average increasing throughout the interval to reach a peak at or near the time of reinforcer delivery. In studies of FI schedules, and other schedules inducing temporal control of responding, the question of absolute or relative response rates (e.g. how much to respond) is rarely treated, and behavioural measures are often presented in relative terms. For example, the response rate at different elapsed times during an FI interval can be expressed as a fraction of the maximum response rate, thus "normalizing" the responding that occurs in terms of the peak response rate for that FI value (e.g. see Lejeune and Wearden, 1991, and Church et al., 1994, for examples). On the other hand, studies of the "how much"

[☆] The data reported here were collected by the first author as part of the requirements for the degree of Ph.D., awarded in 1980 by the then-titled University College of North Wales, Bangor, United Kingdom.

^{*} Corresponding author.

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

¹ The author is now at the University of Huddersfield, United Kingdom.

² The author is now at the University of Wales, Bangor, United Kingdom.

question are dominated by the use of aperiodic reinforcement schedules such as variable-interval (VI) or random-interval (RI). Such schedules set up and deliver reinforcers at highly variable or random times since the previous reinforcer delivery, so have no temporal periodicities for the animal to detect. Responding consequently proceeds at an approximately constant rate over time, but the overall response rate can be related to the overall rate of reinforcement, as in [Herrnstein \(1970\)](#).

The present article discusses data obtained from two-valued mixed-FI schedules, usually with components of unequal probability. In such a schedule, reinforcers are delivered either according to FI S or FI L (where $L > S$; from now on we will just refer to the components as S and L , as in all cases in our study the schedules are FI), where nothing signals to the animal which schedule is in force on any particular interval. Why should mixed-FI schedules be of any interest in addressing issues about timing of behaviour and response output? There are basically two reasons. The first is that the mixed-FI schedule offers a number of procedural possibilities, not otherwise usually available, for dissociating variables that might influence response timing (such as the time at which reinforcers are normally available for responses) from those which influence response output (such as the rate or probability of reinforcement). For example, consider a two-component mixed FI schedule where S is FI 30 s. Here, at least in some intervals, the first response emitted 30 s or more after the start of the interval is reinforced. The 30 s time of potential reinforcement can be kept constant, when the overall rate of reinforcement in the experimental situation is varied (by varying L : see [Whitaker et al., 2003](#), for examples), or the probability of reinforcement at 30 s can be varied, as in the present study. In an ordinary FI schedule, the rate of reinforcement and the FI value are necessarily confounded, and the probability of reinforcement on the FI schedule cannot be varied much without disturbing the temporal periodicity of the schedule (although occasional omission of reinforcers on FI is a frequently used technique, see [Staddon and Innis, 1969](#), for the original study and [Mellon et al., 1995](#), for a modern replication and discussion). On mixed FI, on the other hand, the potential times of reinforcement, S and L , can be maintained while the probabilities of reinforcement at S and L are varied.

The second reason why mixed-FI schedules are interesting is that some recent theories of animal timing either make specific predictions about timing and responding on mixed FI ([Leak and Gibbon, 1995](#); [Machado, 1997](#); [Whitaker et al., 2003](#)), or have particular difficulties in dealing with behaviour on such schedules ([Dragoi et al., 2003](#)).

The present article reports data obtained from two-component mixed-FI schedules where the components were 30 and 60 s. The present study varied the probability of reinforcement according to the two component schedules, so enabling potential dissociation of variables affecting *timing* of responding (e.g. the times of reinforcement, determining *when* to respond) from variables affecting the *rate* of responding (e.g. reinforcement probability, determining *how much* to respond).

Like the data in [Whitaker et al. \(2003\)](#), and [Blomeley et al. \(2004\)](#), the results in the present study have the peculiarity

that they were collected 25 or more years ago (in the present case for the thesis of [Whitaker, 1980](#)), although the experiments use completely standard procedures, and the data are otherwise unpublished. The technology available at the time the experiments were conducted imposed some limitations on the measures that were taken, and some measures of behaviour that might now be common cannot be derived from the data set.

The data in the present article come from an experiment in which three rats were exposed to mixed FI 30 s FI 60 s, where the probability of reinforcement at S (30 s) and L (60 s) was varied over values of 0, 0.1, 0.3, 0.5, 0.7, 0.9 and 1.0. Obviously, the schedules used enable observation of the effects of reinforcement probability, and correlated variables like reinforcement rate, on the timing of behaviour and response output, when times of reinforcement were kept constant at 30 and 60 s. One possibility is that timing and response output are independent: that is, manipulation of reinforcement probability affects one variable (probably response output), but not the other (see [Roberts, 1981](#), for this suggestion). Such an outcome would be consistent with predictions of SET ([Gibbon et al., 1984](#)), although perhaps only by default (as [Whitaker et al., 2003](#), point out) as SET makes no predictions at all about response rate. In contrast, [Killeen and Fetterman's \(1988\) Behavioral Theory of Timing \(BeT\)](#) links the timing process itself to the rate of reinforcement in the experimental situation, so in general timing would be more relatively precise when higher rates of reinforcement were provided. BeT, like SET, is a theory concerned solely with the temporal control of behaviour so, again like SET, makes no predictions about how response output should vary with changes in reinforcement rate or probability. A more recent theory, the *Learning to Time* model (LeT) of [Machado \(1997\)](#) not only makes explicit predictions about behaviour on mixed-FI schedules, but also predicts that response output on such schedules should vary as the probability of reinforcement is changed. In particular, if the probability of reinforcement for responses at S is varied (as in both our experiments) then the rate of responding at S should increase or decrease with increases or decreases in reinforcement probability. We will discuss in detail how the different timing models deal with our results later, but we will begin with a description of performance.

2. Method

2.1. Animals

Three experimentally naïve male hooded rats, 12 weeks old at the start of the experiment, served. They were individually housed and maintained at 80% of their free-feeding weight throughout the study. Water was freely available in the home cages.

2.2. Apparatus

Three Lehigh Valley model 42–45 operant chambers were used. These were located within sound-attenuating cubicles, with an exhaust fan mounted at the back. The ambient noise level was 60 dB. The reinforcers used were 45 mg Noyes pel-

lets. All experimental events were controlled by a Data General Nova 1200 computer, which also recorded responses.

2.3. Procedure

As the animals were naïve, lever-pressing was shaped in the first session, and the rats were allowed to obtain 60 reinforcers, with each lever-press reinforced. Following this, they were exposed to simple FI and mixed-FI schedules. All mixed-FI schedules were mixed FI 30 s FI 60 s, but the difference between mixed FI conditions resided in the probability of reinforcement at FI 30 and FI 60 s. All intervals ended with reinforcer delivery, so the reinforcement probabilities at 30 and 60 s summed to one. We describe the mixed FI schedules in terms of the probability (p) of reinforcement according to FI 30 s, so the probability of reinforcement at FI 60 s was $1 - p$. The number of sessions in each condition is shown in brackets, and sometimes differed slightly between rats. The order of the schedules was: $p = 0.5$ (56–60); $p = 1$ [FI 30 s] (37–38); $p = 0.7$ (30–32); $p = 0.3$ (40–41); $p = 0$ [FI 60 s] (31–32); $p = 0.9$ (36–38); $p = 0.1$ (25–26); $p = 0.7$ (28); $p = 0.1$ (25–26). Sessions lasted for 60 reinforcer deliveries.

3. Results

The data reported here come from the last three sessions of each condition, and when a condition was repeated, data from the two exposures were aggregated together. Within each experiment, the data from different individual rats were averaged together to produce an averaged data sample, but the data from individual rats were also analyzed separately. The data collected were in the form of response rate versus elapsed time functions, and all the data from the mixed-FI schedules came from intervals where the reinforcer was delivered according to FI L . The response rates were collected in 20 bins, each $1/20$ th of L , so bin widths were 3 s. Data from the simple FI schedules were analogous, and the resolution was the same as that for L , so response rates were collected in 10 3 s bins for FI 30 s schedule.

Some care is needed in the interpretation of the phrase “probability of reinforcement” in the present and following sections. In the procedure above, we defined each mixed-FI schedule in terms of the probability of reinforcement at S (30 s in both cases), but probabilities of reinforcement at S and L summed to 1, so the probability of reinforcement at L was $1 - p$. However, these values are a priori probabilities of reinforcement, that is, the probability of reinforcement at S and L before a particular interval has started. All our data come from intervals where responding was reinforced according to FI L , so in fact the actual probability of reinforcement at L in these intervals was always 1.0. A priori probabilities make psychological sense, as nothing signals to the rat whether reinforcers are available for responses according to FI S or FI L , but an additional complexity is that the plots of data in some Figures are different for data relating to S and data relating to L . For example, consider a $p = 0.3$ condition. Here, the a priori probability of reinforcement for responses at S is 0.3, so the a priori probability of reinforcement for responses at L is 0.7. When data from this $p = 0.3$ condition are plotted as

a function of a priori probability of reinforcement (e.g. Fig. 2), then the data point from S is plotted against 0.3 on the x -axis, but the data point from L against 0.7.

As in Whitaker et al. (2003) and Blomeley et al. (2004) the principal method of data analysis was nonlinear regression, carried out by the “nonlinear” sub-program of SPSS versions 10 and 11. The dependent variable was the response rate in each of the 20 bins in each schedule condition, and the independent variable was elapsed time in the interval. This was represented as the mid-point of the successive bins (e.g. 1.5 s, 4.5 s, . . . etc.).

Data from simple FI schedules were fitted by a one-Gaussian function, $kG(m, S.D.)$. Here G indicates a Gaussian function, m its mean, S.D. the standard deviation and k is a scaling constant needed to fit a probability distribution to a response rate function. For all fits to data from simple FI schedules, m was fixed at the schedule value of 30 or 60 s, so the values derived were the standard deviation of the curve (S.D.), and the scaling factor (k), which together determined the peak height of the curve.

Data from mixed FI schedules were fitted by the sum of two Gaussian functions $k_S G(m_S, S.D._S) + k_L G(m_L, S.D._L)$. Here, m_S and m_L are the means of the two distributions, S.D._S and S.D._L their standard deviations and k_S and k_L are their scaling factors. Fitting the two-Gaussian function will be referred to below as a “2-Gaussian” fit, as there are two independent contributions to the response rate versus time function, one coming from schedule S , the other from schedule L . m_L was always forced to be at the real time value of L , as in Whitaker et al. (2003), but m_S was allowed to vary freely. m_L was fixed at L as values above L were rarely experienced by the animals, so were not psychologically meaningful. As will be seen later, this decision did not do violence to the data, which were well-fitted by 1- or 2-Gaussian functions in almost all cases.

Fig. 1 shows the averaged data from the different experimental conditions, and can be used to illustrate the curve-fitting method. Consider first data from the simple FI 30 s and FI 60 s schedules, shown in the lower right panel. These were fitted by 1-Gaussian functions which both had r^2 values of 0.99. For illustration of the 2-Gaussian fit, consider data in the lower left panel of Fig. 1, which come from the 0.9 condition. Inspection of the response rate function shows that it increases from near zero early in the interval, to a peak at or near FI S , only to slightly decline, then rise to another peak at L . It is easy to imagine that a 2-Gaussian function will fit such data well, with one peak near S and another at L , and the line in the panel shows the 2-Gaussian fit obtained. The curve shown produces a number of measures: (a) the peak location near S (m_S), (b) the standard deviations of the two curves (S.D._S and S.D._L), and (c) k_S and k_L , the scaling constants which can be used to calculate the peak heights of the underlying Gaussian curves.

Inspection of the data in the other panels of Fig. 1 suggests that two peaks in the response function may not always be clearly detectable by eye (e.g. the $p = 0.3$ function), but our analytical procedure was consistent. Data from mixed-FI conditions were always initially fitted by the 2-Gaussian function, whether or not 2 response peaks were evident on inspection: that is, we let the nonlinear regression program “decide” whether there were two underlying sources of control of responding, or only one.

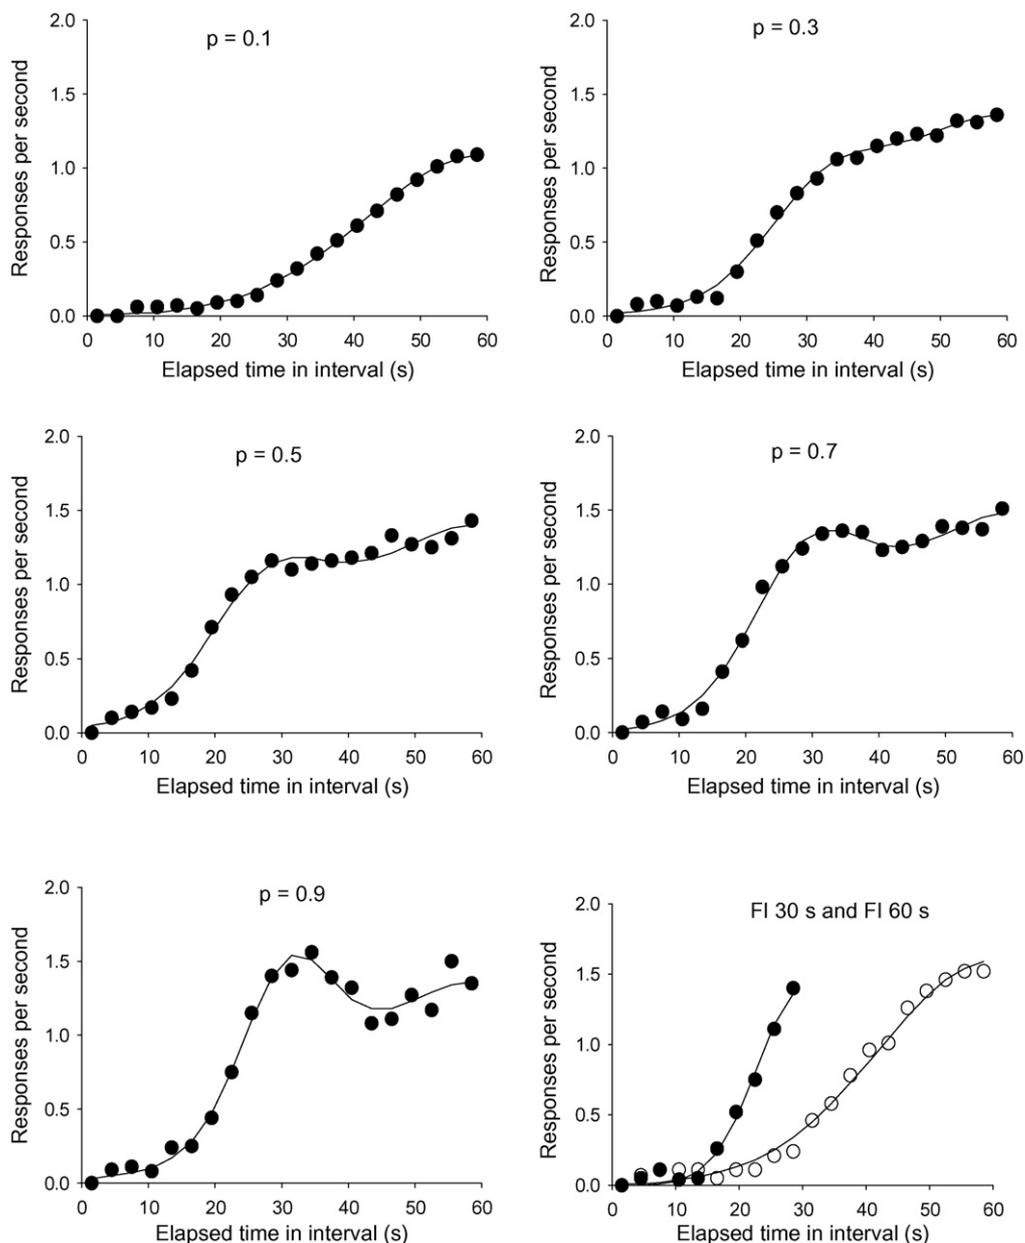


Fig. 1. Response rate vs. elapsed time in the interval functions from the single FI and mixed-FI conditions. Filled circles show the data points, and the curves within each panel shown the best fitting 2-Gaussian or 1-Gaussian fit (see text for details). The bottom right panel shows data from the simple FI schedules used (FI 30 s and FI 60 s), all other panels show data from the mixed FI 30 s FI 60 s schedules, with the probability of reinforcement at 30 s (p) indicated in each panel. For all mixed-FI schedules data come from intervals where responding was reinforced according to FI 60 s.

Data from the $p=0.1$ condition illustrate a case sometimes encountered. Here, a 2-Gaussian fit was tried initially, but no fitted function could be obtained. Rejection of a 2-Gaussian fit in such cases was based on obtaining a regression solution with a *negative* k_S value. This indicates that the putative source of behavioural control normally attributable to S was making a *negative* contribution to total responding, which is psychologically meaningless. When such conditions arose (and they were rare), the data were fitted with a 1-Gaussian function with m_L forced at L : that is, they were treated identically to simple FI L . Such a treatment fitted data well, as the curve in the top left-hand panel of Fig. 1 shows.

Consider now the data points shown in Fig. 1 in more detail. Inspection of the panels showing data from the mixed-FI conditions suggests that increasing the probability of reinforcement for responses at S progressively increased the probability of responding at around S , whereas responding at L was little affected by the probability of reinforcement at L . It should be recalled, however, that the data come from intervals where the response was reinforced according to FI L , so later in the interval the conditional probability of reinforcement at L was 1.0.

The panels also show the results of the nonlinear regression in the form of curves. The curve fitted data well in all cases ($r^2 > 0.99$). Data from the $p=0.3, 0.5, 0.7$ and 0.9 conditions

were well described by 2-Gaussian fits, whereas no 2-Gaussian fit could be obtained at $p=0.1$, which was accordingly treated as FI 60 s.

The curve fitting procedure was applied to the response rate versus time functions produced by the individual rats, and measures descriptive of temporal control from individuals, as well as averaged data, are shown in Fig. 2. All fitted functions to data from individuals had r^2 values of 0.95 or greater.

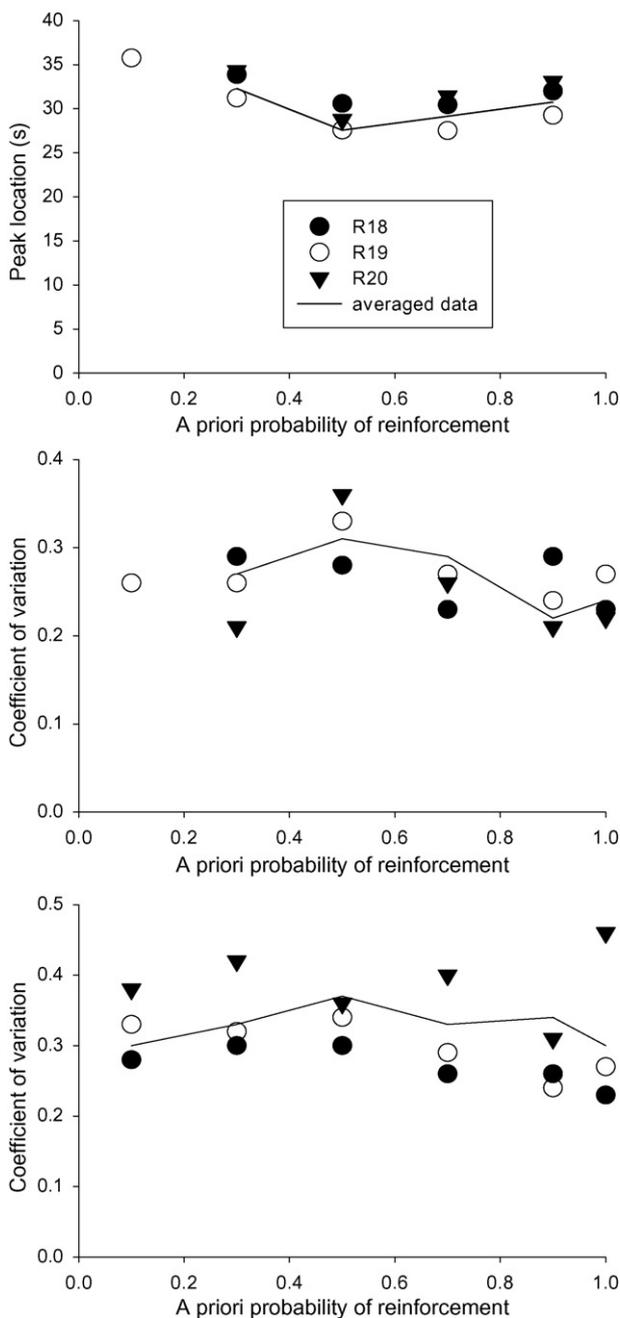


Fig. 2. Results from Gaussian curve-fitting. (Upper panel) Peak location of the peak corresponding to control by the FI 30 s schedule. Line connects results from averaged data, unconnected symbols show data from individuals. (Centre panel) Coefficient of variation of the curve related to control by the FI 30 s component. (Bottom panel) Coefficient of variation of curve related to control by the FI 60 s component. In all panels data are plotted against a priori probability of reinforcement (see text for details).

Although no 2-Gaussian fit could be obtained at $p=0.1$ for averaged data, and for two individuals (R18 and R20), R19 did show two sources of behavioural control at this probability value. Fig. 2 shows two principal measures of temporal control. The first (upper panel) is m_S , the peak location of the Gaussian curve with the lower mean. In all cases, S was 30 s, so the obvious questions are (a) where is the peak actually located and (b) was this location affected by reinforcement probability at S ? Data in the upper panel of Fig. 2 (individuals shown by unconnected symbols, averaged data by a line) show that for all individuals, and in averaged data, the peak location corresponding to a source of behavioural control by S was located close to S , but usually slightly higher than S , about 33 s. Reinforcement probability had no apparent systematic effect on peak location, either in averaged data, or in any individual.

The next two panels of Fig. 2 show the coefficient of variation of the underlying Gaussian curves, the standard deviation divided by peak location. This is a relative measure of temporal sensitivity (with smaller values indicating more sensitive timing: see Lejeune and Wearden, 1991, for discussion), similar to a Weber fraction. The centre panel of Fig. 2 shows the coefficients of variation from the curve appropriate to S , when a 2-Gaussian fit was possible: that is, the values show the precision of the timing of the 30 s component of the mixed-FI schedule. Also included, as $p=1.0$ are values obtained from the simple FI schedule, so we can compare the precision of timing when 30 s is timed singly (simple FI) and as a part of a two-component mixed-FI schedule.

Averaged coefficient of values for timing at or near S were about 0.28, and showed no systematic change with reinforcement probability at S , neither did any individual animal exhibit systematic change in timing sensitivity with changes in reinforcement probability. The sensitivity of timing of the 30 s component of the mixed-FI schedule was, furthermore, not systematically different from the value obtained when 30 s was presented alone in simple FI.

The lowest panel shows coefficient of variation values from the Gaussian curve appropriate to L (60 s), as well as the value obtained when FI 60 s was presented alone ($p=1.0$). The x -axis = 0.9 points for the two individuals for which there was no 2-Gaussian fit come from the 1-Gaussian fit to this condition (which was $p=0.1$ for reinforcers programmed according to FI S). Coefficients of variation from analysis of averaged data were just above 0.3, but showed no clear effect of reinforcement probability at L , neither in the averaged data nor in data from individuals.

Overall, therefore, the timing of behaviour can be described as follows: (a) peak locations were close to S in all cases, although were much more likely to be slightly above S than below it, (b) coefficients of variation of curves corresponding to S and L showed little or no effect of reinforcement probability, although there was a suggestion that S was timed slightly more precisely (i.e. with smaller coefficient of variation) than L , a point we will return to later.

Fig. 3 shows measures of response output. The focus of interest was response rate at around the times of S and L , and this was measured in two ways. One measure (just called response

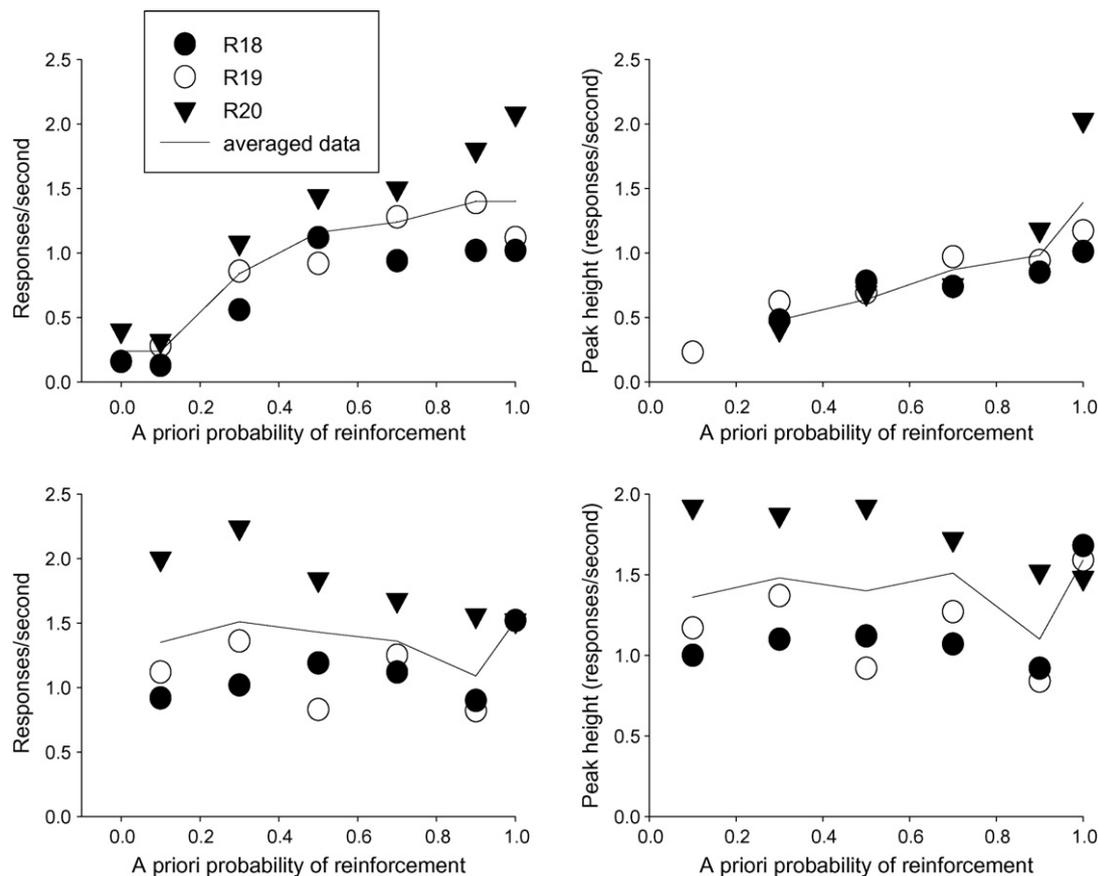


Fig. 3. (Left panels) Response rate in bin 10 (upper panel) and bin 20 (lower panel), plotted against a priori probability of reinforcement. Lines show results from averaged data, unconnected symbols show data from individuals. (Right panels) Peak heights of Gaussian curves corresponding to control by FI 30 s (upper panel) and FI 60 s (lower panel).

rate) was the observed response rate in bins 10 and 20, that is, the bins just before the programmed times of *S* and *L*. The other (peak height) was the calculated peak heights of the Gaussian curves fitted to data (in most cases coming from 2-Gaussian fits to mixed-FI conditions). As will be seen below, the two performance indices yielded very similar pictures of how response output changes with response probability, but the two measures are not the same. Measured response rate in bins 10 and 20 reflects the contribution of *both* sources of temporal control, that corresponding to *S* and that corresponding to *L*, in these bins, whereas the peak height measure separates the two sources of control.

The upper 2 panels of Fig. 3 show results from bin 10 (left panel), or from the Gaussian curve corresponding to FI *S* (right panel). In the left panel, the value at a probability of 0.0 comes from FI 60 s: including this measure enables us to observe whether providing reinforcement with some probability according to FI *S* increases response rate in bin 10 above the rate that would have occurred had the probability of reinforcement at this time been zero. The results from the response rate measure are clear: reinforcement at FI *S* with a 0.1 probability does not appear to increase responding much above what would have occurred with 0.0 probability, whereas higher probability values systematically increased responding in bin 10, with all animals showing clear (although not always monotonic) increases in

response output as the probability increased towards 1.0 (i.e. FI 30 s). The upper right panel shows peak heights calculated from the Gaussian curve. Obviously, no 0.0 probability value can be included here, and only one animal (R20) produced a 2-Gaussian fit at $p=0.1$, but otherwise the results are very similar to those obtained with the bin-10 response measure. Peak heights of the curve at near to *S* increased systematically, and nearly monotonically, for all animals, as reinforcement probability at *S* was increased. Obviously, both response measures show that responding at around *S* was systematically related to the probability of reinforcement at *S*.

The lower two panels show analogous data: response rates from bin 20 (the bin preceding 60 s: left panel), and the peak heights (right panel). In neither case did the a priori probability of reinforcement at 60 s appear to influence response output, in contrast to the situation at 30 s. No individual rat showed any trend of change of response rate as the a priori probability of reinforcement changed.

4. Discussion

Our data set is more extensive than any previously published, but a small number of other observations of the effects of varying reinforcement probability on mixed-FI and related schedules exist. For example, Catania and Reynolds (1968), in their Exper-

iment 5, presented data from pigeons on mixed-FI schedules where L was always 240 s, but S was varied over values of 30, 90, 150 and 210 s. Furthermore, the probability of reinforcement for responses at S was varied over values of 0.05 and 0.5. Their data resembled ours in a number of respects. Firstly, as S and L became closer, the response rate versus elapsed time curve developed from an obviously 2-peak structure to one closer to a single peak (their Fig. 18, p. 360). Secondly, effects of changing probability of reinforcement at S were marked, with response rates around S increasing as reinforcement probability for responses at S changed from 0.05 to 0.5 (Figs. 19 and 20, pp. 361–362). Thirdly, although responding at around S was sensitive to probability of reinforcement for responses at S , responding at L was less sensitive, even though, as in our studies, the probability of reinforcement at S and L always summed to 1.0. Fourthly, there was a suggestion in Catania and Reynolds' data that when reinforcers were available for responses at S , the peak of responding was often slightly later than S (e.g. their Figs. 20 and 21, and the present Figs. 2 and 4). No curve-fitting analyses were performed on Catania and Reynolds' data, but the results were reanalyzed by Machado (1997, p. 252), whose work is discussed further below.

Roberts (1981) introduced the *peak-interval* procedure, which is a variant of mixed FI with the two components being FI and extinction. In Roberts' Experiment 1, for example, food was available on some trials for responding according to FI 20 s or FI 40 s (in different conditions), whereas the other trials (the peak trials) lasted 80 s and terminated without food delivery. Data from the peak trials, in the form of response-rate versus elapsed time in the trial, could be analyzed in various ways to yield measures of peak location, and peak height, as in the present article (although the fitting procedures used were slightly different). In Experiment 1, Roberts varied the probability of reinforcement delivered according to FI 20 s over values of 0.2 and 0.8 (i.e. the remaining trials were peak trials). This manipulation changed peak height very markedly (with peak response rate with the 0.2 probability being about 25% of that with the 0.8 probability, see Roberts' Table 1, p. 245), whereas peak location was unaffected, and close to the programmed time of reinforcement, 20 s. Roberts carried out a number of other manipulations, and generally found that peak location was close to the time of reinforcement (although, as in our data, more likely to be slightly above it than below), and was unaffected by manipulations that kept the time when reinforcement was available constant. Of particular interest are variability data from Roberts' Experiment 1, which are reported in his Table 2, but not otherwise discussed in his article. Although the peak location was little affected by probability of reinforcement at 20 s, the spread of the fitted curve increased with decreasing reinforcement probability: that is, timing was less precise with the lower reinforcement probability. Recalculation of data from Roberts' Table 2 (p. 247), shows that coefficient of variation (standard deviation of fitted curve/peak location) was 0.37 with $p=0.8$, but 0.56 with $p=0.2$.

Some of our results also resemble those found previously by Whitaker et al. (2003), who studied 2-component mixed-FI schedules where the components were always equally likely (i.e. $p=0.5$). For one thing, peak locations were close to S , as this was

varied over conditions, but much more likely to be slightly above S than below it (see their Fig. 7, p. 287). Coefficient of variation ratios for timing of S and L were, however, generally closer to 1.0 than in the present study, indicating greater similarity of precision of timing of S and L in Whitaker et al. (2003) than in the present study.

In many ways, the results obtained in our study are consistent with a "commonsense" view of timing and responding in animals. Q: When do animals respond most? A: They respond most at the time(s) when reinforcers are usually available for responding (i.e. at S and L). Q: How much do they respond? A: They respond at rates determined by the probability of reinforcement for responses at the time(s) when reinforcement is sometimes available. The first of these results is consistent with all theories of temporal control in animals (e.g. Gibbon et al., 1984), the second with response-generating principles related to Matching (e.g. Herrnstein, 1961, 1970), which can be derived in a number of ways (see McDowell, 2004, for the most recent derivation).

However, although our results are consistent with a casual "commonsense" view of behavioural control, contemporary theories of timing experience considerable difficulties in accounting for their details and, in some cases, these difficulties seem to reveal fundamental problems with the assumptions of the theories themselves.

The most recently developed theory of timing is that of Dragoi et al. (2003). This regards timed behaviour as a property emerging from interaction of reinforced and nonreinforced responses, and does not use an internal timer. The details of this model will not be discussed here, except to note that their account experiences particular difficulties in simulating behaviour on mixed-FI schedules. Two peaks in responding on mixed-FI can be predicted when S and L are not very different (Dragoi et al., 2003, p. 141), but not when they differ substantially (as in many examples in Whitaker et al., 2003), and the model also fails to take any account of effects of reinforcement probability at S (although no details are given, Dragoi et al., 2003, p. 141), so apparently cannot in principle explain our data. Another recently developed theory, the Multiple-Time-Scales model of Staddon and Higa (1999); see also Staddon et al. (2002) does not appear to have been applied in any detail to mixed-FI schedules, and it is unclear what its predictions will be, and the same is true for another recent account of animal timing, the "packet theory" of Kirkpatrick (2002).

This leaves three competing accounts which have discussed timing on mixed-FI schedules: *Scalar Expectancy Theory* (SET: Gibbon et al., 1984), the *Behavioral Theory of Timing* (BeT: Killeen and Fetterman, 1988), the *Learning to Time* model (LeT: Machado, 1997). The first two deal only with temporal control, and do not make predictions about response output, the third deals with relative response output at S and L . We will provide a simplified introduction to each theory in turn, then discuss how it might or might not deal with the main features of our data.

SET (Gibbon et al., 1984) derives timed behaviour from a combination of clock, memory, and decision processes. An internal clock of a pacemaker-accumulator-type provides a raw measure of elapsed time (e.g. in an interval of FI), and when a reinforcer is delivered for responses at time t , t is stored in a ref-

erence memory of “important” times. The decision to respond involves a comparison of current elapsed time (as measured by the clock), and a sample of t drawn from the reference memory. When current elapsed time and the sample of t are “close enough” (according to mathematically specified rule), the animal responds (see Church et al., 1994, for details). An important property of SET is the “scalar property” of variance which gives the theory its name. This is the requirement that absolutely different durations are timed with the same relative sensitivity (i.e. the same coefficient of variation). The principal source of (scalar) variance in the model is considered to be the reference memory of times associated with reinforcement, at least in “classical” SET (Gibbon et al., 1984; see discussion in Jones and Wearden, 2003, 2004).

On mixed-FI schedules, SET predicts that responding will peak at S and L (as these times are associated with reinforcement, and thus stored in the reference memory), and this was close to the result obtained in our experiment (although the peak near S tended to be slightly higher than S). It predicts that S and L will be timed with the same coefficient of variation (see Whitaker et al., 2003), whereas in our data (but not in that of Whitaker et al.) we found a tendency for slightly more precise timing of S than L .

SET makes no predictions about how much responding will occur at times when reinforcement is available, so has no predictions to verify or falsify about the effects of reinforcement probability observed in our work. However, Roberts (1981), in an influential article which predated the formal development of SET by Gibbon et al. (1984), proposed that the decision about when to respond was determined by a clock/comparison mechanism (later more fully specified by SET), but that response rate was determined by a “third stage” (Roberts, 1981, pp. 246–247), which determined “how much” to respond, and was influenced by variables such as probability of reinforcement. As mentioned above, Roberts (1981) was only concerned with what we have called peak location, rather than measures of variability. As far as peak location is concerned, our data support Roberts’ views of the “independence” of timing and responding, as peak location was independent of peak response rate.

It is perhaps surprising, given the intensive theoretical focus of SET, that Roberts’ suggestion of a “third stage” determining response rate (perhaps using a reinforcement-rate-sensitive principle, like Herrnstein’s equation, 1970) was not formally developed. Obviously, SET needs some response-generating mechanism to deal with data like those in the present article, and even to make predictions about absolute response rate on simple FI schedules, and a recent article by Wearden and Lejeune (2006) has provided some suggestions as to what the response generating mechanism might be.

Killeen and Fetterman’s (1988) BeT dispenses with the memory and decision processes of SET and, although it uses an internal pacemaker, the function of this pacemaker is quite different from that in SET. Timed behaviour is derived from a series of (usually unobserved) “adjunctive” behaviours, each one of which serves as a discriminative cue for the next, with the measured operant being triggered by one of adjunctive behaviours. The function of the pacemaker is to shift the animal from one

adjunctive state to another, and the pacemaker is not directly “read” by the organism. The function of reinforcement at t is to organise the sequence of adjunctive behaviours, so that the animal is responding at t , and its responses are reinforced. One distinctive feature of BeT is that the rate of the pacemaker is related to the rate of reinforcement in the experimental situation, going faster if the reinforcement rate is higher, and a number of studies have attempted to test this prediction (see Lejeune et al., 1998, for one test, and for a review of others). In general, a faster pacemaker will produce more precise timing, which has some consequences for predictions of behaviour in our experiments, as will be discussed further later.

BeT, like SET, has no difficulty locating peaks of responding on mixed-FI schedules at S and L , since reinforcement at these times will organise adjunctive sequences so that the measured operant occurs at or near these times. In general, BeT would tend to predict more precise timing of L than S (see Leak and Gibbon, 1995; Whitaker et al., 2003, for discussion) whereas we found that, if anything, S was tended to be timed more precisely than L .

Perhaps surprisingly for a “behavioural” theory, BeT makes no specific predictions about how response rate should vary as a function of reinforcement rate or probability, although principles which are similar to those used in BeT, but which predated it (e.g. Killeen, 1979) did make some statement about absolute response rate. In Killeen (1979), the effect of reinforcement was to induce “arousal”, with higher reinforcement rates producing higher arousal, and this arousal was translated into response rate (Killeen, 1979, pp. 40–41). An arousal level proportional to reinforcement rate could be used to derive matching-like relations between reinforcement and response rates (Killeen, 1979, pp. 59–60), so potentially dealing with some of the effects in our data. However, the prediction of absolute response rates never became a full formal development of BeT. It seems that both BeT and SET need further theoretical development before absolute response rates can be predicted, but both have potential “add-on” mechanisms that might be developed to do this (e.g. Roberts’ “third stage” or Killeen “arousal” hypothesis).

LeT (Machado, 1997) is sometimes considered to be a development of BeT (although some theorists disagree, see Hopson, 2003). In LeT, the delivery of a “time marker” (such as reinforcement on an FI or mixed-FI schedule) starts a sequence of “states”, each of which rises and falls in activation level as time elapses, but these states are less clearly identified with overt behaviours than in BeT. At any particular time, a particular state will have a certain level of activation, and a certain “coupling” to the measured operant. This coupling between the measured operant and a state occurs when responding is reinforced in the presence of this state. The coupling between a state and the measured operant is governed by a conditioning-like process, so increases when responses are reinforced when the state is active, and decreases when responses occur when the state is active but are not reinforced. The generation of a response at some time thus depends on a state which is coupled to the response being active at that time.

LeT predicts responding on mixed-FI schedules at around the time of S and L , as during learning the states active at around S

and L are coupled to the measured operant by the occurrence of reinforcement for responses at these times. LeT further predicts that peak response rate at S will be affected by reinforcement probability, as states active at S have their coupling with the measured operant reduced by nonreinforcement, so lower probabilities of reinforcement produce greater nonreinforcement, less coupling, and thus lower response rates. Machado (1997) in fact modelled the data from Catania and Reynolds (1968), discussed above, and simulated their data on effects of reinforcement probability at S (see his Figs. 11–13, pp. 251 and 252).

However, LeT makes different predictions for responding at L . States active at L on mixed-FI schedules are never associated with nonreinforcement: if the interval ends with reinforcement at S , then these states have not yet been activated so can show no change in coupling, whereas if the interval ends with reinforcement at L , the states active at L are consistently reinforced. In fact, Machado's LeT deals particularly well with the data from our experiment, as response rate at S is systematically affected by reinforcement probability, whereas that at L is much less sensitive (something also noted in data from Catania and Reynolds, 1968), as LeT predicts. However, a potential problem for LeT is that it predicts that the response rate at L is *always* higher than at S (something which is contradicted by data, although higher rates at L than S are sometimes observed, see also Whitaker et al., 2003). Thus, LeT predicts relative response output at S and L well in some cases, but seems to lack flexibility, as responding at L is predicted to be (a) always higher than at S and (b) unaffected by reinforcement probability. While these cases do occur, many contrary data exist (e.g. in Whitaker et al., 2003, and the present study).

As Machado (1997) himself notes, the properties of responding at L come from fundamental assumptions concerning the coupling of operant responses to activated states by reinforcement and extinction experiences with those responses, and alternative possibilities cannot be accommodated within the model without changing its basic structure and assumptions radically.

How do BeT and LeT deal with differences in precision of timing, such as the coefficients of variation of the Gaussian curves corresponding to behavioural control by S and L ? Both BeT relates pacemaker rate to reinforcement rate, so increasing reinforcement rate makes the pacemaker run faster, with attendant increases in the precision of timing. LeT is less specific about this issue, although it allows (but does not require) links between λ , the rate of transition between states, and the interval value. This parameter effectively governs the sensitivity of timing of the system, with more sensitive timing occurring with higher λ values. To produce scalar timing (i.e. constant coefficient of variation) BeT requires the pacemaker rate to vary proportionally with reinforcement rate in the experimental situation, and LeT requires λT to be constant, where T is the time to reinforcement from the time marker. So, on simple FI schedules, to achieve scalar timing, the rate of the pacemaker of BeT would have to vary inversely with the FI value, and λ would also have to vary inversely with the FI value (Machado, 1997, p. 251).

In our experiments, varying the probability of reinforcement for responses at S changes overall reinforcement rate very

markedly, so theories which appear to predict changes in sensitivity of timing with changes in reinforcement rate, like BeT does, would predict more sensitive timing with increases in probability of reinforcement at S , but these were not found in our data (e.g. Fig. 2), where coefficients of variation remained roughly constant as reinforcement probability at S changed.

Finally, we briefly discuss a recent account of response rates on FI and related schedules, including mixed-FI schedules, proposed by Wearden and Lejeune (2006). This model is unusual in that it deals exclusively with response output on these schedules, leaving issues of timing to other mechanisms (such as SET), so constitutes something close to Roberts' (1981) "third stage". The basic idea is that response rate on some FI-like schedule is determined by the average delay of reinforcement for responding during the response periods occurring, with delay of reinforcement being related to response rate by a hyperbolic function. Longer FI values engender longer periods of responding before reinforcer delivery, thus on average responses in these periods are reinforced with longer delays at longer FIs than shorter ones, leading to decreasing response rates with increasing FI values, as found in data. The predictions for 2-valued mixed FI schedules are more complicated. Responding before S is reinforced both "immediately" on intervals when FI S is in force, but only after a delay when reinforcement is arranged according to FI L , and the two sources of reinforcement both contribute to response rate at S . Responses in the response period preceding L , on the other hand, are always reinforced "immediately", on intervals when reinforcers are delivered according to FI L .

In the present case, Wearden and Lejeune's model predicts that response rates will change with reinforcement probability at S , but not with reinforcement probability at L . The reason is that as the probability of reinforcement at S decreases, responses in the response period before S are less likely to be reinforced "immediately" and thus more likely to be reinforced after a delay when reinforcers are delivered according to FI L . On the other hand, responses in the response period preceding L are always "immediately" reinforced, so are unaffected by reinforcement probability at L .

Wearden and Lejeune's model can make exact quantitative predictions but, unfortunately, we cannot fit it to data from the present experiment as the durations of the response periods before S and before L were not measured when the experiment was conducted more than 25 years ago, and cannot now be derived from the data which exist. However, the delay of reinforcement model proposed by Wearden and Lejeune appears in good qualitative agreement with the data we present here.

Overall, therefore, the data in our study suggest that timing accuracy (e.g. peak location) and timing sensitivity (e.g. coefficient of variation) are not affected by reinforcement probability at some critical time, nor by the overall rate of reinforcement in the experimental situation. These factors appear to be determined solely by the times of reinforcement. On the other hand, response output at S , but not L , is strongly determined by reinforcement probability in a manner consistent with the delay of reinforcement model of Wearden and Lejeune (2006), itself a variant of models of response output on aperiodic schedules.

References

- Blomeley, F.J., Lowe, C.F., Wearden, J.H., 2004. Reinforcer concentration effects on a fixed-interval schedule. *Behav. Process.* 67, 55–66.
- Catania, A.C., Reynolds, G.S., 1968. A quantitative analysis of responding maintained by interval schedules of reinforcement. *J. Exp. Anal. Behav.* 11, 327–383.
- Church, R.M., Meck, W.H., Gibbon, J., 1994. The application of scalar timing theory to individual trials. *J. Exp. Psychol.: Anim. Behav. Process.* 20, 135–155.
- Dragoi, V., Staddon, J.E.R., Palmer, R.G., Buhusi, C.V., 2003. Interval timing as an emergent learning property. *Psychol. Rev.* 110, 126–144.
- 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 Timing and Time Perception*, vol. 423. New York Academy of Sciences, New York, pp. 52–77.
- Herrnstein, R.J., 1961. Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exp. Anal. Behav.* 4, 267–272.
- Herrnstein, R.J., 1970. On the law of effect. *J. Exp. Anal. Behav.* 13, 243–266.
- Hopson, J.W., 2003. General learning models: timing without a clock. In: Meck, W.H. (Ed.), *Functional and Neural Mechanisms of Interval Timing*. CRC Press, Boca Raton, FL, pp. 24–60.
- Jones, L.A., Wearden, J.H., 2003. More is not necessarily better: examining the nature of the temporal reference memory component in timing. *Q. J. Exp. Psychol.* 56, 321–343.
- Jones, L.A., Wearden, J.H., 2004. Double standards: memory loading in temporal reference memory. *Q. J. Exp. Psychol.* 57B, 55–77.
- Killeen, P.R., 1979. Arousal: Its genesis, modulation, and extinction. In: Zeiler, M., Harzem, P. (Eds.), *Reinforcement and the organisation of behaviour*. Wiley, Chichester, pp. 31–78.
- Killeen, P.R., Fetterman, G.J., 1988. A behavioral theory of timing. *Psychol. Rev.* 95, 274–295.
- Kirkpatrick, K., 2002. Packet theory of conditioning and timing. *Behav. Proc.* 57, 89–106.
- 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., 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.
- Lejeune, H., Wearden, J.H., 1991. The comparative psychology of fixed-interval responding: some quantitative analyses. *Learn. Motiv.* 22, 84–111.
- Machado, A., 1997. Learning the temporal dynamics of behavior. *Psychol. Rev.* 104, 241–265.
- McDowell, J.J., 2004. A computational model of selection by consequences. *J. Exp. Anal. Behav.* 81, 297–317.
- Roberts, S., 1981. Isolation of an internal clock. *J. Exp. Psychol.: Anim. Behav. Process.* 7, 242–268.
- 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., Innis, N.K., 1969. Reinforcement omission on fixed-interval schedules. *J. Exp. Anal. Behav.* 12, 689–700.
- 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., Lejeune, H., 2006. The stone which the builders rejected”: delay of reinforcement and response rate on fixed-interval and related schedules. *Behav. Process.* 71, 77–87.
- Whitaker, J.S., 1980. Temporal control of animal operant performance. Unpublished Ph.D. Thesis. University College of North Wales, Bangor.
- 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.