Available online at www.sciencedirect.com
SCIENCE DIRECT
Behavioural Processes 71 (2006) 77–87
BEHAVIOURAL PROCESSES
www.elsevier.com/locate/behavproc

Review

"The stone which the builders rejected. . .": Delay of reinforcement and response rate on fixed-interval and related schedules

J.H. Wearden[a,*], Helga Lejeune[b]

[a] University of Keele, School of Psychology, Keele, Staffordshire ST5 5BG, UK
[b] University of Liège, Faculty of Psychology, Boulevard du Rectorat, Sart-Tilman, 4000 Liège, Belgium

Received 6 May 2005; received in revised form 16 July 2005; accepted 17 August 2005

Abstract

The article deals with response rates (mainly running and peak or terminal rates) on simple and on some mixed-FI schedules and explores the idea that these rates are determined by the average delay of reinforcement for responses occurring during the response periods that the schedules generate. The effects of reinforcement delay are assumed to be mediated by a hyperbolic delay of reinforcement gradient. The account predicts that (a) running rates on simple FI schedules should increase with increasing rate of reinforcement, in a manner close to that required by Herrnstein's equation, (b) improving temporal control during acquisition should be associated with increasing running rates, (c) two-valued mixed-FI schedules with equiprobable components should produce complex results, with peak rates sometimes being higher on the longer component schedule, and (d) that effects of reinforcement probability on mixed-FI should affect the response rate at the time of the shorter component only. All these predictions were confirmed by data, although effects in some experiments remain outside the scope of the model. In general, delay of reinforcement as a determinant of response rate on FI and related schedules (rather than temporal control on such schedules) seems a useful starting point for a more thorough analysis of some neglected questions about performance on FI and related schedules.

Keywords: Fixed-interval schedules; Mixed-fixed-interval schedules; Running rates; Peak rates; Delay of Reinforcement; Temporal Control

Contents

Like an astronomer observing the birth of stars in condensing clouds of gas, the contemporary reader of Skinner's *The Behavior of Organisms* (Skinner, 1938) can watch the field of reinforcement schedule research blaze into life before their very eyes. This is not, however, to say that Skinner's early concerns exactly match those that were to be important later. We discuss the treatment of time-based reinforcement schedules in *The Behavior of Organisms* elsewhere in more detail (Lejeune et al., in press), but a unusual feature of Skinner's early treatment of one of the best-known of his inventions, the fixed-interval (FI) schedule, is that his principal interest was not in the temporal patterning of responding that such a schedule produced, but in the response rates it generated. On an FI schedule, the first response occurring t seconds or more from the previous food delivery is itself reinforced, and this food delivery restarts the time requirement of the schedule, t.

---

[*] Corresponding author.
E-mail address: j.h.wearden@psy.keele.ac.uk (J.H. Wearden).

0376-6357/$ – see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.beproc.2005.08.006

With respect to FI schedules, Skinner's main concern in his early work was to explain learned behaviour in terms of two constructs, "reflex strength" and the "reflex reserve", the latter being defined as "a certain amount of *available activity*, which is exhausted during the process of repeated elicitation and of which the *strength* of the reflex is at any moment a function" (Skinner, 1938, p. 26), response rate being "the principal measure of the strength of the operant" (p. 58 — more recent elaborations of the reflex reserve can be found in Killeen, 1988; Nevin, 1992; Nevin and Grace, 2000, for example). In *The Behavior of Organisms*, what we now know as the FI schedule is described as "periodic reconditioning", a term clearly linked to the idea of reflex elicitation and response depletion by non-reinforcement.

In recent years, FI schedules, and variants such as the peak procedure (Church et al., 1984; Roberts, 1981) have been extensively studied and analysed with respect to the temporal control of behaviour that they produce; that is, the patterning of responses in time, most strikingly exemplified by the within-interval pattern of responding on FI, the "scalloped" or "break and run" pattern found in innumerable studies since 1938 (see Ferster and Skinner, 1957, for literally hundreds of examples).

Temporal control of behaviour on FI and other schedules with temporal reinforcement periodicities or temporal constraints has been extensively investigated and a number of competing models of the timing of responses on FI and other schedules exist (e.g. Dragoi et al., 2003; Gibbon et al., 1984; Killeen and Fetterman, 1988; Machado, 1997; Staddon et al., 2002). In the present article, we will largely leave issues of temporal control (i.e. "when to respond?") aside, and instead concentrate on Skinner's original interest, the rate of responding on FI (i.e. "how much to respond?"). Questions about response rate on FI and related schedules, if not completely neglected, have attracted only a tiny fraction of the experimental and theoretical analysis that temporal control has enjoyed, with the result that many fundamental issues remain unresolved.

Most of the present article is an exploration of the idea that considerable progress in understanding response rate on FI-like schedules can be made using the idea of delay of reinforcement for responses, although it should be acknowledged at the outset that this idea does not account for all existing data and many puzzles remain. Some early attempts to account for response patterning on FI solely in terms of delay of reinforcement (e.g. Dews, 1970) have been severely criticized. The basic argument was that a delay of reinforcement process could bypass the need for a timing mechanism (e.g. an internal clock) as an account of FI behaviour. After considerable training on FI animals pause or respond at low rates immediately after reinforcer delivery, then response rates increase as time elapses in the interval to a peak just before reinforcer delivery (the "scallop", for example, Dews, 1978) or, in other cases, responding emerges abruptly at some high rate after a pause then continues to the end of the interval ("break and run", Schneider, 1969). A delay of reinforcement explanation of such behaviour would note, correctly, that responses occurring early in the interval are reinforced only after a long delay, whereas those occurring later are reinforced much more immediately. This is, of course, true, but presupposes that the organism can distinguish between responses early and later in the interval, i.e. it presupposes a mechanism for temporal control rather than delay of reinforcement being a replacement for one. However, as we hope to convince the reader in this article, delay of reinforcement makes much more sense as an explanation of *response rate* on FI, given that some mechanism for temporal control (which we will not discuss in any detail here) already exists.

1. Response rates on FI

"Response rate" on FI is not a unitary concept, as three types of rates can be distinguished. Firstly, there is "overall response rate", which is the total number of responses on a given interval or session divided by the duration (usually in minutes) of that interval or session. The second type, the "running rate", is computed by dividing the number of responses by the duration of the interval or session minus the duration of the post-reinforcement pause(s). It thus describes response rates after the post-reinforcement pause has ended. The third type, "terminal rate", is usually computed taking into account only responses emitted during some late segment of the interval (for example, the last tenth), just before reinforcer delivery. In a well known variant of the FI schedule, the peak procedure, "peak rate" (the rate at the peak of the response rate function) can be considered as very similar to the terminal rate on FI. In the peak procedure (Roberts, 1981), most trials of an experimental session proceed as normal discrete-trial FI trials, with trials separated by an inter-trial interval, and the start of the interval is signalled by onset of a stimulus. However, some trials, the "peak trials" from which data are taken, last much longer than the normal FI trials, and terminate without reinforcer delivery, so response rate on these trials can be measured at times both above and below the FI value in force. The response rate versus time function from peak trials can be analysed to determine the location and spread of the peak of responding, as well as the peak height.

The core of this paper deals with mainly running rates and terminal or peak response rates. However, for simplicity, the term response rate, without additional qualification, will often be used.

Skinner (1938, Figure 29, p. 120) presented the number of responses cumulated over 24 one-hour sessions on FI schedules of 3, 6, 9, and 12 min, where a single rat was used at each FI value, and his results, converted into response rates, are drawn in the upper part of Fig. 1. Obviously, overall response rates declined with increasing FI value, or with decreasing reinforcement rate. The lower two panels of Fig. 1 show data from Spencer (1978: see also Spencer, 1981), which are strikingly similar to Skinner's original results. The data are averages over four rats of running rates on FI values ranging from FI 15 to 480 s (see Lowe et al., 1979, for a report of more data from this experiment). The centre panel shows an obvious decrease in running rate as the FI values became longer, and the lowest panel shows the same data plotted against the reinforcement rate the FI schedules arranged. The shape of the increase is the familiar rectangular hyperbola

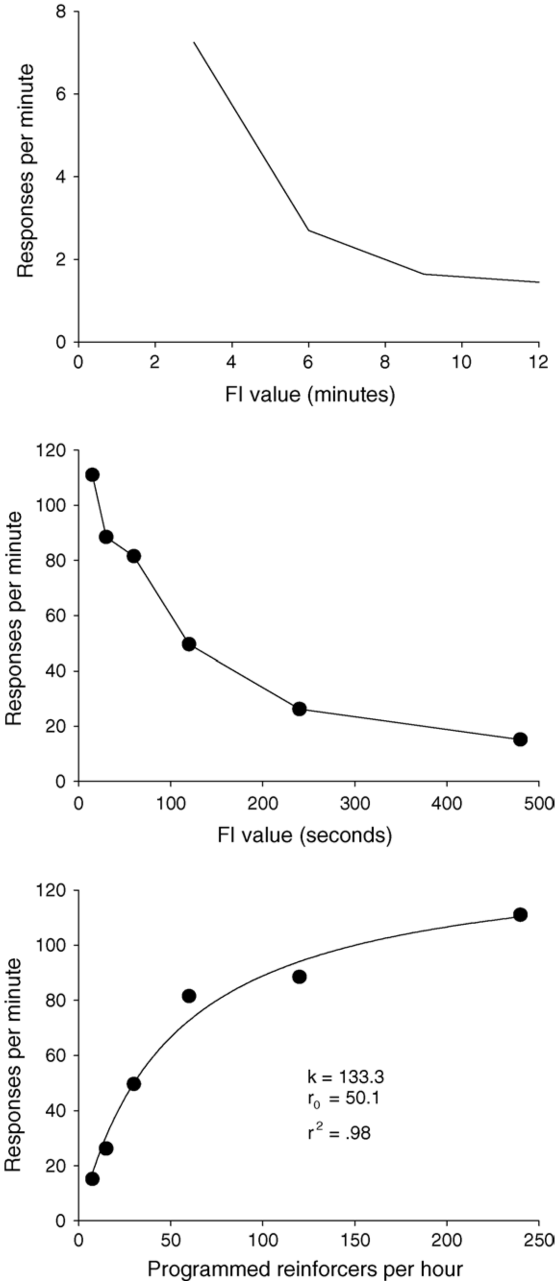

Fig. 1. Upper panel: response rates from four rats (tested separately on FI values of 3, 6, 9, and 12 min), plotted against FI value. Centre panel: Averaged data from four rats tested on FI values ranging from 15 to 480 s. Running rates (response rate after the termination of the post-reinforcement pause) are plotted against FI value. Lower panel: the data from the centre panel rearranged to show response rate versus programmed reinforcement rate. The function fitted to data is Herrnstein's equation (Eq. (1) in the text).

of Herrnstein's equation (Herrnstein, 1970) which is

$$R = \frac{kr}{(r + r_0)} \tag{1}$$

where R is response rate and r is reinforcement rate, with k and r_0 being constants. The usual interpretation of the parameters is that k reflects the maximum response rate of which the animal is capable in that experimental situation, with r_0 representing the rate of "extraneous" reinforcers, that is, all sources of reinforcement present in the experimental situation other than that provided for the operant response. This equation is shown fitted to the data, with the k, r_0, and r^2 (proportion of variance accounted for) parameters shown in the bottom panel. To provide the fit we used programmed rates of reinforcement, as the delivered rates of reinforcement the schedules provided were not available in numerical form in Spencer (1978). Spencer (1981) also showed that Herrnstein's equation provided a good fit to data from individual animals, and he also included examples of a hyperbolic relation between running rates and reinforcement rates on FI obtained from pigeons.

2. Acquisition of FI performance and "confinement"

Rate of response on FI is not the only neglected issue concerning this schedule: another is the acquisition of performance, although this has attracted some interest both in "classical" work (Ferster and Skinner, 1957; Skinner, 1938), and more recently in work by Machado and colleagues (Machado, 1997; Machado and Cevic, 1998). Acquisition of FI performance raises a number of questions that need to be addressed, and Figs. 2 and 3 show some of them. Fig. 2 shows data from four rats placed directly on FI 60 s after initial lever press training, with results coming from an unpublished study conducted at the University

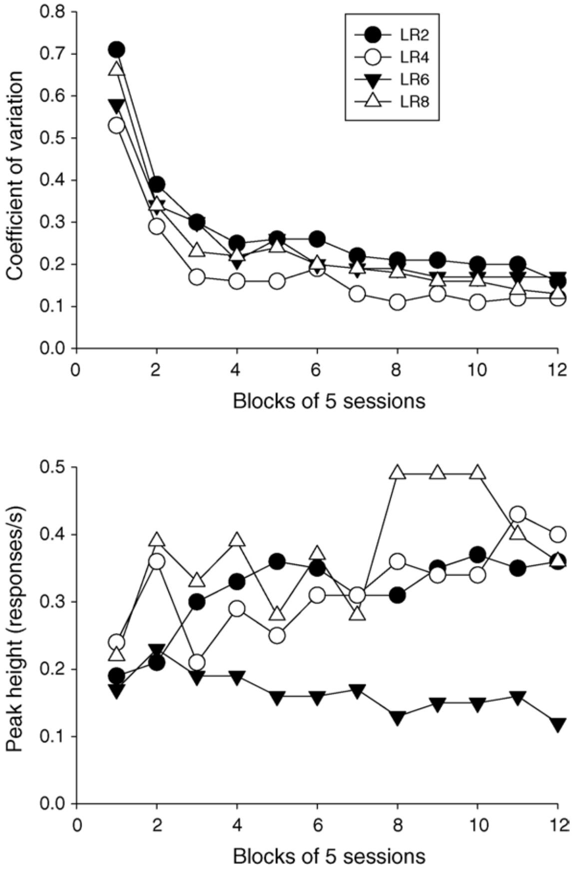

Fig. 2. Data from 12 five-session blocks of acquisition of FI 60 s. Upper panel: coefficient of variation of fitted Gaussian curve (see text for details). Lower panel: peak height of fitted curve (peak response rate). Both measures are plotted against session blocks.

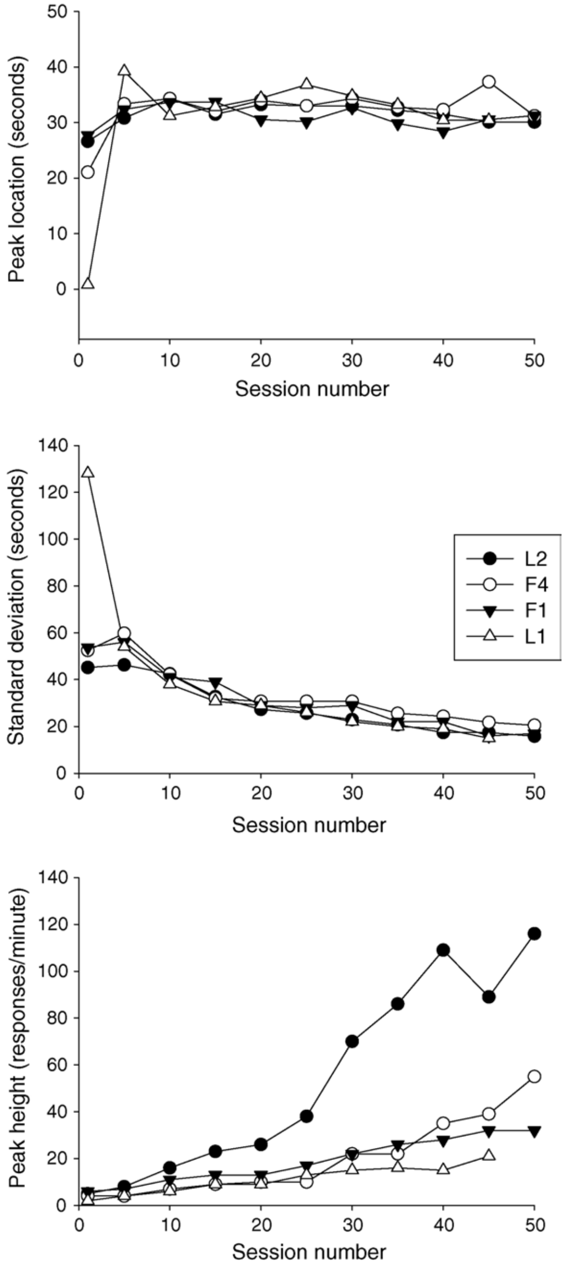

Fig. 3. Response measures derived from Blomeley (1984) over 50 sessions (45 sessions for rat L1) of acquisition of responding on mixed-FI 30 s FI 150 s. Upper panel: Location of peak of Gaussian curve fitted to response rate versus elapsed time function. Centre panel: standard deviation of fitted curve. Lower panel: peak height (peak response rate) of fitted curve.

of Liège. The response rate versus elapsed time in the interval functions were fitted by the left-hand-side of a Gaussian curve (see Lejeune and Wearden, 1991, for examples), with the peak of the curve being forced at the FI value. Data were collected from 60 consecutive sessions, and Fig. 2 shows the coefficient of variation of the fitted curve (the standard deviation of the curve divided by the mean: essentially a measure of the distribution of responding throughout the interval) in the upper panel, and the peak height of the fitted curve (effectively the terminal rate just before the time of reinforcement). Two features are obvious: firstly, coefficients of variation declined with increasing training (initially very rapidly, then more slowly) and secondly, this decrease in the coefficient of variation (which in effect shows that the period of responding is becoming shorter with increasing training) is accompanied by an increase in response rate (at least for three of the four rats shown). We will call this effect "confinement": the shorter the period during which responding occurs on FI, the higher the rate of response at which responding proceeds. So, for example, early in training the animals exhibit high coefficients of variation, indicating that responding is dispersed throughout the interval, and this is associated with low terminal response rates, whereas later the coefficient of variation is smaller, and the responding accordingly more concentrated at the end of the interval, with a higher terminal rate.

The data in Fig. 2 trace only part of the evolution of temporal control and response rate on FI, as the peak of the curve was forced by the fitting procedure to be at the FI value, but data in Fig. 3 permit us to observe a more complete picture. The data come from an unpublished study in Blomeley (1984), which used a method described in more detail in Blomeley et al. (2004). The schedule used was mixed-FI 30 s FI 150 s, where the two components were equally likely, so the first response occurring either 30 or 150 s after the previous reinforcer delivery was equally likely to be reinforced. At the start of the interval, nothing signalled to the animal whether the reinforcer would be scheduled according to FI 30 or 150 s. One difference from "normal" mixed-FI (e.g. Whitaker et al., 2003) was that the different schedules were associated with different levers, so reinforcement according to FI 30 s was scheduled on one lever, reinforcement according to FI 150 s on the other. Data were measured from the "FI 30 s" lever, but in intervals where the response was reinforced according to FI 150 s, so both the ascending and descending limbs of the response rate versus time function could be observed; see Blomeley et al. (2004) for some examples of response functions from an experiment with a very similar procedure.

Gaussian curves were fitted to the response rate versus elapsed time in the interval function, and now we can observe three indices of performance: the elapsed time when responding peaks (peak location), the spread of the response rate function (standard deviation of the curve), and the rate of responding at the peak location (the peak rate). Fig. 3 shows these indices for four rats: peak location (upper panel), standard deviation (centre panel), and peak height (effectively peak rate, lower panel).

It is clear from Fig. 3 that (a) the peak location very quickly centred on a value close to the usual time of reinforcement on the lever (30 s), (b) the period of responding gradually reduced over the whole 50 sessions, initially showing a rapid decrease, then a more gradual reduction, and (c) the response rate at the peak increased steadily over the whole 50 sessions, increasing 20-fold or more in some animals. The correlations between standard deviation and peak height over the sessions of acquisition shown were strongly negative for all animals (from −0.75 to −0.93 for individuals). Thus, the data from this experiment also show "confinement": the shorter the period in which responding is confined (indicated by the value of the standard deviation), the higher the rate at which it proceeds.

The examples above suggest that the animals learn the time of reinforcement fairly quickly, within a few sessions, in effect a few hundreds of reinforcer deliveries, but that responding becomes increasingly confined to elapsed times close to the time of reinforcement (and this progressive change takes much longer), with response rate in this increasingly confined period increasing, sometimes very markedly, as experience with the schedule accumulates. As mentioned above, the concern of the present article is more or less exclusively with issues concerning response rate rather than temporal control on FI, so for the remainder of the work we will try to develop some principled account of responding on FI.

3. The stone which the builders rejected: delay of reinforcement

The principle that we shall use for the remainder of the article is very simple: we assume that response rate on FI is governed by the average delay of reinforcement for responses during a response period, although what constitutes a "response period" is sometimes not completely straightforward, and such responses can have more than one source of reinforcement, as we will see later. Given that we use the idea of response periods on FI, our account is essentially concerned with "running rates", that is, responding after the pause has terminated. Running rates are always higher than overall response rates, but the relationship between both rate measures is more complex than might be expected. Whereas high overall rates are always correlated to high running rates, the reverse is not true. Indeed, low overall rates might be associated with very high running rates if a high-rate bout of responding occurs after a long post-reinforcement pause. As a consequence, overall or running rates are only moderately correlated with pause indices, such as the time to the first response or the quarter life (time taken to produce the quarter of the total number of responses; see Gollub, 1964; Elsmore, 1971). Similar conclusions were reached by Dukich and Lee (1973) using a mixed-FI schedule. The analysis of the relationship between pause indices, overall and running rates on FI would require interval-by-interval measures, as it has long been known that the number of responses per interval as well as the duration of the post-reinforcement pause can vary considerably in absolute value from one interval to the next (for example, Dews, 1970, Figs. 2–7, p. 55; Mechner et al., 1963; Shull, 1971, Fig. 1, p. 228). However, rat data from Gollub (1964) suggested that the moderate correlations between overall rates and the quarter life hold both for session-by-session (Table 3, p. 342) and individual interval analyses (Table 2, p. 339).

To model delay of reinforcement effects we will use a simple version of Mazur's hyperbolic delay equation (Mazur, 1987):

$$R = \frac{A}{(1 + cP)} \tag{2}$$

Here, R is response "strength" (which is assumed to be transformed in some very simple way into response rate), P, the average delay to reinforcement during some response period, and A and c are constants. A is usually interpreted as the maximum response rate possible when delay is zero, and c is a constant reflecting the effects of delay. When c is very small, delays of reinforcement have little effect on responding, whereas when c is larger, increasing delay of reinforcement decreases response rate markedly.

We will use Eq. (2) to provide at least qualitative predictions of changes in running rate with changes in FI value and some other variables, but some words of caution are in order. Firstly, Eq. (2) is only the simplest version of hyperbolic delay equation, and more complex equations can be used. Secondly, we will generally keep calculations very simple to illustrate the mechanics of the model, so we will not usually attempt to fit the actual running rates obtained in experiments, just the pattern of results obtained when experimental parameters are changed. Thirdly, we must admit at the outset that although Eq. (2) performs surprisingly well in some novel situations, there are cases in the literature which seem to require a more elaborate, or perhaps completely different, treatment. However, since the theoretical analysis of response rates on FI has been seriously neglected, a simple model which makes some initial progress seems a reasonable starting point.

One thing that is obvious from Eq. (2) is that the phenomenon of "confinement" noted during acquisition (e.g. Figs. 2 and 3) is immediately predicted. Other things being equal (e.g. constant A and c values) the shorter the average response period (P) the higher the running rate produced. This does not depend on the actual values of A and c chosen, providing these remain constant across conditions which are compared, unlike some other effects to be discussed later. Figs. 2 and 3 showed examples of progressively shorter response periods occurring as animals acquired temporal control on FI schedules, where the schedule value remained constant. Another way of changing the duration of the response period on FI is, obviously, to change the absolute value of the schedule itself. Now, longer FI values will generate longer response periods even in steady-state conditions (e.g. Lowe et al., 1979), so the principle of confinement should predict changes in response rate as schedule values are varied on FI.

How does Eq. (2) perform as an account of running rate on simple FI schedules, effects like the results shown in the panels of Fig. 1? To determine this, we need first to calculate the duration of the response period on FI (the "run" after the post-reinforcement pause has terminated), and there are a number of possibilities. One suggestion for simple FI schedules might be to assume that the response period was determined by the FI value minus the post-reinforcement pause, with the mean pause being a power function of the FI value (pause = at^b), where a and b are constants and t is the FI value (as found by Lowe et al., 1979). Typical values might be $a = 1.5$ and $b = 0.8$, which produces response periods which are increasing proportions of the FI value as this increases, the usual result found in data (Lowe et al., 1979), and we will initially use these for our calculations.

Fig. 4 shows response "strength" on FI (expressed here as the percent of maximum running rate) determined by Eq. (2). A was kept constant at 100, and c was varied over values of 0.25, 0.125, and 0.0625. Response strengths are shown for FI values of from 15 to 480 s, converted into reinforcers per hour. For sim-

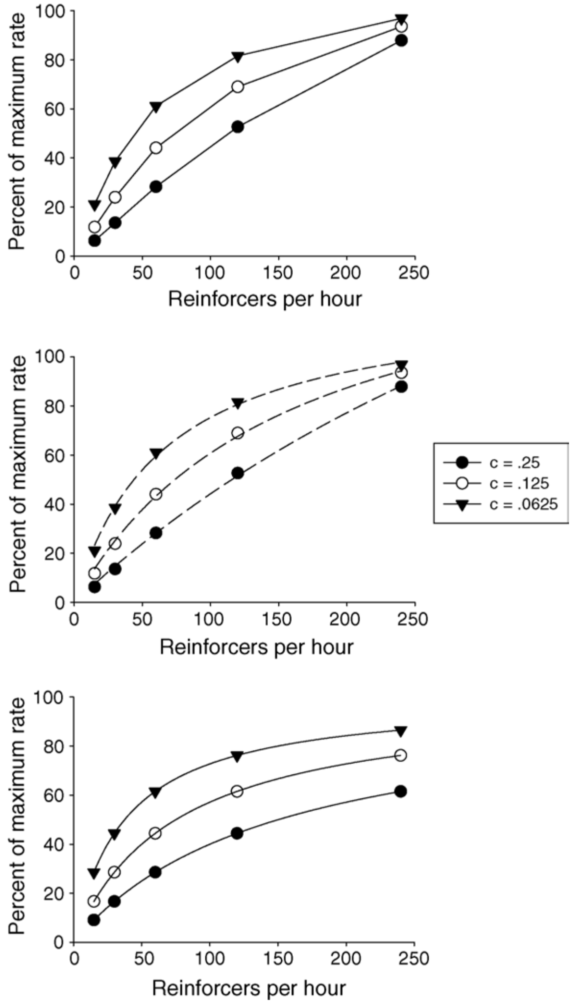

Fig. 4. Upper panel: simulated response strengths based on Eq. (2), plotted against the rate of reinforcement delivered by the FI schedule, assuming power function post-reinforcement pause. Centre panel: the same results as the upper panel, but with Herrnstein's equation fitted. Bottom panel: simulated response strengths based on Eq. (2) (symbols), assuming post-reinforcement pause proportional to interval length. Lines show fits of Herrnstein's equation.

plicity we assume that the programmed rates of reinforcement accurately reflect the obtained rates. The upper panel shows calculated response strengths assuming a power function pause. Obviously, response strengths increased with rate of reinforcement, and were generally higher when the effects of delay of reinforcement were smaller. The response strength versus reinforcement rate functions had a curved, apparently hyperbolic, shape, and the centre panel of Fig. 4 shows Herrnstein's equation (Eq. (1): dashed lines) fitted to the simulated results. It is clear on inspection of the figure that the response strengths calculated using a power function for post-reinforcement pause and Eq. (2) to determine output produced results which obeyed Herrnstein's equation, and in fact the smallest r^2 value for the fits in the lower panel of Fig. 4 was 0.997. The bottom panel shows simulated strengths (symbols), and fits of Herrnstein's equation (lines), derived from a simulation where it is assumed that the pause is always two-third of the interval duration. Once again, response strength plotted against reinforcement rate had a hyperbolic shape, and was well fitted by Herrnstein's equation ($r^2 > 0.999$).

The idea that average delay to reinforcement on FI is an important determinant of response strength on this schedule, and that the effects of delay are mediated by a hyperbolic delay function like Eq. (2) thus gives a good *prima facie* account of running rate on simple FI schedules. Not only should longer FI values (which generate longer response periods) give rise to lower running rates, but the form of the running rate versus reinforcement rate function is predicted to be indistinguishable from that required by Herrnstein's equation, which data suggest that it is (e.g. Spencer, 1981). Note that the current treatment always predicts that increasing the FI value should decrease running rate, providing that A and c are kept constant. It is not necessary to assume that the response period is an increasing proportion of the FI as the FI value increases (i.e. assuming a power function pause); simple proportionality between the response period and FI value will also generate increasing running rates when FI duration decreases (bottom panel of Fig. 4). As long as the response period increases in some ordinal way as the FI value increases, running rates will decrease as the FI value grows. Eq. (2) thus always generates running rates on simple FI schedules which obey the principle that higher rates of reinforcement generate higher rates of response, as is obvious from its structure.

4. Response rates on mixed-FI schedules

On a mixed-FI schedule with two equiprobable components, responses are reinforced according to either FI S or FI L (where $S < L$), with only one schedule being in force in any interval. At the start of the interval, nothing signals to the animal which schedule is operative. Whitaker et al. (2003) reported an extensive study of the performance of rats on mixed-FI schedules with two equiprobable components and, from the point of view of temporal control, many of their results followed commonsense. For example, consider mixed-FI 30 s FI 240 s, with data coming from the 240 s interval. What might be expected? The animals have learned that reinforcers are either delivered for responses at around 30 or 240 s, but not at other elapsed times in the interval, so a clear prediction is that responding increases from zero or low rates early in the interval to a peak around 30 s, which might be followed by a decline, then a further rise in responding towards a second peak at 240 s. This is exactly the pattern of results obtained (Whitaker et al., 2003, Fig. 1) and, in general, when S and L are not "too close together" (i.e. an L:S ratio of 3:1 or 4:1) two peaks in responding are observed in data, one close to S, the other at L. Whitaker et al.'s article contained data from many other conditions, as well as extensive discussion of the implications of the data for theories of temporal control, but we will not describe these further here.

The data from mixed-FI schedules were usually well-fitted by the sum of two Gaussian curves, one with a peak close to S, the other peaking at L, but how much did the animals respond at S and L? Now, a commonsense answer to this question fits the data much less well. The obvious expectation would be that peak rates (e.g. peak height of fitted curves) would be higher for S than for L, as the S schedule is associated with a higher rate of reinforcement than L, but this was not always found, even when S and L were very different. In Whitaker et al.'s Fig. 1, for example, data are shown from mixed-FI 30 s FI 240 s, and the peak heights of the Gaussian curves (i.e. peak rates at 30 and 240 s) were very similar when data were averaged over four rats. Data from individual rats (shown in Whitaker et al.'s Fig. 2) presented a more complex picture. One rat showed a response peak at around 30 s that was much higher than at 240 s, two rats had peaks that were about equal in height, and a final rat had a higher peak rate at 240 s than 30 s. This latter result, in particular, seems to violate the principle that schedules associated with higher rates of reinforcement produce higher rates of response.

In general, in much of the data in Whitaker et al. (2003) the peak rates at S and L were often much more similar than the values of S and L might suggest they would be. On simple FI schedules, higher running rates at L than S are mathematically forbidden by our Eq. (2), providing that the response period in L is longer than that in S, which we will assume it always is, and there is little evidence that such "reversals" occur except by chance (e.g. the present Fig. 1). Why do mixed-FI schedules provide instances of similarity of peak rates at S and L, or even "reversals" of the expected pattern of rates?

To illustrate a possible explanation using Eq. (2), we will simplify the calculations by assuming a simple linear relation between response period and FI value, namely that the animal always responds for one-third of the FI value. Consider the case of mixed-FI 30 s FI 240 s. The average response periods for these two schedules are 5 and 40 s (both 0.5 of 1/3 of the FI value), but an additional consideration is that the response period associated with the FI 30 s schedule is reinforced "immediately" (i.e. according to FI 30 s) only on half the intervals, while being reinforced after a much longer delay on the other half (i.e. after $L - S$, $240 - 30 = 210$ s plus half the response period, 5 s). The response period associated with FI L is, on the other hand, always "immediately" reinforced (i.e. according to FI 240 s) if it occurs. If it does not occur (i.e. the interval is one in which the reinforcer is arranged according to FI 30 s), then the "strength" of the later response period is unaffected because the interval ends after the FI 30 s reinforcer has been dispensed. This effect tends by itself to "bias" responding towards the longer FI, but there is also an effect of the delay sensitivity parameter, c in Eq. (2).

To calculate the response strengths at FI 30 and 240 s, we used a modification of Eq. (2). The "immediate reinforcement" effect at FI 30 s assumes an average delay of 5 s. The "delayed reinforcement" of the response period around FI 30 s is assumed to occur after 215 s, (i.e. delayed by $L - S$ plus 5 s). These two sources of strength add together, but the total must be divided by 2, to take account of the fact that responding according to FI 30 s is "immediately reinforced" on half the intervals, and reinforced after a delay on the other half. The response strength at 240 s was calculated using a 40 s average delay (i.e. half the presumed response period), but responding at this time was always "immediately reinforced". To conduct the calculations, A in Eq. (2) was kept at 100, and c was varied between values of 0.25 and 0.0156. Note that smaller values imply less sensitivity to the effects of delay. Fig. 5 shows the ratio of the "response strengths" at 30 and 240 s as a function of c: here values greater than 1 mean higher peak rates at 30 s, values less than 1, higher peak rates at 240 s.

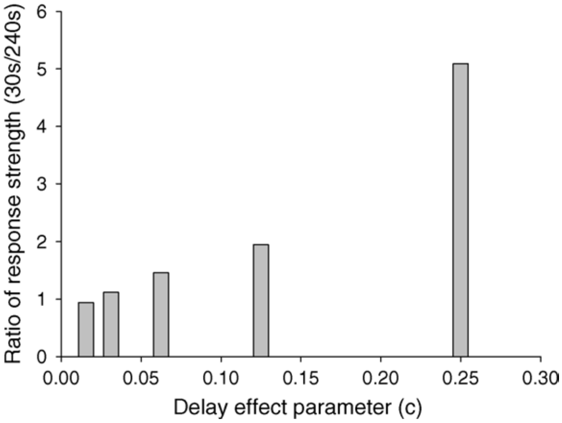

Fig. 5. Ratio of response strengths calculated from Eq. (2), for response periods on the separate components of mixed-FI 30 s FI 240 s. Results are shown as a function of the delay effect parameter (c in Eq. (2)).

It is obvious on inspection of Fig. 5 that the "strength" ratio depends strongly on c: high sensitivity to delay will produce higher rates at 30 s than 240 s, whereas smaller c values make rates much more similar (i.e. closer to a ratio of 1.0), and "reversals" are even possible (i.e. higher rates at 240 s than 30 s), when c is very small. Thus, our treatment of response rate using Eq. (2), while forbidding "reversals" on simple FI schedules, permits them on mixed-FI schedules, and in general is consistent with the result that the response rates at the times of reinforcement of the two components of the mixed-FI schedule are more similar than might be expected from the rates of reinforcement associated with each schedule alone.

The present account is "permissive" but not completely "predictive": it allows a range of possible outcomes for response rates on mixed-FI schedules but exactly which result is obtained in any particular case depends on the individual animal's sensitivity to delay (and the way that response "strength" is translated into observed response rate), which is unknown.

What happens to response rates in data when one of the components of the mixed-FI schedule is kept constant and the other varied, and what does the delay of reinforcement approach predict in these circumstances? The upper panel of Fig. 6 shows some results from Whitaker et al. (2003), Experiment 3a.

In this experiment the short component (S) of a 2-valued mixed-FI was kept at 30 s, and the longer component (L) varied across conditions over values of 45, 60, 120, and 240 s. Inspection of the averaged data from four rats (Whitaker et al., 2003, Fig. 5, p. 285) shows that only the 60, 120, and 240 s cases produced two distinct response peaks (one near S and one at L), so we will take data from these conditions. The focus of interest here is the ratio of the peak heights at S and at L, as L was varied. The upper panel of Fig. 6 shows that the S/L peak height ratio increased as L increased: in other words, the height at S became relatively greater than the height at L as L increased. A

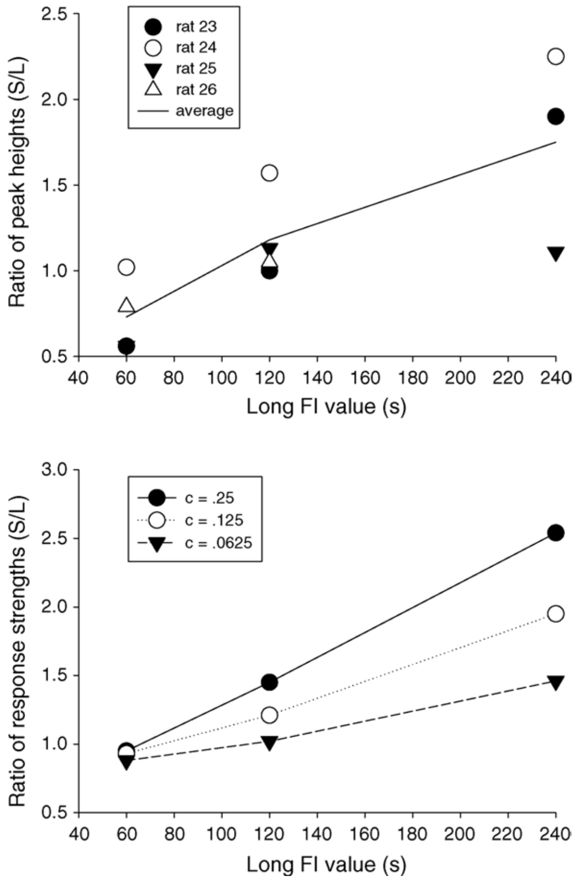

Fig. 6. Upper panel: ratios of peak heights (height at shorter FI/height at longer FI) from Whitaker et al. (2003), Experiment 3a. FI *S* was always 30 s, and results from FI *L* values of 60, 120, and 240 s are shown. One data point is omitted (see text for details). Symbols show data from individual rats and the line the average. Lower panel: relative response strengths (*S*/*L*) from the delay of reinforcement model with FI *S* of 30 s, and FI *L* varied over values of 60, 120, and 240 s. The delay effect parameter, *c*, was varied over values of 0.25, 0.125, and 0.0625.

data point from rat 26 at mixed-FI 30 s FI 240 s was omitted as the ratio produced (8.27) was abnormally large. It did, however, fit the general trend shown in Fig. 6. The upper panel of Fig. 6 also shows the average of the four rats, and on average the peak height ratio went from less than 1.0 (i.e. a “reversal”, a peak height greater at *L* than *S*) when *L* was 60 s, to a greater peak height at *S* when *L* was 120 and 240 s.

The delay of reinforcement model was used to simulate these conditions, by calculating the response strengths at *S* and *L*, using three different delay effect parameters (*c* in Eq. (2)), and the results are shown in the bottom panel of Fig. 6. The response strength ratio increased as *L* increased, from a value below 1.0 when *L* was 60 s to values considerably above 1.0 when *L* was 240 s, thus mirroring exactly the pattern of results shown in the data.

Overall, therefore, on 2-valued mixed-FI schedules with equally probable components, the delay of reinforcement model captures some of the complexity of the response rate data obtained. Peak heights can be equal, lower, or higher at *L* than *S*, even when *S* and *L* are very different, and changes in the relative value of *S* and *L* appears to change peak heights in a way consistent with the model (e.g. Fig. 6).

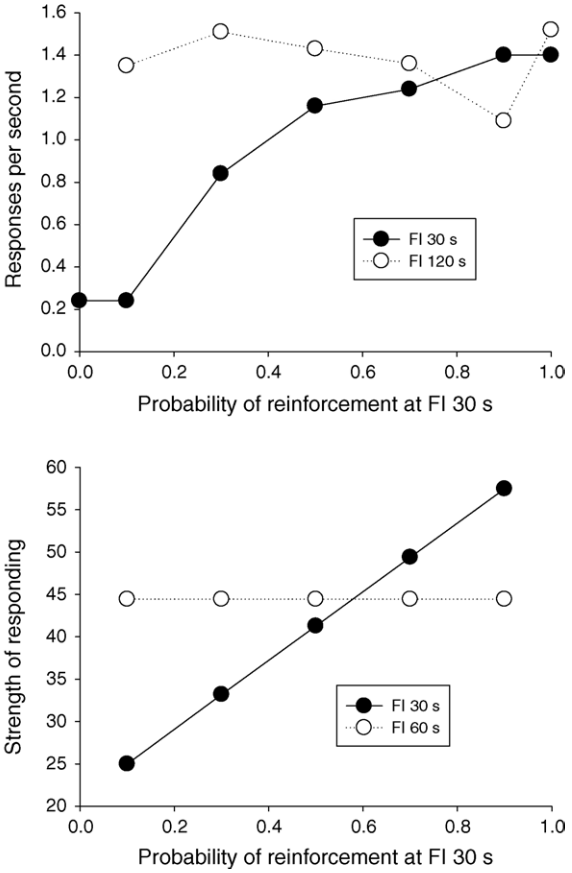

Fig. 7. Response strengths for the separate components of mixed-FI 30 s (filled circles) FI 60 s (unfilled circles), as the probability of reinforcement according to the two components is varied. Results are shown plotted against the probability of reinforcement according to FI 30 s.

5. Probability of reinforcement on mixed-FI schedules

Two-valued mixed-FI schedules allow potential separation of factors relating to temporal control (i.e. at which elapsed times since food responses are reinforced) from those relating to response rate (such as probability of reinforcement at these times). For example, on some mixed-FI schedule FI *S* FI *L*, the probability of reinforcement for responses at *S* and *L* can be manipulated when *S* and *L* remain constant. Whitaker (1979) conducted a series of experiments of this sort, and the upper panel of Fig. 7 shows some data from one of his conditions. The schedule used was mixed-FI 30 s FI 60 s, with data coming from the FI 60 s intervals, collected into 6 s bins. The response rate data are terminal rates, i.e. rates from the bins just before 30 s and just before 60 s, and the probability of reinforcement for responses at 30 s was varied over values of 0 (i.e. FI 60 s), 0.1, 0.3, 0.5, 0.7, 0.9, and 1.0 (i.e. simple FI 30 s). The probability of reinforcement according to FI 60 s was just 1 minus the probability at FI 30 s.

The striking feature of the behaviour obtained was that terminal rates at around 30 s were markedly affected by the probability of reinforcement according to FI 30 s, increasing with increasing reinforcement probability at FI 30 s, whereas terminal rates at around 60 s remained roughly constant. We used the same logic as for equiprobable two-valued mixed-FI schedules, discussed above, to simulate this case, with the response period at FI 30 s being affected by both "immediate" and "delayed" reinforcement, with the probability of each depending on the probability of reinforcement at FI 30 s, with the FI 60 s response period being only reinforced "immediately". The lower panel of Fig. 7 shows that this calculation produced results in striking qualitative correspondence with the data: response "strength" at FI 30 s increased systematically with increasing reinforcement probability at this value (as the "immediate" reinforcement effect became stronger and stronger), whereas responding at FI 60 s was unaffected.

We should acknowledge, however, that some data in Whitaker (1979) did not show the effect of constant terminal rate at the longer FI component when the probability of reinforcement according to the shorter FI was varied, but the data in Fig. 7, and the delay model calculation, are also in accord with some earlier data from mixed-FI schedules, such as those collected by Catania and Reynolds (1968) who likewise found that varying the probability of reinforcement according to FI *S* markedly affected responding around *S*, but had much less of an effect around *L*.

6. Discussion

The present article is intended to provide a brief exploration of the simple idea that the response rate on simple and mixed-FI schedules can be understood in terms of the average delay of reinforcement for responses during the response periods that occur on the schedules. This simple idea has some striking successes: it predicts the "confinement" effect that occurs during acquisition of FI, it accounts for increases in response rate with reinforcement rate on simple FI schedules, and even predicts the correct form for this increase, it forbids "reversals" of rate effects on simple FI but permits them on mixed-FI schedules, and also qualitatively fits some data on effects of varying the relative values of the two components making up the mixed-FI schedule, and changes in reinforcement probability on mixed-FI schedules.

On the other hand, there are some data from mixed-FI schedules that the model cannot deal with, and a more general problem is that of deriving exact response rates from the response strength measures that the model generates. A simple direct multiplication by some scaling constant is one obvious possibility, but there are other more complex potential transformations, which are not dealt with here. In addition, the model has some simplifying features which are almost certainly too simple. For one thing it uses a very simple version of hyperbolic reinforcement delay (Eq. (2)), and for another it assumes that the average delay of reinforcement for responses during some response period is always half the length of the response period. This latter assumption is tantamount to assuming that the local rate of response during a response period is constant (i.e. a "break and run" pattern of responding), whereas if responding accelerates throughout a response period (i.e. a "scalloped" pattern of responding), then there will be more responses in the second half of the period than the first, thus the experienced delay will be on average less than half the response period. However, given that response rates on FI and mixed-FI schedules have received only sketchy theoretical treatment up until now, a simple model seemed to us the best starting point.

Although the issue of response rate on FI schedules has been neglected compared with questions related to temporal control, the neglect has not been absolute, and work by Machado (1997) and the "packet theory" of Kirkpatrick (2002) have both attempted to deal with some aspects of response rate. Machado's theory is radically different from our account in many ways, but the difference that we want to emphasize here is that we presume that temporal control and response rate on FI are determined by different mechanisms, whereas in Machado's account a common process generates both. Machado (1997) presumes that timing on FI is the result of a series of successive "states" (perhaps identified with adjunctive behaviours, as in Killeen and Fetterman, 1988, although this is not always clear) each of which can become associated with the operant response if the response is reinforced when the state is active. So, for example, on a simple FI schedule, states that occur at the time of reinforcer delivery become associated with the response, and those occurring earlier have any association with the response that they might have had gradually extinguished. This theory thus correctly predicts an increasing tendency to respond as time in the interval passes, as is observed in data. Machado's model can also deal with responding on mixed-FI schedules, as states active at the time of *S* and the time of *L* both become associated with the operant. However, states active near *S* have their associations sometimes strengthened (when reinforcement is delivered according to FI *S*) and sometimes weakened (when reinforcement is delivered according to FI *L*), because responding at or near *S* is not reinforced on this sort of trial. The result of this is that states active at *S* have their association with the reinforcer partially extinguished, whereas states active near *L* never have their associations weakened. This leads to the prediction that response rates are *always* higher at *L* than at *S*, a violation of the usual result that schedules associated with higher rates of reinforcement generate higher rates of response. Although higher rates at *L* than *S*, what we have called "reversals", can be observed in data, they are less common than higher rates at *S* than *L*, or similar rates at both, and our account allows all these possibilities. Machado's model can also deal with acquisition of FI responding (see also Machado and Cevic, 1998), and the increasing association of certain states with the operant response as acquisition progresses may also predict the phenomenon which we have called "confinement".

An even more recent account of animal performance on FI and other schedules, the "Packet theory" of Kirkpatrick and Church (Kirkpatrick, 2002; Kirkpatrick and Church, 2003) also deals with some aspects of response rate on FI-like schedules. Packet theory assumes that responses emitted by animals occur in "packets", which are periods containing a group of responses, usually separated from other packets (although packets can overlap, giving the impression of a continuous period of responding). Both the probability of emission of a packet of responses at any particular time, and the "contents" of the packet (i.e. the number and rate of response contained), can vary according to experimental contingencies. Packet theory has a very wide application (see, for example, Kirkpatrick and Church, 2004) and has many aspects that are not relevant to our discussion here.

Packet theory provides an account of both temporal control and response rate. On a schedule like FI (in fact, most operant work using packet theory uses a magazine-entry response and response-independent fixed-time [FT] schedules, but for simplicity we will assume here that these conditions are closely analogous to FI), the probability of emission of a packet is controlled by the average conditional expected time to reinforcement which is translated into an expectation of reinforcement at any time in the interval, *t*. To use Kirkpatrick's (2002, p. 93) own example, on FI 90 s the expected time to food is 90 s at the start, 89 s after 1 s, and so on, so expectancy increases linearly as time in the interval elapses. This expectancy controls the probability of emission of a packet, so one prediction is that response rate increases with increasing elapsed time in the FI, whereas it remains constant over time (as the expectation of food over time is constant) on an aperiodic schedule like a random-interval schedule.

To simulate response rates, a number of additional assumptions can be made. One is that the number of packets per interval is on average constant, whatever its length, so packet theory naturally produces the effect of decreasing response rates with increasing interval length (Kirkpatrick and Church, 2003, p. 13), as the probability of packet emission per unit time decreases as the interval lengthens, although it is less clear that the form of the response rate versus reinforcement rate predicted is in accord with Herrnstein's equation itself. To predict the obtained response rate on FT schedules ranging from FT 45 to 360 s Kirkpatrick (2002, p. 97) kept the number of packets per interval constant, but also varied another parameter (the probability of firing a reactive packet), which is apparently related to response rate, so it is unclear whether both parameters need to be adjusted to simulate real response rates.

Although packet theory can deal with response rate on FI, it is less clear that it can deal with the phenomenon which we call confinement. Our Fig. 3 suggests that rats learn the time of reinforcement very quickly, thus their expected times to reinforcement as the interval elapsed would presumably be stable from the early sessions of training. Nevertheless, as the response period decreases, response rate increases. The decrease in the duration of the response period has a natural explanation if we assume that some threshold for initiating responses becomes stricter as training proceeds, and this shortening of the response period inflates response rate by our delay of reinforcement mechanism.

Data from mixed-FI schedules also appear to provide difficulties for packet theory, at least in its simple form. Consider, for example mixed-FI 30 s FI 240 s, where the two components are equiprobable. At the start of any interval the average time to reinforcement is 135 s ([30 + 240]/2), but data show response peaks at 30 and 240 s, with a low rate at around 135 s (e.g. Whitaker et al., 2003, Fig. 2, p. 282), so average expected time to reinforcement does not provide a good account of temporal control on these schedules. Once the shorter interval in the mixed-FI schedule has passed, however, then expected time to reinforcement may give a good account of the general features of response emission as response rate would increase on average as time passed, towards a peak at the longer of the two FI values, as is found in data. It also appears that varying the probability of reinforcement for the different components of a 2-valued mixed-FI schedule produces results that packet theory cannot deal with without modification, as this manipulation changes the average expected time of reinforcement, so should change the pattern of temporal control observed, whereas the data show peaks at the two FI values, with changes in the peak heights (but not peak location) when reinforcement probability changes, at least for the shorter interval (see our Fig. 7, for example). Packet theory may also have difficulty predicting the response rates occurring at the times of the two FI values in the mixed-FI schedule, particularly the "reversals" (i.e. higher rate at *L* than *S*) which do occur in data, as presumably a principle like that of keeping the number of packets in an interval constant never allows higher rates of response on schedules associated with lower rates of reinforcement.

Although we have drawn attention to apparent difficulties of packet theory with some of the data we discuss in the present article, we readily admit that modifications of the theory might overcome most of these apparent problems. For example, on mixed-FI schedules, the two components might be timed separately, rather than giving rise to an average expected time to reinforcement, and the "decision" component of the packet theory model might be used to deal better with the phenomenon of confinement.

7. Conclusions

The present article had two general aims. One was to draw attention to some potentially interesting, and relatively neglected, problems concerning response rate on FI and related schedules. Another was to explore the idea that the use of a simple principle, that of average delay of reinforcement for responses occurring during the response periods that such schedules generate, can make some progress in understanding at least some of these phenomena. Among our successes were the prediction of the Herrnstein's equation relation for responding on FI, "confinement", and some aspects of responding on mixed-FI schedules, particularly allowing "reversals" of response rate to occur, and predictions of some of the effects of reinforcement probability on mixed-FI. In our treatment, we have ignored questions of temporal control and concentrated exclusively on response rate (the reverse of the SET position which deals with temporal control but not rate), whereas some other recent accounts (like *Learning to Time*, Machado, 1997, and packet theory [Kirkpatrick, 2002]) provide integrated accounts of temporal control and response rate. Whatever the ultimate outcome of debates between the different theoretical models developed to study animal timing, our article will have been successful if it stimulates some interest in the problem which so occupied Skinner in 1938, that of how to understand the rate of responding that occurs on FI schedules of reinforcement.

Acknowledgement

Preparation of this article was supported by a Joint Project grant from the Royal Society.

References

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

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., 1984. The application of scalar timing theory to individual trials. J. Exp. Psychol.: Anim. Behav. Process. 20, 135–155.

Dews, P.B., 1970. The theory of fixed-interval responding. In: Schoenfeld, W.N. (Ed.), The theory of reinforcement schedules. Appleton Century Crofts, New York, pp. 43–61.

Dews, P.B., 1978. Studies on responding under fixed-interval schedules of reinforcement. Part II. The scalloped pattern of the cumulative record. J. Exp. Anal. Behav. 29, 67–75.

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.

Dukich, T.D., Lee, A.E., 1973. A comparison of measures of responding under fixed-interval schedules. J. Exp. Anal. Behav. 20, 281–290.

Elsmore, T.F., 1971. Independence of post-reinforcement pause length and running rate on FI pacing reinforcement schedules. Psychon. Sci. 23, 371–372.

Ferster, C.B., Skinner, B.F., 1957. Schedules of Reinforcement. Appleton Century Crofts, New York.

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, 423: Timing and Time Perception. New York Academy of Sciences, New York.

Gollub, L.R., 1964. The relation among measures of performance on fixed-interval schedules. J. Exp. Anal. Behav. 7, 337–343.

Herrnstein, R.J., 1970. On the Law of Effect. J. Exp. Anal. Behav. 13, 243–266.

Killeen, P., 1988. The reflex reserve. J. Exp. Anal. Behav. 50, 319–331.

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

Kirkpatrick, K., 2002. Packet theory of conditioning and timing. Behav. Process. 57, 89–106.

Kirkpatrick, K., Church, R.M., 2003. Tracking the expected time to reinforcement in temporal conditioning procedures. Learn. Behav. 31, 3–21.

Kirkpatrick, K., Church, R.M., 2004. Temporal learning in random control procedures. J. Exp. Psychol.: Anim. Behav. Process. 30, 213–228.

Lejeune, H., Richelle, M., Wearden, J. H., in press. About Skinner and time: behavior-analytic contributions to research on animal timing. J. Exp. Anal. Behav.

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

Lowe, C.F., Harzem, P., Spencer, P.T., 1979. Temporal control of behavior and the power law. J. Exp. Anal. Behav. 31, 333–343.

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

Machado, A., Cevic, M., 1998. Acquisition and extinction under periodic reinforcement. Behav. Process. 44, 237–262.

Mazur, J.E., 1987. An adjusting procedure for studying delay of reinforcement. In: Commons, M.L., Mazur, J.E., Nevin, J.A., Rachlin, H. (Eds.), Quantitative Analysis of Behavior, vol. 5, Effects of Delay and Intervening Events on Reinforcement Value. Hillsdale, Erlbaum, NJ, pp. 55–71.

Mechner, F., Guevrekian, L., Mechner, V., 1963. A fixed-interval schedule in which the interval is initiated by a response. J. Exp. Anal. Behav. 6, 323–330.

Nevin, J.A., 1992. An integrative model for the study of behavioral momentum. J. Exp. Anal. Behav. 57, 301–316.

Nevin, J.A., Grace, R.C., 2000. Behavioral momentum and the law of effect. Behav. Brain Sci. 23, 73–130.

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

Schneider, B.A., 1969. A two-state analysis of fixed-interval responding in the pigeon. J. Exp. Anal. Behav. 12, 677–687.

Shull, R.L., 1971. The sequential patterns of post-reinforcement pauses on fixed-interval schedules of food. J. Exp. Anal. Behav. 15, 221–232.

Skinner, B.F., 1938. The Behavior of Organisms: An Experimental Analysis. Appleton Century Crofts, New York.

Spencer, P.T., 1978. Some determinants of the postreinforcement pause on temporally-defined schedules. Unpublished Ph.D. thesis, University College of North Wales, Bangor.

Spencer, P.T., 1981. Temporal control of behavior and the law of effect: a description of fixed-interval performance. Behav. Anal. Lett. 1, 325–329.

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.

Whitaker, J.S., 1979. 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.