The structure of random ratio responding in humans

Phil Reed

Swansea University

Correspondence address: Phil Reed,

Department of Psychology,

Swansea University,

Singleton Park,

Swansea, SA2 8PP, U.K.

Short title: Human schedule performance

Cite as: Reed, P. (2015). The structure of random ratio responding in humans.. *Journal of Experimental Psychology: Animal Learning and Cognition* 41(4), 419-431.
doi:10.1037/xan0000081

Abstract

Four experiments explored the rate and structure of human responding on random ratio (RR) schedules of reinforcement, using three different methods of analysis. All experiments demonstrated that responding on RR schedules was higher with intermediate ration values – rates being higher on an RR-30 schedule than on either a RR-10 or RR-60 schedule. All experiments identified two distinct types of responding generated by RR schedules: 'response-initiation' responding and 'within-burst' responding. In contrast to overall rates, response-initiation responding was greatest on RR-10 schedules, and less pronounced as the ratio value increased, but response rates did not vary greatly within the burst (being higher on longer ratio schedules when there was a difference). These findings are consistent with previously obtained data from nonhumans, and suggest human schedule performance should be considered as a joint product of two different aspects of responding: response-initiation responding which is sensitive to rates of reinforcement; and within-burst responding which is less sensitive to the molar characteristics of the schedule in operation. The presence of two distinct forms of responding may require reinterpretation of many human schedule-controlled phenomena that rely on assumptions about the similarity of all response to one another.

Key words: random ratio, schedule, response initiation, within burst responding, schedules of reinforcement, humans.

The nature of performance controlled by various schedules of reinforcement has been shown to be highly consistent across many species (see Zeiler, 1977, for a review). For example, rates of responding on random ratio (RR) schedules typically are higher than those generated on random interval (RI) schedules, even when the rate of reinforcement is equated on the two types of schedule (e.g., Ferster and Skinner, 1957; Peele, Casey, and Silbeberg, 1987; Zuriff, 1970). The factors that control such schedule performance in nonhumans have also been well investigated, and these factors encompass both 'molecular' factors, such as the reinforcement of different inter-response times (IRTs; see Morse, 1996; Peele et al., 1987), and 'molar' factors, such as the rate of reinforcement (see Baum, 1973), and the response-reinforcer feedback function (see McDowell and Wixted, 1986). The factors suggested by molecular and molar theories should not necessarily be considered as incompatible with one another, but will tend to operate at different times and under different circumstances from one another (see Reed, 2001; 2007; 2015).

In addition to the factors that control free-operant schedule performance, there has been some exploration of the structure of responding on various schedules in nonhuman subjects (see Shull, 2011, for a review). Using a variety of different procedures and techniques to analyze schedule behavior, two different components of responding have emerged with relative consistency – 'response-initiation' and 'within-burst' behaviors (Brackney, Cheung, Neisewander, and Sanabria, 2011; Killeen, Hall, Reilly, and Kettle, 2002; Mellgren and Elsmore, 1991; Pear and Rector, 1979; Reed, 2011; Shull, Gaynor, and Grimes, 2001; Sibley, Nott, and Fletcher, 1990). Furthermore, it is suggested that these two response types are differentially sensitive to various aspects of the contingency (Brackney et al., 2011; Reed, 2011; Shull et al., 2011).

Response-initiation responses commence a 'bout' (or 'burst') of responding (Sibley et al., 1990); these responses are sensitive to the overall rate of reinforcement on an interval schedule, and, somewhat relatedly, their rate of emission varies inversely with the number of responses required by a ratio schedule (see Shull et al., 2001; Shull and Grimes, 2003). However, the relationship of within-burst response rates to overall rates of reinforcement is less clear, and may depend on the manner in which they are assessed, with some procedures appearing not to find great variance in the rates at which within-burst responses are emitted (e.g., Shull et al., 2001), while other studies finding that they are emitted faster as the size of the burst increases (see Reed, 2011).

In contrast to the investigation of the impact of schedules of reinforcement on nonhuman performance, relatively little is known about human schedule-controlled behavior. In general, humans tend to emit greater response rates on RR compared to RI schedules (see Bradshaw and Reed, 2012; Matthews, Shimoff, Catania, and Sagvolden, 1977; Raia, Shillingford, Miller, and Baier, 2000; Reed, 2001). However, human responding on schedules of reinforcement can also show behavioral patterns that differ from nonhuman performance (e.g., Leander, Lippman, and Meyer, 1968; Lowe, 1979; Weiner, 1970), which may be due to a wide variety of factors. For instance, the type of reinforcement employed (Lowe, Harzem, and Bagshaw, 1978), the delivery style of the reinforcers (Matthews et al., 1977; Raia et al., 2000), and the physical exertion necessitated by the experimental task (Matthews et al., 1977; Reed, 2001), may all impact on schedule performance. These differences may also be due to the presence and influence of verbal behavior in humans that has been shown to impact behavior (see Baron and Galizio, 1983; Bradshaw and Reed, 2012; Lowe, Beasty, and Bentall, 1983; Skinner, 1969). Moreover, there is some evidence to

suggest that personality types in humans, such as schizotypy (e.g., Randell, Ranjith-Kumar, Gupta, and Reed, 2009), or depressive traits (e.g., Alloy and Abramson, 1979; Dack, McHugh, and Reed, 2009; Reed, Frasquillo, Colkin, Liemann, and Colbert, 2001), may also impact on human schedule performance. However, when these factors are addressed by the implementation of appropriate procedures, such as giving minimal verbal instructions, adopting concurrent tasks designed to suppress the formation of verbal rules, and introducing costs for responding (see Raia et al., 2000), humans typically respond at higher rates on RR than on RI schedules (see Reed & Bradshaw, 2012; Dack et al., 2009; Raia et al., 2000).

To further understand the nature of schedule-controlled human behavior, the current experiments explored the 'micro-structure' of human performance on RR schedules – bringing this investigation into line with recent studies of nonhuman performance (e.g., Brackney et al., 2011; Reed, 2011; Shull et al., 2001; Shull, 2011). Although there is debate regarding the degree to which human performance is sensitive to variation in interval schedules (see Leander et al., 1968; Lowe, 1979), human performance has been shown to be consistently sensitive to the properties of ratio schedules. The evidence relating to the impact of RR schedules comes from a variety of experimental procedures, including probability learning studies (see Estes, 1976, for a review), causal judgment experiments (e.g., Alloy and Abramson, 1979; Dack et al., 2009; Wasserman, Chatlosh, and Neunaber, 1983), and reports of schedule-maintained performance (see Bradshaw and Reed, 2012; Dack et al., 2009; Reed, 2001). For this reason, the current experiments employed a variety of RR schedules, and aimed to investigate both whether the structure of responding on these schedules is similar in humans to nonhumans in terms of the potential existence of different components of responding (see Reed, 2010; Shull et al., 2001), and whether

any putative response components in humans would be controlled by similar factors to those noted in nonhumans.

As noted above, there are a number of different procedures that can be adopted to study such response structures on free-operant schedules, and the current manuscript investigated whether similar results would occur when the different approaches were employed. To this end, Experiments 1 and 2 compared two different methods of analysing the data – Experiment 1 employed the log survivor method (Shull, 2011), and Experiment 2 employed the cut-off method (Mellgren and Elsmore, 1991), of investigating the micro-structure of responding. Both of these studies were conducted on steady-state responding, whereas Experiment 3 was an examination of these response patterns in acquisition. Experiment 4 adopted a dual manipulanda approach to this issue (see Reed, 2011).

The main aim of this study was to explore the structure of human responding on RR schedules of reinforcement. Although this is a primarily empirically driven objective, the micro-structure of human responding has, thus far, not been established, and determining this structure would be a useful addition to knowledge regarding such human learning; especially if the factors that control their rates of emission could also be established for this species. However, if the two different types of response components were demonstrated for humans, then this would also have some theoretical implications for the interpretation of previous results using ratio schedules in humans. For example, in studies of human causal judgment, while judgements of causal effectiveness typically follow the probability of an outcome following a response (e.g., Alloy and Abramson, 1979; Dack et al., 2009; Wasserman et al., 1983), this is not always the case (see Reed, 1999), and understanding the structure of responding on ratio schedules might help to understand which types of responses are

being judged in terms of their causal relationship to the outcomes. Similarly, an empirical demonstration of the existence of two forms of responding for human schedule performance might also allow subsequent investigation of the influence of factors such as explicit verbal control of responding, which is regarded as important for humans (see Bradshaw and Reed, 2012). For example, do explicit instructions impact on all responses, or just on response-initiation responses? To work towards these ends, the current study explored human responding on ratio schedules.

Experiment 1

Much of the data relating to the structure of nonhuman schedule performance has been derived from an analytic procedure developed by Shull et al. (2001; see also Brackney et al., 2011; Killeen et al., 2002, Sibley et al., 1990, for similar variants of this approach). The frequency distribution of the inter-response time (IRTs) is created, and the percentage calculated of IRTs emitted in a particular time bin as a proportion of all IRTs not yet emitted. These survival percentages are then turned into logs, and a 'log survivor plot' is generated from these data (see Shull, 2011; Shull et al., 2001). In a log survivor plot, the slope between any two points is an indicator of the relative decline in the proportion of the IRTs per opportunity between those points: the steeper the slope, the higher the relative rate of responding during the interval (see top panel of Figure 1 for an idealized version of such a plot showing the two arms of the plot).

Such log survivor plots produce distinct patterns of data for nonhumans, and have been used to explore rates of response-initiation, and numbers of responses per burst (Brackney et al., 2011; Shull et al., 2001). If responding is emitted at a constant

rate across a session, then the IRT distribution would appear as a single decreasing straight line (see Shull, 2011, for a discussion). However, Shull et al. (2001; Brackney et al., 2011; Killeen et al., 2002) found that the slope of log survival plots for rats were not uniform, but comprised a sharply decreasing initial portion, followed by a portion with a shallower negative gradient (termed a 'broken-stick' appearance).

In contrast to the data on the effects of RI schedules on these different types of responses outlined above (see also Shull, 2011), there is relatively little information for the effects of RR schedules, and none with humans as subjects. From somewhat different but related investigations, it can be suggested that the rate of response-initiation increases with shorter ratio values (Reed, 2011; Shull, 2004), with rates of within-burst responding tending to follow the overall response rates (i.e. showing an inverted-U function with ratio value; Ferster & Skinner, 1957; Reed, 2011; 2015; Reed and Hall, 1988). To determine whether such an analysis would yield similar results from human participants, the current study employed this technique to analyze steady-state responding in humans maintained by RR-10, RR-30, and RR-60 schedules.

Method

Participants

Twelve participants (4 males and 8 females), aged between 19-35 years, were recruited, and received Psychology Department subject-pool credits but no financial payment. All participants had normal or corrected-to-normal vision, and were naïve to the experiment's purpose. As previous studies have shown that individuals scoring highly in terms of depression and schizotypy can show atypical patterns of schedule performance (see Dack et al., 2009; Randell et al., 2009), psychometric tests were

employed to exclude participants with high scores on these dimensions. In this experiment, two participants were excluded on the basis of having a high depression score, leaving 10 participants in the study.

Apparatus

The experimental task was presented using Visual Basic (6.0) on a laptop computer with a 15.6-inch screen. The program presented an RR schedule (10, 30, or 60) to the participants. On a particular schedule, each response (a spacebar press) had an equal probability of reinforcement (i.e., 1/10, 1/30, or 1/60). Each participant began the experiment with 40 points, displayed in a box, under the word "points", in the middle of the screen, approximately one third of the way from the bottom of the screen. A colored square (either blue, purple, or yellow), approximately 8cm wide x 3cm high, was displayed in the middle of the screen, approximately one third from the top of the screen. Reinforcement consisted of 60 points being added to the 'points' box. Each response subtracted 1 point from the 'points' box, which aimed to prevent a lack of performance regulation in humans that can occur when there is no cost for a response (Bradshaw and Reed, 2012; Reed, 2011).

Measures

Oxford Liverpool Inventory of Feelings and Experiences - Brief Version (O-LIFE(B); Mason, Linney, and Claridge, 2005) measures schizotypy and has a Cronbach α between .62 to .80 (Mason et al., 2005). A score of greater than 6 on the Unusual Experiences scale (one standard deviation above the mean, Mason et al., 2005) was taken as a cut-off point for high levels of this trait associated with atypical schedule performance (see Randell et al., 2009).

Beck's Depression Inventory (BDI; Beck, Ward, Mendelson, Mock, and Erbaugh, 1961) assesses depression, and has a Cronbach α between .73 and .92 (Beck, Steer, and Garbin, 1988). A score of higher than 10 was taken as a cut-off for high levels of depression, and potentially producing atypical schedule performance (see Dack et al., 2009).

Procedure

Participants were tested individually in a quiet room containing a desk, a chair, and a computer. Participants read the written instructions, and completed the BDI and the O-LIFE(B) scales. After this, the following instructions were presented on the computer screen:

"*This experiment is concerned with multi-tasking abilities. You have two tasks to complete simultaneously. In the first task, you must count backwards, out-loud, in 7s, from the number 26,758. The second task is to score as many points as possible by pressing the space bar on the computer. The coloured shape may be important. To receive points, sometimes you might need to press the spacebar quickly and at other times you might need to press slowly. The person with the best score on both tasks will receive a £50 [name of company] token."*

Each participant was then exposed to all three schedule types (RR-10, RR-30, and RR-60). Each schedule was presented once to each participant, with each schedule exposure lasting 10min, with a 30s inter-component interval. Each different schedule was signaled by the presence of a different colored rectangle on the screen. The particular colors used to signal the schedules, and the order of schedule presentation, was randomized. Each response subtracted one point from the 'points' box displayed on the screen. This response cost procedure was adopted to make each

response meaningful to the participant (see Bradshaw and Reed, 2012; Reed, 2001). Reinforcement consisted of the addition of 60 points to the 'points' box.

During the time in which they were performing on the schedules, the participants had to perform the counting backwards task (see Andersson, Hagman, Talianzadeh, Svedberg, and Larsen, 2002). They were given a five-digit number, and were asked to count backwards, out-loud, in 7s. This procedure was adopted in an attempt to minimize the potential role of verbal rule formation in influencing participants' performance on the schedule (see Bradshaw and Reed, 2012; Leander et al., 1968; Lowe, 1979; Raia et al., 2000). In order to enhance task adherence, a recording device was placed prominently on the desk in front of the participant, and they were told that their answers to the counting task would be analyzed and scored later.

Results and Discussion

The responses emitted during the last 5min of exposure to each of the three schedule types were analyzed, as they were taken to represent terminal performance on each schedule. The overall response rates for the three schedules (total responses emitted during the last 5min divided by 5) were: RR-10 = 160.10 ($\pm$ 86.96); RR-30 = 204.13 ($\pm$ 79.90); and RR-60 = 203.08 ($\pm$ 84.72). As these data (and all data reported here) were reasonably normally distributed (no Shapiro-Wilk test conducted on any analyzed variable reached a $p < .01$ rejection level), they were subject to parametric analyses. A one-way, repeated-measures analysis of variance (ANOVA) revealed a statistically significant effect of schedule condition, $F(2,18) = 4.32$, $p < .05$, *partial* $eta^2 = .324$, *95% CI* = .000:.540. Paired t-tests revealed significant differences between the RR-10 and RR-30 conditions, $t(9) = 3.60$, $p < .01$, and between the RR-

10 and RR-60 conditions, $t(9) = 2.12$, $p < .05$. This pattern of results is broadly consistent with patterns of responding seen in nonhuman RR schedules (see Reed & Hall, 1988).

The overall rates of reinforcement (total numbers of reinforcement obtained in the last 5min of schedule exposure divided by 5) were: RR-10 = 16.01 (± 8.67); RR-30 = 6.80 (± 2.66); and RR-60 = 3.89 (± 1.41). An ANOVA revealed a significant difference between the conditions, $F(2,18) = 24.51$, $p < .001$, *partial eta*^{2} = .731, *95% CI* = .414:.824. Paired t-tests revealed statistically significant differences between the: RR-10 and RR-30 conditions, $t(9) = 4.54$, $p < .001$; RR-10 and RR-60 conditions, $t(9) = 5.16$, $p < .01$; and RR-30 and RR-60 conditions, $t(9) = 5.94$, $p < .001$.

A log survivor analysis was employed to examine the pattern of responding emitted during the last 5min of each of schedule. The IRT data for each rat were analyzed by temporal bin (100ms), and the number of responses in each bin was calculated. The number of responses emitted during a 100ms bin was turned into a percentage of the number of responses that were not already emitted in a preceding bin. These percentage data were transformed to logarithms, and log survivor plots for all schedules were generated for each participant. These data were fitted to an equation: $P(IRT>t) = (1-p_{(D)})^{e-wt} + p_{(D)}^{e-bt}$; described by Shull (2011), on which the following description is based: $P(IRT>t)$ represents the proportion of IRTs longer than t (i.e., the proportions that make up the overall log survivor plot). The left hand exponential expression describes the survivor plot for within-bout response times, and right hand expression describes the survivor plot for response-initiation responses. The term $p_{(D)}$ indicates the proportion of all responses that are bout-initiations, and $1-p_{(D)}$ indicates the proportion of all responses that are within-bout. The term e represents the base of natural logarithms; w represents the within-bout response rate; b

represents the bout-initiation rate (determined with respect to between-bout pause time); and *t* represents elapsed time since the last response (i.e., the bin values). Assuming that the equation adequately describes the properties of the IRT distribution, the numerical estimates of the key components of response rate can be based fitting this equation to the data for each subject, and employing the best fitting values of *b* (bout-initiation rate), and *w* (within-bout response rate).

Figure 1 about here

The bottom panel of Figure 1 displays the group-mean log survival plots for the three schedules, which show a reasonable approximation to those generated from the study of nonhumans. Visual inspection of these data reveals that there was a 'broken stick' appearance, in that there were two distinct gradients for the plots. The descending limbs of these plots were longer for the longer RR values, suggesting more responses were emitted as the RR value increased (Shull et al., 2001). The slope of the post break points also appeared somewhat steeper for the longer RR schedules, suggesting higher rates for these schedules. This finding has also been noted for nonhumans using procedures with some ratio-like aspects (see Reed, 2011). These descriptions were analyzed numerically by fitting the equations noted above to the data (see below).

Figure 2 about here

Figure 2 displays the group-mean response-initiation rates over the last 5min of exposure to each schedule (solid bars). These rates declined as the ratio value

increased, with the linear trend being significant, $F(1,9) = 9.13$, $p < .01$ *partial eta*^{2} = .504, *95% CI* = .025:.723, but not the quadratic trend, $F < 1$. This finding has been found in experiments exploring the impact of various schedule parameters on nonhuman RR responding using a different procedure to that employed in the current study (Reed, 2011), but was not noted in studies of nonhuman responding that have employed this analytic procedure (Bowers, Hill, and Palya, 2008). The fact that the current study, and that of Reed (2011), both found that response initiation rates varied with the ratio size might reflect the large differences in reinforcement rate noted between the RR schedule conditions employed in these studies, which has been shown to predict rates of response initiation (Shull, 2011).

Figure 2 also shows the group-mean within-burst rates (diagonal lines), which were higher in the RR-30 and RR-60 schedules than in the RR-10 schedule. Both the linear trend, $F(1,9) = 4.50$, $p < .05$, *partial eta*^{2} = .333, *95% CI* = .014:.652, and the quadratic trend, $F(1,9) = 3.89$, $p < .05$, *partial eta*^{2} = .302, *95% CI* = .001:.604, were significant. Paired t-tests revealed significant differences between the RR-10 and RR-30, $t(9) = 3.43$, $p < .01$, and RR-10 and RR-60, $t(9) = 2.03$, $p < .05$, conditions. This finding was also noted by Reed (2011) in a study of nonhuman RR performance.

Thus, analyzing human RR performance using the log survival technique (Shull et al., 2001; Sibley et al., 1990) produced broadly similar patterns of results to those found for nonhumans: overall response rates were higher at intermediate ratio values (i.e. RR-30) than at lower or higher ratio values (Reed and Hall, 1988); response initiation rates were greater on shorter RR schedules (Reed, 2011); and within-burst rates were greater with larger ratio values (which is not always noted; cf. Bowers et al., 2008; Shull et all., 2001).

Experiment 2

As the results from Experiment 1 were among the first to document the nature of the structure of human RR schedule responding, Experiment 2 sought to replicate these findings using a different procedure to ensure that the findings were of some generality. As noted previously, the log survival plot method adopted in Experiment 1 does not adopt a priori rules regarding the classification of performance into response-initiation and within-burst components (see Shull, 2011; Sibley et al., 1990), which is an advantage. However, it should be noted that the modelling approach also has some limitations. Firstly, the selected model might not fit the data, and may need alteration (see Brackney et al., 2011; Kessel and Lucke, 2008). Secondly, the modelling approach can only easily identify distribution-level parameters (e.g., mean IRT values, variance values), and not whether a particular response is a within-bout or bout-initiation. Thirdly, reliable parameter estimates require rather large data sets (i.e., lots of IRTs). Given the limitations of the log survivor analysis, Experiment 2 explored the structure of human RR responding, and to determine if similar results would be achieved using a different analytic methodology.

To this end, the analytic method described by Mellgren and Elsmore (1991; see also Tolkamp and Kyriazakis, 1999) was adopted. This method sets out a priori rules for the categorization of responses into those that are assumed to be response-initiation (i.e., those following some pre-specified length of pause from responding), and those that are within-bout responses (i.e., those that are emitted in close temporal proximity to one another). As suggested by Mellgren and Elsmore (1991), a range of cut-off points was adopted to ensure the results were not the product of fortuitous

choice. Additionally, each subject's performance was analyzed using the log-survivor analysis outlined above and the temporal bin at which the gradient changed (i.e. the point at which two arms of the lines crossed) was obtained to estimate their own cut-off point, which could vary from rat to rat.

This method of analysis also allows successive responses to be classed into one of four transition types: 'non-responding to non-responding' (nn), if two consecutive intervals of the cut-off length elapse without a response; 'non-responding to responding' (nr), if an absence of responding for the cut-off period is followed by a response; 'responding to non-responding' (rn), if a response is followed by a period of not responding greater than the cut-off; and 'responding to responding' (rr), if two responses are emitted within the cut-off value of each other. This allows the probability of the transitions from non-responding to be calculated (e.g., the probability of an 'nr' transition is the number of 'nr' responses divided by the number of 'nn' plus the number of 'nr' responses). This is the probability of a response-initiation response, and should be greatest on the RR-10 schedule and smallest on the RR-60 schedule. Similarly, it allows the probabilities of transitions from responding to responding to be calculated, and these should be greatest on the RR-60 schedule. A similar analysis of schedule performance has been suggested by Jenkins (1970) in order to more fully understand the nature of schedule-controlled performance, but has not been adopted in many previous studies of free-operant performance in any species.

Method

The study consisted of 35 volunteer participants (12 males and 23 females), aged between 18-23 years, and were recruited as described in Experiment 1. Five

participants were excluded on the basis of their BDI or O-LIFE scores, and 3 participants did not complete the study within 60min. This left 27 participants in total (9 in each group). Participants were tested as described in Experiment 1, except that each participant only experienced one schedule type, creating three groups: Group RR-10, RR-30, and RR-60. Training lasted until the participant had received 100 reinforcers.

Results and Discussion

Group-mean overall rates of response were calculated across the last 5 reinforcers to reflect terminal performance: RR-10 = 195.96 (± 79.46); RR-30 = 247.27 (± 208.68); and RR-60 = 145.07 (± 70.45). There was no significant linear trend, $F < 1$, *partial eta*^{2} = .012, *95% CI* = .009:.182, but a significant quadratic trend, $F(1,24) = 3.94$, $p < .05$, *partial eta*^{2} = .141, *95% CI* = .000:.381.

Figure 3 and Table 1 about here

The group-mean response-initiation rates over the final 5 reinforcers were calculated by using four different cut-off points (0.5s, 0.75s, and 1s), and then dividing the number of responses with an IRT of more than this cut-off by the summed time for all of those responses. These response-initiation rates are shown in the top of Figure 3. The response rate pattern for each cut-off is broadly similar and all differ from the overall rate pattern discussed above. Response-initiation rates were greatest for Group RR-10, followed by Group RR-30, and were lowest for Group RR-60 (80.31 ± 33.23). The top panel of Table 1 displays the linear (all significant) and quadratic (all non-significant) trend tests and effect sizes.

The within-burst response rates for the last 5 reinforcers were calculated by dividing the number of those responses by the summed IRTs for those responses. The grouo-mean data are shown in the bottom of Figure 3 for each cut-off value. Each cut-off value produced a similar pattern of results, and all mirrored those for the overall response rates (highest for Group RR-30). The results of the linear (all non-significant) and quadratic (all significant) trends, as well as effect sizes are shown in the lower part of Table 1. This pattern of results was also noted in the current Experiment 1, and has also been observed in some studies of nonhuman RR performance using different procedures (Reed, 2011).

Figure 4 and Table 2 about here

Figure 4 displays the group-mean probabilities for nr (initiation-responses) and rr (within-burst responses) transitions for the three groups, over the last 5 reinforcers, based on each of the cut-off criteria. Only nr and rr probabilities are shown, as the nn and rn probabilities are just the inverse of these former two values, respectively. All cut-off points produced a broadly similar pattern of results: nr transitions (top panel) were numerically more likely for Group RR-10, but there was little reliable difference in the rr probabilities (bottom panel) between the groups. Table 2 displays the results from the linear and quadratic trend tests conducted on these data as well as the effect sizes.

These results broadly replicated those from Experiment 1 and many features of previously established patterns of responding on RR schedule for nonhumans (see Reed, 2011; Reed & Hall, 1988). As with Experiment 1, initiation rates were higher with shorter RR values, and there were numerically greater probabilities of transition

to responding from non-responding with lower ratio schedules. These data again show human responding on RR schedules to be broadly similar to that of nonhumans, and the fact that these results emerged with a different procedure to that described in the current Experiment 1 show that these findings are relatively robust across different analytic techniques. The duration of the cut-off point used did not impact greatly the pattern of data noted, suggesting that, within the current limits, the selection of a cut-off point foes not greatly impact the outcomes.

Experiment 3

Experiments 1 and 2 examined the manner in which RR schedules control human performance during steady-state performance, which is the typical aspect of free-operant performance studied in nonhumans (see Brackney et al., 2011; Reed, 2011; Shull et al., 2001). In contrast, Experiment 3 examined the structure of human RR responding during initial exposure to the contingency. This aspect of schedule-controlled behavior has not been extensively studied (if at all) in any species, and so it would be of some interest to investigate the development of these response structures. Additionally, performance during initial schedule exposure was thought to be important to explore in humans in order to establish if these patterns were learned, and not present at the start of the study, which may be possible given the previously demonstrated impact of pre-formed verbal rules on human performance (see Leander et al., 1968; Lowe, 1979; Weiner, 1970). If the schedule contingencies are operating to shape behavior in humans, it might be expected that the typical patterns of RR performance described in the preceding two current studies would emerge over the course of training. To these ends, human behavior on three RR schedule values (RR-

10, RR-30, and RR-60) was studied to maintain compatibility with previous work, and the procedure adopted in Experiment 2 was employed to explore how human RR performance develops over early training. Given that the actual value of the cut-off was shown not to impact the results, a value of 1s was chosen for this study, as the focus was in acquisition were overall rates may be relatively low. As the intention was to adopt a reinforce-by-reinforcer approach to the analysis of behavior, the log-survivor technique was thought inappropriate due to the relatively few data points between each successive reinforcer making equation fitting problematic.

Method

Thirty participants (10 males and 20 females), aged 20-43 years, were recruited as described in Experiment 1 (these participants were different to those in Experiment 2). The participants were randomly and equally divided into the three schedule conditions. However, 3 participants were excluded for scoring above cut-off on the BDI, leaving 27 participants in total. The procedure was identical to that described in Experiment 2, except that the participants only received 8 reinforcers.

Results and Discussion

Figure 5 about here

Figure 5 presents group-mean overall responses rates across each successive reinforcer (number of responses emitted for the reinforcer divided by the time taken to emit these responses). The overall response rates were similar in the three schedules at the start of the experiment, but increased over the course of acquisition at different

rates across the three schedules. By the final (8th) reinforcer, Group RR-30 was responding faster than Group RR-10, with Group RR-60 displaying the slowest rate of response. A two-factor mixed-model ANOVA, with group (RR-10, RR-30, RR-60) as a between-subject factor, and reinforcer number as a within-subject factor, revealed no significant main effect of group, $F < 1$, but a statistically significant main effect of reinforcer, $F(7,168) = 16.88$, $p < .001$, *partial eta*^{2} = .413, *95% CI* = .279:.486, and a significant interaction between the two factors, $F(14,168) = 2.21$, $p < .01$, *partial eta*^{2} = .156, *95% CI* = .111:.189. Simple effect analyses revealed no significant simple effect of group on Reinforcers 1 to 5, but significant differences between the groups on Reinforcers 6-8, smallest $F(1,168) = 2.39$, $p < .05$. Tukey's HSD tests, conducted separately on response rates for Reinforcers 6-8, revealed the same pattern of results on all trials; the RR-30 schedule had a higher response rate than the RR-10 schedule, which, in turn, had a higher rate than the RR-60 schedule, all *p*s < .05.

Figure 6 about here

The top of Figure 6 reveals the group-mean response-initiation rates. All groups started with approximately equal response-initiation rates as one another. These rates increased over successive reinforcers for Group RR-10, and, to a lesser extent, for Group RR-30, but not for Group RR-60. A two-factor mixed-model ANOVA (group x reinforcer) revealed significant main effects of group, $F(2,24) = 8.78$, $p < .001$, *partial eta*^{2} = .421, *95% CI* = .092:.598, and reinforcer, $F(7,168) = 5.25$, $p < .001$, *partial eta*^{2} = .180, *95% CI* = .059:.251, and a significant interaction between the two factors, $F(14,168) = 3.05$, $p < .001$, *partial eta*^{2} = .203, *95% CI* = .046:.245. Simple effect analyses revealed no simple effect of group on Reinforcer 1,

but significant group differences on Reinforcers 2-8, smallest $F(1,168) = 6.57$, $p < .001$. Tukey's HSD tests revealed the same pattern of results for Reinforcers 2 to 8: Group RR-10 schedule had a higher response-initiation rate than Group RR-30, which, in turn, had a higher response-initiation rate than Group RR-60, all *p*s < .05.

The bottom of Figure 6 displays the group-mean within-burst response rates for the three groups across successive reinforcers. These rates increased slightly with training for all groups, but the groups did not differ greatly from one another. A two-factor mixed-model ANOVA (group x reinforcer) revealed a significant main effect of reinforcer, $F(7,168) = 3.88$, $p < .001$, *partial eta*^{2} = .139, *95% CI* = .029:.205, but neither the main effect of group, nor the interaction, were significant, both *p*s > .10. That these data did not follow the patterns seen in the previous two experiments reported here suggests that this aspect of schedule performance takes some time to develop.

Figure 7 about here

Figure 7 shows the group-mean probabilities for different types of response transitions for each successive reinforcer. Initially, nr transitions (response initiations) were not likely to occur, but this probability increased with each successive reinforcer. This pattern was more pronounced for Groups RR-10 and RR-30, than for Group RR-60. A two-factor ANOVA (group x reinforce) revealed no significant main effect of group, $F(2,24) = 1.65$, $p > .20$, but a significant main effect of reinforcer, $F(7,168) = 10.91$, $p < .001$, *partial eta*^{2} = .610, *95% CI* = .225:.741, and a significant interaction between the factors, $F(14,168) = 2.33$, $p < .01$, *partial eta*^{2} = .162, *95% CI* = .016:.198. Simple effect analysis revealed no significant simple effect

of group on Reinforcers 1 to 6, inclusive, all *p*s > 0.20, but there were significant group differences on Reinforcers 7 and 8, smallest $F(2,24) = 3.91$, $p < .05$. Tukey's HSD tests conducted on Reinforcers 7 and 8 both revealed that Group RR-10 differed from Group RR-60, $p < .05$, but no other differences were significant.

Inspection of Figure 7 also shows that rr transitions (within-burst responding) became increasingly more likely with each reinforcer, and by the end of training this was more pronounced for Groups RR-30 and RR-60 than for Group RR-10. However, a two-factor ANOVA (group x reinforcer) revealed that while there was a significant main effect of reinforcer, $F(7,168) = 9.61$, $p < .001$, *partial eta*^{2} = .286, *95% CI* = .151:.363, neither the main effect of group nor the interaction were significant, *F*s < 1. These data mirror those for the rates discussed above, and suggest that probability of transitions do change with training for humans, and response-initiation responses are more likely on RR-10 schedules than for the other schedules.

These results replicated for acquisition some of the steady-state patterns of RR responding previously observed in nonhumans, and also in humans in the current Experiments 1 and 2. By the end of the experiment, the predicted pattern of overall response rate was noted, with response rates being higher at intermediate ratio values (at least, over this range of values; see Reed, 2011; Reed and Hall, 1988). Analysis of the structure of responding demonstrated that initiation rates were higher on the RR-10 than on the RR-60 schedule, with greater probabilities of transition to responding from non-responding on the shorter schedule. These findings were similar to those noted in the preceding two steady-state studies for humans, and also to nonhuman steady-state responding (see Reed, 2011). However, some of the findings noted in the previous studies relating to within-burst patterns of responding were not noted in acquisition (i.e., the rate of within-burst responding). This suggests that these aspects

of behavior take some time to establish, and that the human participants did not bring preformed rules of responding to the study.

Experiment 4

The forms of analyses adopted for the preceding three experiments allow investigation of many aspects of human RR performance. However, there are a number of problems involving the identification of response-initiation or within-burst responses based on either post hoc or a priori criteria that might not reflect the actual performance generated (see Shull, 2011). Although the results from the preceding studies show convergence with one another, and with previous studies of nonhuman performance, the final experiment reports the data from a study designed to remedy some of these problems (see Reed, 2011). This approach involves an experimental procedure based on that suggested by Mechner (1992; see also Pear and Rector, 1979). Two discrete manipulanda are provided to the participant, a response to one manipulandum marks the start of a response, which is then conducted on another manipulandum. This clearly demarks the 'burst-initiation', from 'within-burst', responses (see Reed, 2011).

To this end, human participants were presented with two squares displayed on a computer screen. At the start of a session, one of the squares (the 'burst-initiation' square) was filled with a color. A response (mouse click) to this square was taken to mark the start of a response, and it extinguished the color in that square. The second square was then filled with a color, and that square was operative for the schedule to be completed. If the participants ceased responding for a period of time prior to the schedule being completed, this was taken as the burst being terminated (Mellgren and

Elsmore, 1991; Reed, 2011). The color in the response-burst lever was extinguished, the square ceased to be operative, and the response-initiation square was filled with its color again. This way, the numbers of response-initiating, and within-burst responses, can be clearly demarked from one another. Of course, the length of time without a response that determines the end of a burst is an arbitrary (albeit pre-experimentally defined) criterion. A value of 1s was chosen for all of this experiment based on previous data reported above in the current studies.

Using this procedure, it was hoped to establish whether the observations made from the post hoc analysis of IRT data by Shull et al. (2001; see also Shull and Grimes, 2003), and noted above in Experiments 1, 2 , and 3, would be corroborated. To this end, three groups were studied, each group responding on a different RR schedule of reinforcement: RR-10, RR-30, and RR-60. The data were analyzed to establish overall rates and response structures on these schedules.

Method

Participants

Thirty participants (10 male and 20 female), aged 19-27 years were recruited as described in Experiment 1. Six participants were excluded from the study due to high BDI or O-LIFE(B) scores, and two did not complete the experimental task. This left each group with 8 participants.

Apparatus

The apparatus was as described in Experiment 1, except that, instead of one square being displayed on the screen, there were two squares displayed on the screen. These squares measured 2cm x 2cm, and were located at the bottom of the screen –

one square a third of the way in from the left, and the other square a third of the way in from the right. The squares had a black outline, but were not otherwise filled (the background screen color being visible through the squares). The squares could be illuminated in either green or yellow. Each square was always either green or yellow for a particular participant, but which square received which color was random across participants. The response-initiation square was always the square on the left of the screen.

Procedure

The participants were tested individually as described in Experiment 1. After completing the psychometric tests, they were presented with the following instructions on the computer screen:

"*This experiment is concerned with multi-tasking abilities. You have two tasks to complete simultaneously. In the first task you must count backwards, out-loud, in 7s from 26,758. The second task is to score as many points as possible by clicking in the two coloured boxes displayed in the computer screen. You may need to click on the boxes in a particular order to score points. To receive points, sometimes you might need to press the spacebar quickly and at other times you might need to press slowly. The person with the best score on both tasks will receive a £50 [name of company] token."*

The participants then responded for 20min on the computer task. In these sessions, initially the left square was filled with a color, and the right square was not filled. A response (mouse click) to the left square extinguished the color for that square, and filled the right square with its color. While the left square was filled, a response to the right square had no programmed consequences. After the right square

was filled, responses to that square could be reinforced according to a particular schedule: RR-10, RR-30, or RR-60 (in all cases, the probability of each response being reinforced was set at 1/RR value). Failure to respond to the right square for 1s while the square was filled with color resulted in the extinction of the color in that square, and responses to that square having no programmed consequences. The left square would then be illuminated, and the sequence started again.

Results and Discussion

The data from the last 10 reinforcers, taken to reflect steady-state performance, were analyzed. The mean obtained numbers of reinforcers per min for the three groups were: RR-10 = 7.19 (± .62); RR-30 = 3.81 (± .49); RR-60 = 1.67 (± .28). An ANOVA conducted on these data revealed a significant effect of group, $F(2,21) = 261.30$, $p < .001$, *partial eta*^{2} = .961, *95% CI* = .911:.974. Tukey's HSD tests revealed that the pairwise comparisons between all groups were statistically significant, *p*s < .01.

Figure 8 about here

Figure 8 (solid bars) displays the group-mean overall responses per min (all responses to illuminated squares for the total session) for the three groups. Group RR-30 produced the highest rate of response, followed by the RR-60 group, and both of these groups had a higher rate of response than the RR-10 group. An ANOVA that revealed a significant effect of group, $F(2,21) = 20.46$, $p < .001$, *partial eta*^{2} = .661, *95% CI* = .337:.773. Tukey's HSD tests revealed that the difference between the RR-10 and RR-30, and the RR-10 and RR-60, groups were significant, *p*s < .05.

Figure 8 also shows (diagonal-lined bars) the mean 'burst-initiation' rates (responses to the response-initiation square per min after the time spent responding in the bursts and the 1s periods without a response were removed from the session time). These data show a different pattern to the overall response rates, with the RR-10 schedule having the highest rate of response-initiation responses, followed by the RR-30, and then the RR-60 schedule. An ANOVA revealed a statistically significant effect of group, $F(2,21) = 121.89$, $p < .001$, *partial eta*^{2} = .921, *95% CI* = .820:.947. Tukey's HSD tests revealed that all pairwise comparisons were statistically significant, all *p*s < .05.

The group-mean rates of within-burst responding for the three groups (from the first response to the 'burst' square to the last, and excluding the 1s time of not responding at the end of the bursts) were: RR-10 = 148.97 (± 36.80); RR-30 = 167.14 (± 33.85); RR-60 = 115.59 (± 22.58). An ANOVA revealed a statistically significant effect of group, $F(2,21) = 5.45$, $p < .05$, *partial eta*^{2} = .342, *95% CI* = .021:.544, and Tukey's HSD tests revealed that both the RR-10 and the RR-30 schedule each differed from the RR-60 schedule. These 'within-burst' rates of responding conformed to the rates of overall responding that might be expected in typical single-manipulandum schedules for these ratio values (see Reed and Hall, 1988).

In summary, the present experiment replicated the key features of human RR performance seen in the preceding steady-state studies presented in the current series: overall rates (and within-burst rates) demonstrated an inverted-U relationship to ratio size, but response-initiation rates demonstrated an inverse relationship to RR size (which might be related to the rate of reinforcement experienced on that schedule). In addition, these results are consistent with the analysis of the effects on 'burst-

initiating' responses, and 'within-burst' responses produced by reports of nonhuman responding (see Reed, 2011).

General Discussion

The current studies explored the structure of human responding maintained by RR schedules, both in steady-state using a variety of procedures and analytic techniques (Experiments 1, 2, and 4), and during acquisition (Experiment 3). The main aim of these studies was to produce new empirical knowledge regarding this response micro-structure. However, it was also hoped to be able to separate the factors that control any emerging different types of responses, and to relate this information to a number of uninvestigated areas such as the influence of language on human responding, and the potential role of these response types in determining the outcomes of human causality judgment experiments.

All of the current studies, whether using a single manipulandum (Experiments 1 to 3) or a dual manipulandum (Experiment 4) found overall responding to be higher on the intermediate ratio values employed compared to smaller or larger values(at least across the values currently studied), which is also found for nonhumans on such schedules (e.g., Reed and Hall, 1988). This consistency across the current experiments suggests that the effects are relatively independent of the types of experimental procedure and analysis used to explore them. However, performance of nonhumans on free-operant schedules has been found to comprise both response-initiation and within-burst components, and these responses have been noted to be controlled differentially by various aspects of the schedule (see Brackney et al., 2011; Pear & Rector, 1979; Reed, 2011; Shull et al., 2011). The current studies

also found evidence for the existence of two separate components of responding in humans: responses that initiated the response-bout; and those that formed part of the response burst. These effects have not previously been explored for humans.

As with nonhumans in previous reports, these components of human schedule behavior appeared to be controlled differentially by different aspects of the contingency to one another. Irrespective of the overall rates of response, response response-initiation responses were more frequent on lower ratio values (e.g., RR-10) than on higher ratio values (e.g., RR-60). This mirrors some findings from nonhuman studies that used different procedures to those adopted here (e.g., Bower et al., 2008; Reed, 2011), and is also consistent with the view that response-initiation responses are sensitive to overall rates of reinforcement (see Shull, 2011). In contrast, within-burst responses in humans appear not to be sensitive to reinforcement rate; the current experiments tending to find that the rate of within-burst responding was similar to the overall rates of response. That these results were quite consistent across the studies, and also were found to emerge with training rather than being present from the start of schedule exposure (Experiment 3), suggests that there are multiple elements to human schedule performance that are very similar to those found in nonhumans.

The current findings may have some implications for the further study of human schedule performance in two areas: the impact of language, and the impact of schedules on causal judgments. The sensitivity of human performance to reinforcement rate, which has so far been difficult to establish (see Lowe, 1979), may be better seen in a study of response-initiation rates. The ability to distinguish between response-initiation and within-burst behaviors may also allow further investigation of the impact of a range of factors that alter human performance. For example, it has been suggested that verbal rules impact on human schedule

performance (see Bradshaw and Reed, 2012; Hayes et al., 1986), but it is unclear exactly how this impact is mediated. By allowing investigation of whether such rules impact more on the response-initiation or within-burst responding, it may be possible to parse the nature of this linguistic control. Shull et al. (2001) have suggested that the response-initiation responses may be more sensitive to molar schedule properties, and within-burst responses to molecular factors, and it may also be that the former are more sensitive to explicit linguistic control than the latter.

The potential presence of two distinct components in human schedule performance may also help to re-assess human probability learning and causal judgment studies – and might help to explain some deviations from a simple outcome probability rule in predicting judgments of causality (see Dack et al., 2009; Reed, 1999; 2001). Experiments conducted in these areas have assumed that probabilities of outcomes given a response can be relatively easily calculated. This may still be the case, as the procedures used here are not identical to those employed in these other contexts (see Reed, 1999; 2001, for discussion of the difference between schedules of reinforcement and schedules of outcomes). However, it may also be that many findings in the causal and probability judgment literature that are not easily explained by simple associative rules may be accommodated if response-initiation and within-burst responses are considered as separate components, and not related in the same way as one another to the outcomes. For example, Reed (2001; see also Dack et al., 2009) has noted that response emitted on ratio schedules are rated as more causally effective than those emitted on interval schedules, even when the probabilities of an outcome given a response are equated. It may be that, as the behavior on RR schedules has fewer response-initiation responses than those on an RI schedule (Shull et al., 2001), the experienced probabilities of an outcome given a response are not

equivalent on the two schedules, as the participants are not treating all responses as equivalent to one another.

Leaving aside these theoretical implications, the current data suggest strong similarities between human and nonhuman performance on RR schedules of reinforcement: a similarity found at both the overall and the structural levels of responding. Such similarities, at least in overall performance, are not novel (e.g., Bradshaw and Reed, 2013; Randell et al., 2009; Raia et al., 2000), but are not always noted (see Lowe, 1979). It should be noted that the current procedure adopted a number of modifications suggested by the previous literature in order to achieve this similarity. For example, response costs were introduced, as well as the use of reinforcers that translated into tangible outcomes (prizes). Without these modifications it has been shown to be difficult to obtain strong schedule control over human behavior (see Bradshaw and Reed, 2012; Raia et al., 2000; Reed, 2001). In addition, the current study attempted to suppress engagement in the formation of verbal rules that can interfere with human schedule performance (see Bradshaw and Reed, 2013; Hayes et al., 1986). It should be pointed out that it is unclear whether this modification had an effect, as no control group was used in which it was not employed (as this was not the purpose of the study). Nor is it clear whether the verbal instructions employed in these studies (which are typical of the field), impacted on the participants' performance. Nevertheless, the presence of this 'language-suppressing' procedure in the current studies should be noted. Finally, the current studies excluded participants who demonstrated high psychometrically-defied levels of depression, or schizotypy, as both of these traits have been shown to interfere with obtaining typical nonhuman schedule performance in humans (see Randell et al., 2009). It should also

be noted that the wide range of effect sizes in some of the studies might indicate a relatively low power for some analyses.

In summary, these studies demonstrate that human schedule performance on RR schedules is highly similar to nonhumans at both the overall and structural level. The existence of putatively different components may open a number of lines of investigation onto human performance that may prove fruitful in illuminating the controlling factors for human operant behavior.

References

Alloy, L.B., & Abramson, L.Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser?. *Journal of Experimental Psychology: General*, **108**, 441.

Andersson, G., Hagman, J., Talianzadeh, R., Svedberg, A., & Larsen, H. C. (2002). Effect of cognitive load on postural control. *Brain Research Bulletin*, **58**, 135-139.

Baron, A. & Galizio, M. (1983). Instructional control of human operant behaviour. *Psychological Record*, **33**, 495-520.

Baum, W.M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, **20**, 137-153.

Beck, A.T., Steer, R.A., & Garbin, M.G. (1988). Psychometric properties of the Beck Depression Inventory: Twenty-five years of evaluation. *Clinical Psychology Review*, **8**, 77-100.

Beck, A. T., Ward, C. H., Mendelson, M., Mock, J., & Erbaugh, J. (1961). An inventory for measuring depression. *Archives of General Psychiatry*, **4**, 561–571.

Bowers, M.T., Hill, J., & Palya, W.L. (2008). Interresponse time structures in variable-ratio and variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, **90**, 345-62.

Brackney, R.J., Cheung, T.H., Neisewander, J.L., & Sanabria, F. (2011). The isolation of motivational, motoric, and schedule effects on operant performance: a modeling approach. *Journal of the Experimental Analysis of Behavior*, **96**, 17-38.

Bradshaw, C. A. & Reed, P. (2012). Relationship between contingency awareness and

human performance on random ratio and random interval schedules. *Learning and Motivation,* **43**, 55-65.

Dack, C., McHugh, L., & Reed, P. (2009). Generalization of causal efficacy judgments after evaluative learning. *Learning & Behavior*, **37**, 336-348.

Estes, W.K. (1976). The cognitive side of probability learning. *Psychological Review,* **83**, 37.

Ferster, C.B. & Skinner, B.F. (1957). *Schedules of Reinforcement*. New York: Appleton-Century-Crofts.

Hayes, S.C., Brownstein, A.J., Zettle, R.D., Rosenfarb, I. & Korn, Z. (1986). Rule-governed behavior and sensitivity to changing consequences of responding. *Journal of the Experimental Analysis of Behavior,* **45,** 237-256.

Jenkins, H.M. (1970). Sequential organization in schedules of reinforcement. *The theory of reinforcement schedules. New York: Appleton-Century-Crofts,* 63-109.

Kessel, R., & Lucke, R. (2008). An analytic form for the response rate analysis of Shull, Gaynor, and Grimes with applications and extensions. *Journal of the Experimental Analysis of Behavior*, **90**, 363–386.

Killeen, P.R., Hall, S.S., Reilly, M.P., & Kettle, L.C. (2002). Molecular analyses of the principal components of response strength. *Journal of the Experimental Analysis of Behavior*, **78**, 127-160.

Kulubekova, S., & McDowell, J.J. (2008). A computational model of selection by consequences: Log survivor plots. *Behavioural Processes.* **78**, 291–296.

Leander, J.D., Lippman, L.G. & Meyer, M.M. (1968). Fixed interval performance

as related to subjects' verbalizations of the reinforcement contingency. *The Psychological Record,* **18**, 469-74.

Lowe, C.F. (1979). Determinants of human operant behavior. In M. D. Zeiler & P. Harzem (Eds.). *Advances in Analysis of Behaviour: Vol. 1. Reinforcement and the Organisation of Behaviour* (pp. 71-93). Chichester, UK: Wiley.

Lowe, C.F., Beasty, A., & Bentall, R.P. (1983). The role of verbal behavior in human learning: Infant performance on fixed-interval schedules. *Journal of the Experimental Analysis of Behavior*, **39**, 157-164.

Lowe, C.F., Harzem, P. & Bagshaw, M. (1978). Species differences in temporal control of behavior. II: Human performance. *Journal of the Experimental Analysis of Behavior,* **29**, 351-61.

Mason, O., Linney, Y. & Claridge, G. (2005). Short scales for measuring schizotypy. *Schizophrenia Research*, **78**, 293-296.

Matthews, B. A., Shimoff, E., Catania, A. C. & Sagvolden, T. (1977). Uninstructed human responding: Sensitivity to ratio and interval contingencies. *Journal of the Experimental Analysis of Behavior*, **27**, 453-67.

McDowell, J.J., & Wixted, J.T. (1986). Variable-ratio schedules as variable-interval schedules with linear feedback loops. *Journal of the Experimental Analysis of Behavior*, **46**, 315-329.

Mechner, F. (1992). *The revealed operant: A way to study the characteristics of individual occurrences of operant responses*. Cambridge, MA: Cambridge Center for Behavioral Studies.

Mellgren, R.L., & Elsmore, T.F. (1991). Extinction of operant behavior: An analysis based of foraging considerations. *Animal Learning & Behavior*, **19**, 317–325

Morse W.H. (1966). Intermittent reinforcement. In W.K. Honig (Ed.) *Operant*

behavior: Areas of research and application (52-109). New York: Appleton Century-Crofts; 1966. pp. 52–109.

Pear, J.J., & Rector, B.L. (1979). Constituents of response rate. *Journal of the Experimental Analysis of Behavior*, **32**, 341-362.

Peele, D.B., Casey, J. & Silverberg, A. (1984). Primacy of interresponse-time reinforcement accounting for rates under variable-ratio and variable-interval schedules. *Journal of Experimental Psychology: Animal Behavior Processes,* **10**, 149-67.

Raia, C.P., Shillingford, S.A., Miller, H.L. & Baier, P.S. (2000). Interaction in procedural factors in human performance on yoked schedules. *Journal of the Experimental Analysis of Behavior,* **74**, 265-81.

Randell, J., Ranjith-Kumar, A. C., Gupta, P., & Reed, P. (2009). Effect of schizotypy on responding maintained by free-operant schedules of reinforcement. *Personality and Individual Differences*, **47**, 783-788.

Reed, P. (1999). Effect of perceived cost on judgments regarding the efficacy of investment. *Journal of Economic Psychology*, **20,** 657-76.

Reed, P. (2001). Schedules of reinforcement as determinants of human causality judgments and response rates. *Journal of Experimental Psychology: Animal Behavior Processes*, **27**, 187-95.

Reed, P. (2007). Response rate and sensitivity to the molar feedback function relating response and reinforcement rate on VI+ schedules of reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes,* **33**, 428.

Reed, P. (2011). An experimental analysis of steady-state response rate components on variable ratio and variable interval schedules of reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, **37**, 1-9.

Reed, P. (2015). Rats show molar sensitivity to different aspects of random-interval-with-linear-feedback-functions and random-ratio schedules. *Journal of Experimental Psychology: Animal Learning and Cognition*.

Reed, P., Frasquillo, F., Colkin, C., Liemann, V., & Colbert, S. (2001). Interference with judgements of control and learning as a result of prior exposure to controllable and uncontrollable feedback during concept-learning tasks. *The Quarterly Journal of Experimental Psychology*, **54**B, 167-183.

Reed, P., & Hall, G. (1988). The schedule dependency of the signalled reinforcement effect. *Learning and Motivation*, **19**, 387-407.

Shull, R.L. (2004). Bouts of responding on variable interval schedules: Effects of deprivation level. *Journal of the Experimental Analysis of Behavior*, **81**, 155–167.

Shull, R.L. (2011). Bouts, changeovers, and units of operant behavior. *European Journal of Behavior Analysis*, **12**, 49-72.

Shull, R.L., Gaynor, S.T., & Grimes, J.A. (2001). Response rate viewed as engagement bouts: Effects of relative reinforcement and schedule type. *Journal of the Experimental Analysis of Behavior*, **75**, 247–274.

Shull, R.L., & Grimes, J.A. (2003). Bouts of responding from variable-interval reinforcement of lever pressing by rats. *Journal of the Experimental Analysis of Behavior*, **80**, 159-171.

Shull, R.L., Grimes, J.A., & Bennett, J.A. (2004). Bouts of responding: The relation between bout rate and the rate of variable-interval reinforcement. *Journal of the Experimental Analysis of Behavior.* **81**, 65-83.

Sibley, R.H., Nott, H.M.R. & Fletcher, D.J. (1990). Splitting behaviour into bouts. *Animal Behaviour*, **39**, *63-69.*

Skinner, B.F. (1969). *Contingencies of reinforcement*. New York: Appleton Century Crofts.

Tolkamp, B.J., & Kyriazakis, I. (1999). To split behaviour into bouts, log-transform the intervals. *Animal Behaviour*, **57**, 807-817.

Wasserman, E.A., Chatlosh, D.L., & Neunaber, D.J. (1983). Perception of causal relations in humans: Factors affecting judgments of response-outcome contingencies under free-operant procedures. *Learning and Motivation*, **14**, 406-432.

Weiner, H. (1970). Human behavioural persistence. *Psychological Record*, **20**, 445-56.

Zeiler, M.D. (1984). The sleeping giant: Reinforcement schedules. *Journal of the Experimental Analysis of Behavior*, **42**, 485-493.

Zuriff, G.E. (1970). A comparison of variable-ratio and variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, **13**, 369-74.

Acknowledgements

Grateful thanks are due to Richard Shull for his generous help with the analyses of these data, to Gary Freegard for his technical support, and to Lisa A. Osborne for her support and comments. Correspondence regarding this article should be sent to: Phil Reed, Department of Psychology, Swansea University, Singleton Park, Swansea, SA2 8PP, UK (e-mail: p.reed@swansea.ac.uk).

Table 1: Experiment 2. Statistical results for the analysis of linear and quadratic trend son response-initiation and within-burst rates of responding (see Figure 3). Top panel = response-initiation rates. Bottom panel = within-burst rates.

Cut-off	Analysis	F(1,24)	Partial eta^2	95% CI
Individual	Linear	4.82*	.170	.000 - .410
	Quadratic	1.07	.043	.000 - .258
0.5s	Linear	4.39*	.155	.000 - .395
	Quadratic	< 1	.029	.000 - .231
0.75s	Linear	3.79*	.136	.000 - .377
	Quadratic	< 1	.028	.000 - .231
1s	Linear	4.73*	.165	.000 - .405
	Quadratic	< 1	.030	.000 - .235
Cut-off	Analysis	F(1,24)	Partial eta^2	95% CI
Individual	Linear	< 1	.016	.000 - .202
	Quadratic	4.50*	.158	.000 - .399
0.5s	Linear	< 1	.005	.000 - .162
	Quadratic	4.15*	.147	.000 - .388
0.75s	Linear	< 1	.027	.000 - .228
	Quadratic	4.69*	.164	.000 - .404
1s	Linear	< 1	.030	.000 - .236
	Quadratic	4.20*	.149	.000 - .390

*p < 0.05; **p < 0.01

Table 2: Experiment 2. Statistical results for the analysis for linear and quadratic trends on probabilities of N-R (response-initiation) and R-R (within-burst) transitions (see Figure 4). Top panel = N-R probabilities. R-R probabilities.

Cut-off	Analysis	F(1,24)	Partial eta^2	95% CI
Individual	Linear	4.89*	.169	.000 - .410
	Quadratic	< 1	.014	.000 - .197
0.5s	Linear	3.49	.127	.000 - .367
	Quadratic	< 1	.032	.000 - .278
0.75s	Linear	4.54*	.159	.000 - .400
	Quadratic	< 1	.013	.000 - .194
1s	Linear	4.23*	.150	.000 - .390
	Quadratic	< 1	.018	.000 - .208

Cut-off	Analysis	F(1,24)	Partial eta^2	95% CI
Individual	Linear	< 1	.013	.000 - .195
	Quadratic	< 1	.006	.000 - .169
0.5s	Linear	< 1	.026	.000 - .227
	Quadratic	1.29	.051	.000 - .271
0.75s	Linear	1.12	.045	.000 - .261
	Quadratic	< 1	.036	.000 - .246
1s	Linear	2.76	.103	.000 - .340
	Quadratic	< 1	.020	.000 - .214

Figure 1. Top panel: Idealized log survivor plot. Bottom panel = Experiment 1. Group-mean log survival plots for the three schedules.

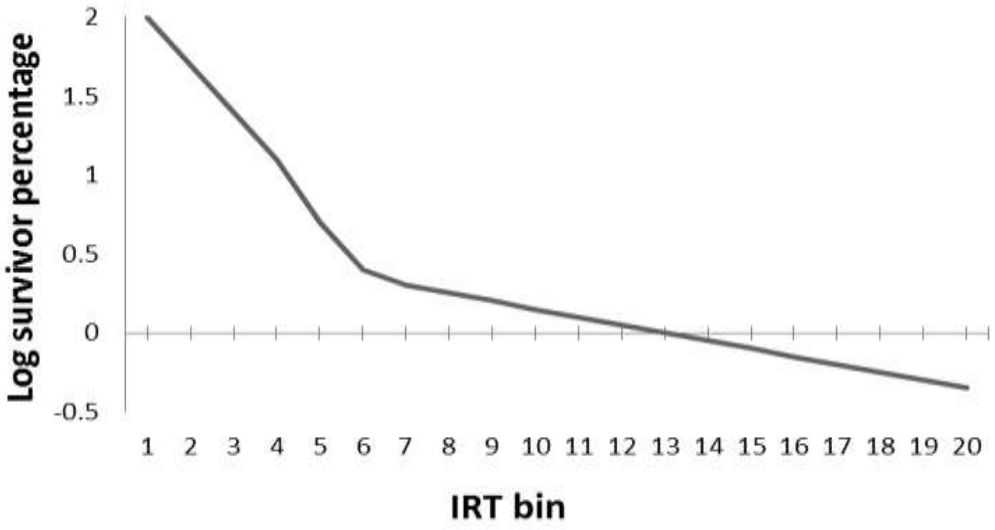

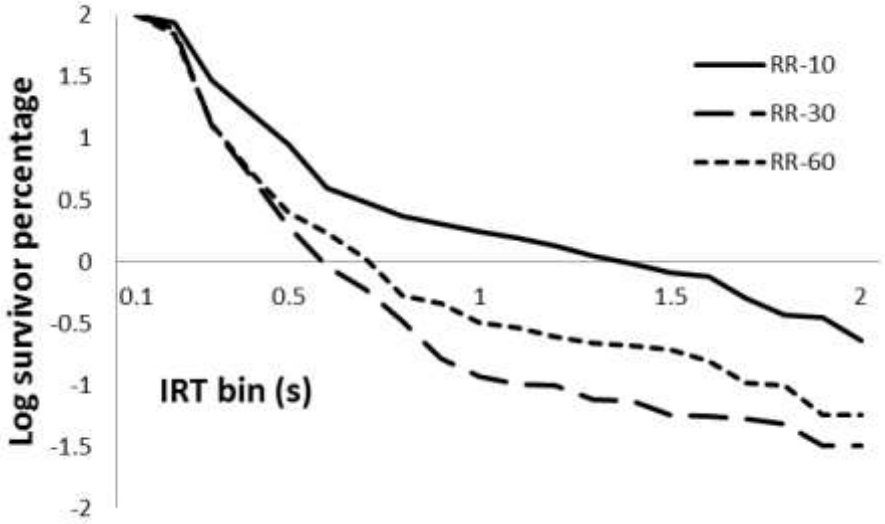

Figure 2. Experiment 1. Solid bars = mean response-initiation rates over the last 5min of exposure to each schedule. Diagonal bars = mean within-burst rates of response. Error bars = 95% confidence intervals.

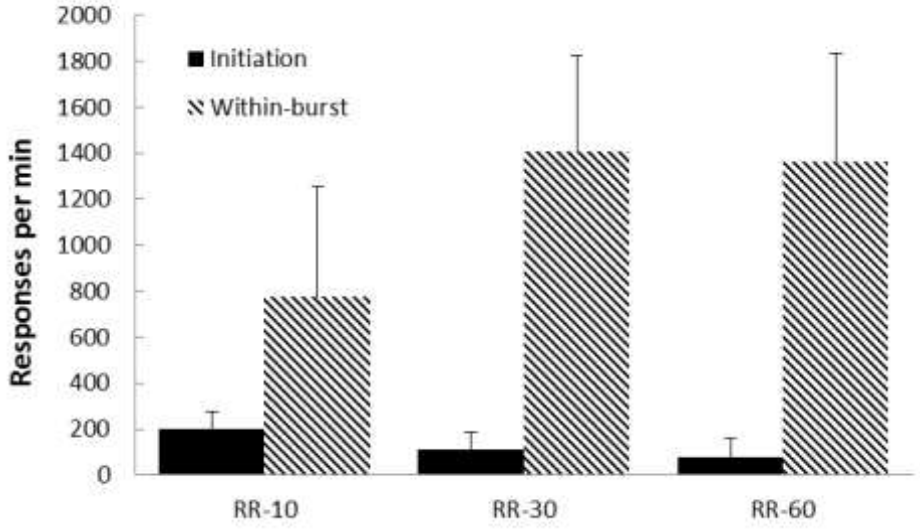

Figure 3. Experiment 2. Top panel = group-mean response-initiation rates for the three groups, averaged across the last 5 reinforcers. Bottom panel = group-mean within-burst rates for the three groups, averaged across the last 5 reinforcers. Error bars = 95% confidence intervals.

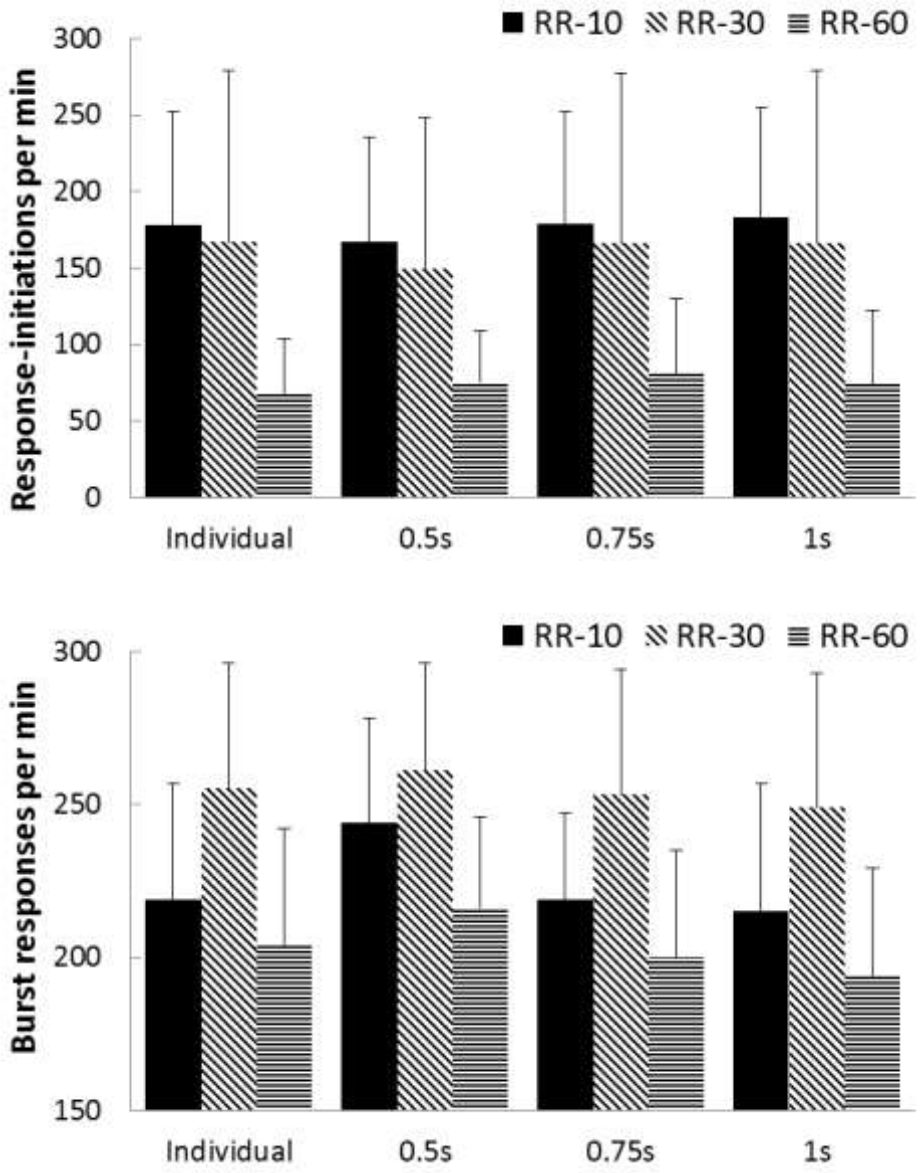

Figure 4. Experiment 2. Top panel = group-mean probabilities for nr (initiation-responses) transitions for the three groups over the last 5 reinforcers of training. Bottom panel = group-mean probabilities for rr (within-burst responses) transitions for the three groups over the last 5 reinforcers of training. Error bars = 95% confidence intervals.

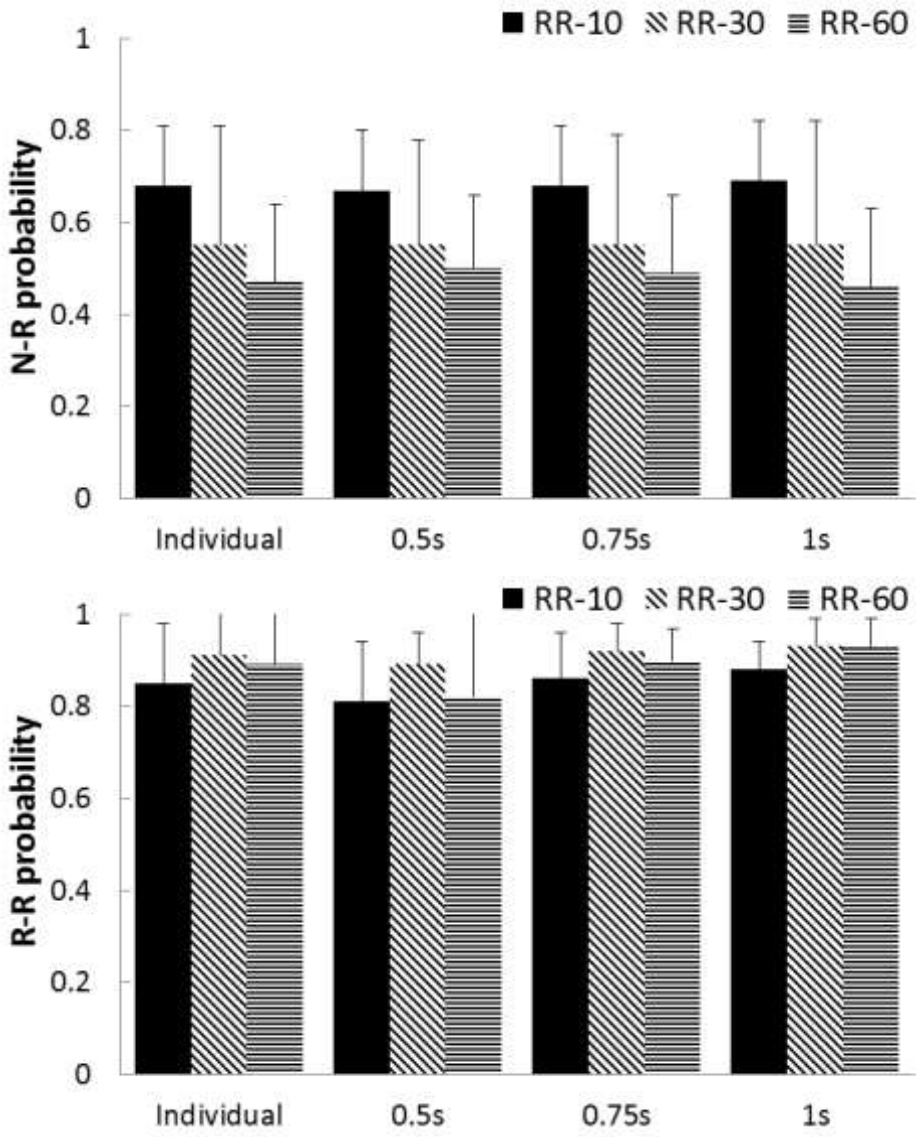

Figure 5. Experiment 3. Group-mean overall responses rates across each successive reinforcer.

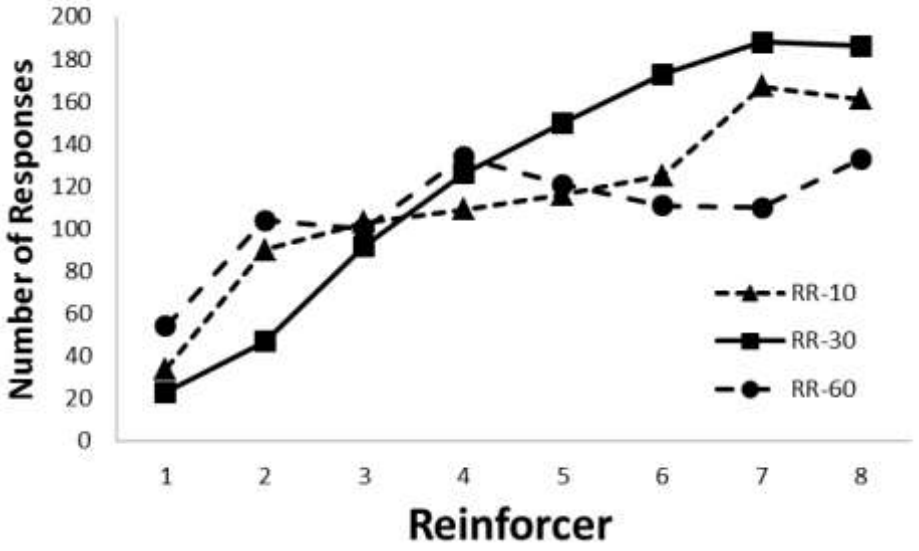

Figure 6. Experiment 3. Top panel = group-mean response-initiation rates across each successive reinforcer. Bottom panel = group-mean within-burst response rates for the three groups across successive reinforcers.

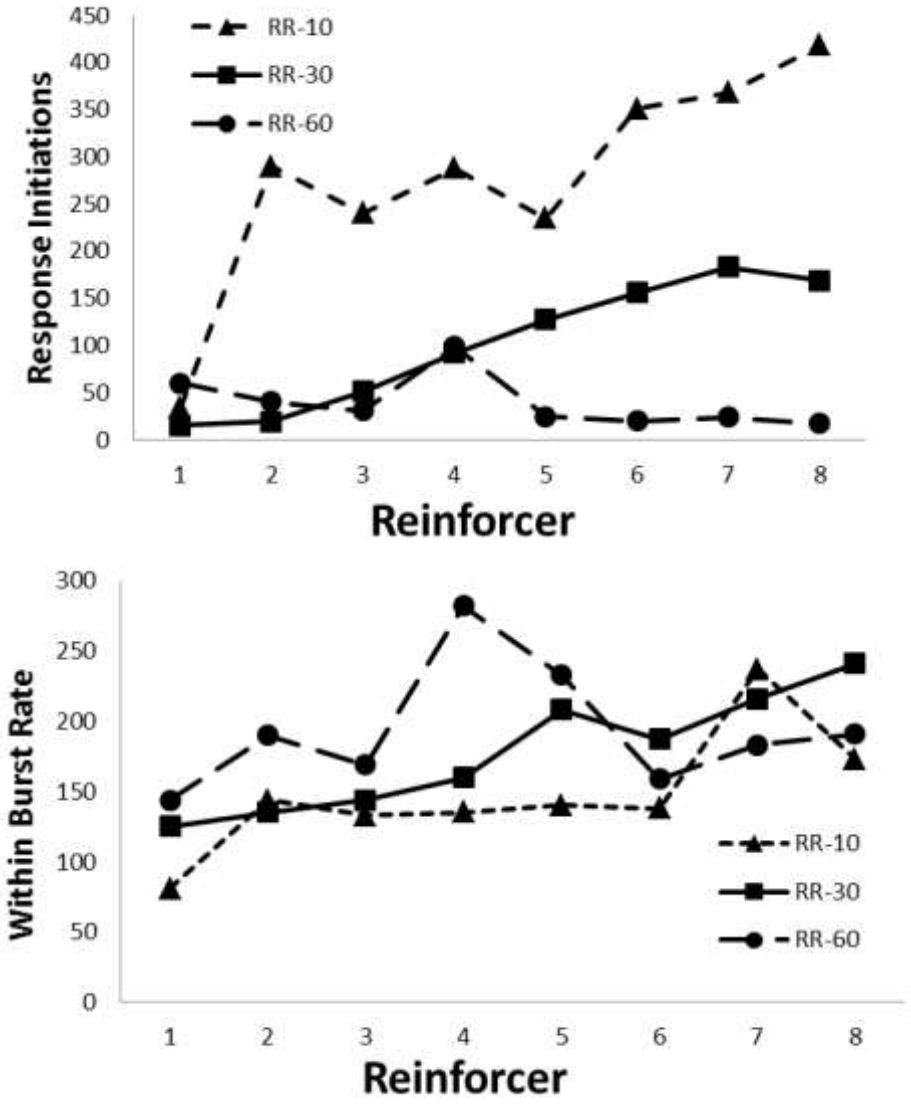

Figure 7. Experiment 3. Group mean probabilities for nr (initiation-responses) and rr (within-burst responses) transitions for the three groups over successive reinforcers.

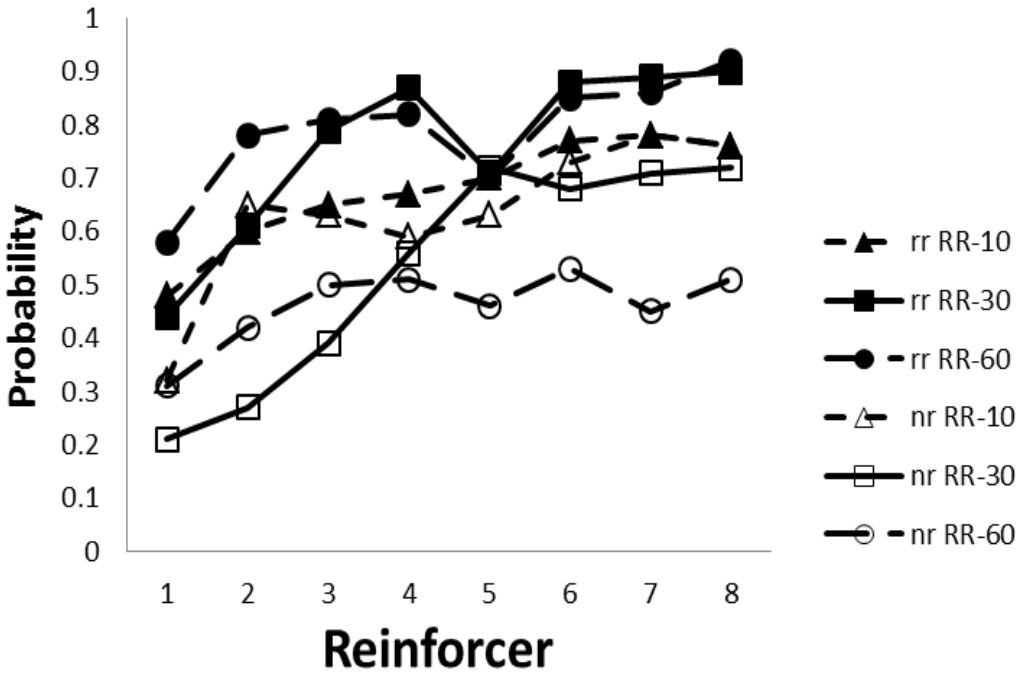

Figure 8. Experiment 4. Solid bars = group-mean overall responses per min for the three groups. Diagonal bars = group-mean response-initiation for the three groups. Error bars = 95% confidence intervals.