



City Research Online

City, University of London Institutional Repository

Citation: Jennings, D., Alonso, E., Mondragon, E., Franssen, M. & Bonardi, C. (2013). The Effect of Stimulus Distribution Form on the Acquisition and Rate of Conditioned Responding: Implications for Theory. *Journal of Experimental Psychology: Animal Behavior Processes*, 39(3), pp. 233-248. doi: 10.1037/a0032151

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/5186/>

Link to published version: <https://doi.org/10.1037/a0032151>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

The effect of stimulus distribution form on the acquisition and rate of conditioned responding: Implications for theory.

Dómhnall J. Jennings¹, Eduardo Alonso², Esther Mondragón³, Mathijs Franssen⁴ and Charlotte Bonardi⁴

¹Institute of Neuroscience, Newcastle University

²City University

³Centre for Computational and Animal Learning Research

⁴School of Psychology, University of Nottingham

Correspondence:

Charlotte Bonardi, School of Psychology, University of Nottingham, University Park, Nottingham, NG7 2RD, email: charlotte.bonardi@nottingham.ac.uk

Dómhnall Jennings, Institute of Neuroscience, Newcastle University, Henry Wellcome Building, Framlington Place, Newcastle upon Tyne, NE2 4HH, email: domhnall.jennings@newcastle.ac.uk.

Running head: Fixed and variable duration CSs

Abstract

In four experiments rats were conditioned to an auditory conditioned stimulus (CS) that was paired with food, and learning about the CS was compared across two conditions in which the mean duration of the CS was equated. In one the CS was of a single, fixed duration on every trial, while in the other the CS duration was drawn from an exponential distribution, and hence changed from trial to trial. Higher rates of conditioned responding to the fixed than to the variable stimulus were observed, in both between- (Experiment 1) and within-subject designs (Experiments 2 and 3). Moreover, this difference was maintained when stimuli trained with fixed or variable durations were tested under identical conditions (i.e. with equal numbers of fixed and variable duration trials) - suggesting that the difference could not be attributed to performance effects (Experiment 3). In order to estimate the speed of acquisition of conditioned responding, the scaled cumulative distribution of a Weibull function was fitted to the trial-by-trial response rates for each rat. In the within-subject experiments specific differences in the pattern of acquisition to fixed and variable CS were shown; a somewhat different pattern was found when ITI was manipulated (Experiment 4). The implications of these findings for theories of conditioning and timing are discussed.

Key words: fixed and variable stimulus duration, CR acquisition, Pavlovian conditioning, conditioned responding

Introduction

It has been argued that Pavlovian conditioning involves encoding the relationship between events that are temporally contiguous, and the interaction of temporal factors with the learning process has been a focus of conditioning research since Pavlov's initial reports (e.g. 1927). Such work has established that both acquisition and rate of conditioned responding are superior when the duration of the CS is short, and the intertrial interval (ITI) duration long. The ratio of the CS duration (T) and ITI duration (I) – the so called I/T ratio - is considered to be the best determinant of the number of trials needed to acquire the conditioned response (CR), even when the absolute values of T and I vary dramatically (Gibbon, Baldock, Locurto and Terrace, 1977¹). Moreover, Perkins et al. (1975) reported that the rate of autoshaped pecking increased with increasing ITI duration but decreased with increasing CS duration (see also Lattal, 1999; Holland, 2000; Terrace, Gibbon, Farrell & Baldock, 1977).

Observations of this type have prompted the development of models that seek to account for conditioning in terms of temporal factors (e.g. Gibbon & Balsam, 1981; Gallistel & Gibbon, 2000, 2002; Jenkins, Barnes & Barrera, 1981). These models primarily explain *timing* - the animals' ability to show a gradual increase in responding over the course of a fixed duration CS that peaks at the point of US delivery – but can be extended to explain the acquisition of conditioned responding. For example, according to one of the earlier models of this type, Gibbon & Balsam's Scalar Expectancy Theory (SET; 1981 - cf. Gibbon, 1977), the degree to which the rate of reinforcement during the CS exceeds the overall rate in the background determines the "strength of associative responding" (Gibbon & Balsam, 1981 p.225). As reinforcement rate during an interval is inversely related to its duration, SET can explain the observed relationship between the strength of the CR and the I/T ratio - a

¹ Some models substitute cycle time (C) for I, where C is the intertrial interval between subsequent US deliveries, and thus equals I+T.

prediction that is shared by other models (e.g. Kirkpatrick & Church, 1998; Kirkpatrick, 2002; see also Machado, 1997). Some authors, however, place greater emphasis on explaining the effect of I/T ratio on the rate of CR acquisition, rather than CR strength. For example, according to Rate Estimation Theory (RET: Gallistel and Gibbon, 2000; 2002), the relative rate of reinforcement during the CS and the background - the I/T ratio - determine the point in training at which the CR is acquired. A recent development of this type of model is that proposed by Balsam and Gallistel (2009), who interpret the conditioning procedure as one which reduces the uncertainty of when a US will occur, using a mathematical definition of information (Shannon, 1948). They argue that the rate of acquisition is determined by the degree to which the onset of the CS reduces the expected time to reinforcement - defined as the "informativeness" of the CS; the longer the ITI relative to the CS, the more informative the CS becomes and the faster is the acquisition of the CR. Because such theories seek to explain how two different aspects of behaviour can be accounted for within a single model, they can be referred to as hybrid models of learning, in order to distinguish them from models that address only conditioning or timing.

In contrast, theories specifically developed to explain conditioning have been slow to accommodate the effects of temporal factors on conditioned behaviour (e.g. Pearce & Hall, 1980; Mackintosh, 1975; Rescorla & Wagner, 1972; Wagner, 1981 but see Sutton & Barto 1987, 1990; Vogel, Brandon & Wagner, 2003). These accounts rely on association formation to explain acquisition of the conditioned response, and have no principled or quantitative account of the effects of I/T ratio on either the emergence or the final level of conditioned responding, or timing of US occurrence. In view of this failure of associative models to explain timing effects (although see e.g. Sutton & Barto 1987; Vogel, Brandon & Wagner, 2003), some have argued that it would be more parsimonious to adopt a hybrid theory to provide a unified account of conditioning and timing (e.g., Church & Broadbent, 1990). But from an associative perspective, current formulations of hybrid theories have serious shortcomings. For example,

most have difficulty explaining cue competition phenomena, which mediate the ability to selectively associate events that are truly correlated (e.g. Kirkpatrick, 2002). Moreover, the explanations they do offer rely on the assumption that effects such as overshadowing and blocking are complete (e.g. Balsam & Gallistel, 2009; Gallistel & Gibbon 2000, 2002), and as a result are forced to attribute the observation that cue competition effects *appear* to be graded in magnitude to an artefact created by averaging over a number of subjects (e.g. Balsam & Gallistel, 2009). In summary, weaknesses in both classes of theory suggest a need to identify principled reasons for adopting one or other of these two theoretical perspectives.

One clear point of difference that might help distinguish these approaches stems from the way in which they assume that information about the environment is extracted. Hybrid accounts often appear to require that information about environmental events is integrated over a number of trials, in order to obtain accurate estimates of (for example) rates of reinforcement, which are then used to compute parameters reflecting the global properties of the environment that can be used to guide behaviour. In essence such models gather information about the environment, and use it to make a decision about when to respond - and in this sense may be termed decision models. For example, both SET and RET assume that animals obtain information about the I/T ratio by tracking the accumulated durations of the CS and the ITI as training proceeds, and use them to compute the rates of reinforcement in these two intervals. Similarly, in order to compute the informativeness of the CS (cf. Balsam & Gallistel, 2009), estimates of the reinforcement rate during the CS and the background are required. Thus it is a summary measure of, or *cumulative* CS duration across trials that determines behaviour according to these models - and although the time window that is required to make these estimates is not specified, the strong implication is that the characteristics of individual CS/US pairings – such as the CS duration on a particular trial – are not necessarily critical (see e.g. Bouton & Sunsay, 2003).

The present set of experiments explores whether this question of whether trial by trial

information is in fact critical, by examining both speed of acquisition and level of conditioned responding with CSs that varied in their temporal characteristics. More specifically, we compared learning about a CS that was of a fixed duration with learning about one whose duration varied from trial to trial, but whose mean duration matched that in the fixed duration condition. Thus, whereas the fixed CS gives precise information as to when the US will occur, by comparison the variable CS does not, as its onset does not help the animal predict when the US will occur. The fact that the mean duration was the same means that hybrid models of the type discussed above will not, as a rule, discriminate between these two conditions (although see discussion of Balsam and Gallistel (2009) below). As a result, models predicting that the I/T ratio determines the rate of conditioned responding predict that fixed and variable CSs will produce similar response rates, (Gibbon & Balsam, 1981; Kirkpatrick & Church, 1998; Kirkpatrick, 2002; see also Machado, 1997), while models which assert that the I/T metric determines the rate of acquisition of the CR (e.g. Balsam and Gallistel 2009; Gallistel & Gibbon, 2000) also predict acquisition will be identical in the two conditions. To the best of our knowledge the sole potential exception is a recent model proposed by Balsam and Gallistel (2009); as mentioned above, these authors argue that acquisition rate increases with the "informativeness" of CS onset - where informativeness is proportional to the C/T ratio. Nonetheless, these authors note that although informativeness determines one component of the Shannon information provided by CS onset about the time of food delivery, arranging for the CS to have a fixed rather than a variable duration adds further to the information provided by CS onset. Accordingly they suggest that it is also in the spirit of their approach to ask whether a fixed duration CS will be learned about more quickly than one of variable duration (Ward et al., 2012).

In contrast, associative accounts assume that learning occurs on a trial-by-trial basis; thus differences in CS duration on each trial may have a profound influence on what is learned, even when cumulative CS duration (in terms of averages) is equated. For example, if

one adapts the standard associative model slightly, and assumes that each CS is composed of smaller elements that are presented in sequence, then it is possible to intuit how conditioning to fixed and variable duration CSs might differ. Let us assume that the mean duration of both fixed and variable CSs is 2 units; but while the fixed CS is 2 units on every trial, the variable is either 1, 2 or 3 units in duration. In addition, for the sake of simplicity let us assume that only the final unit, that is contingent with US delivery, acquires associative strength. In the fixed case, as unit 2 is the only unit contingent with US delivery, it is the only unit to acquire associative strength; as it maintains the same temporal relation to the US on every trial, it will eventually reach asymptote. In contrast, units 1, 2 and 3 of the variable CS are all adjacent to the US on some trials but not on others; unit 1 will be reinforced on 33% of trials and nonreinforced on 66%; unit 2 will be reinforced on 33% and nonreinforced on 33%, and unit 3 will be reinforced on 33% of trials, and is never nonreinforced. Such effects could produce differences in both the rate at which acquisition to fixed and variable CSs occurs, and also the level of responding; a more detailed discussion of the precise predictions made by associative models will be taken up in the general discussion.

Intuitions notwithstanding, empirical studies examining these issues have not provided evidence for a consistent difference in either the rate of responding or the speed of acquisition to fixed and variable CSs. With respect to rate of conditioned responding, Kirkpatrick and Church (1998) reported a "subtle" (and statistically significant) superiority in conditioned responding to a fixed CS (Kirkpatrick & Church 1998; see also Jennings, Alonso, Mondragón & Bonardi 2006, 2011) while others have reported no difference (Kamin, 1960; Low & Low, 1962; Patterson, 1970), and the results of a series of studies by Libby & Church (1975) were inconclusive. With respect to rate of acquisition of the CR, to our knowledge only one study has compared learning about these different distribution forms (Ward et al., 2012 Experiments 1 and 3); no difference in acquisition to fixed and variable CSs was found (although *IT* ratio manipulations did have the predicted effect on CR acquisition). The present experiments

extend this work. First, we attempt to replicate the higher levels of responding to the fixed duration CS reported by Kirkpatrick and Church, and rule out potential artefacts. We also examine the issue of whether there are differences in speed of CR acquisition to fixed and variable CSs, in order to examine the relationship between the temporal form of the CS and the way in which the CR is acquired.

Experiment 1

Two groups of rats were conditioned to a clicker which signalled a food pellet. For one (Group F) the CS had a mean duration of 60s, whereas for the other we followed the procedure employed by Kirkpatrick and Church (1998), and arranged for it to have a variable duration drawn from an exponential distribution with an arithmetic mean² of 60s (Group V). Since the cumulative duration of the CS within each pair of groups was equated, hybrid timing models making predictions about the rate of conditioned responding assert that it should be identical in the two conditions, while those dealing primarily with the rate of acquisition would either predict no difference (e.g. Balsam & Gallistel, 2009; RET; Gallistel & Gibbon, 2000) or, with added assumptions, faster acquisition in the fixed condition (Balsam & Gallistel, 2009). We assessed differences in the level of conditioned responding by examining response rates throughout the course of training. We also examined differences in acquisition speed, by fitting the scaled cumulative distribution of a Weibull function to the trial-by-trial response rates for each rat. This method has been argued to be theoretically unbiased (cf. Gallistel, Fairhurst & Balsam, 2004; Harris, 2011), and yields optimal fit parameters which allow interpretation of different features of the acquisition function. It also gives other estimates of the speed of

² Use of either the geometric or harmonic mean yields average durations for the variable CS that are *lower* than that of the fixed, leading SET, for example, to predict *less* responding on fixed than on variable trials. As our results suggested the opposite, we viewed this as a conservative strategy.

acquisition, such as the number of trials it takes for responding to reach 10% (*onset latency* - which can be seen as the start of acquisition), and the *dynamic interval* - the number of trials required to go from 10% to 90% of the individual's asymptotic rate of responding.

Method

Subjects:

Subjects were 32 male Lister hooded rats (Harlan UK) with a mean free-feeding weight of 321g (range: 275-355 g); the experiment was run in two replications, with 16 rats in each replication. They were deprived to 85% of their ad lib weight before the start of the experiment, and maintained at this level (with adjustments for natural growth rate) by being fed a restricted amount of food at the end of each session; they were housed in pairs in plastic tub cages with sawdust bedding. Water was freely available in the home cages. They were maintained on a 12-hour light/dark cycle, the light period starting at 7am; the temperature was maintained at 21°C ($\pm 1^\circ$), and the humidity at 60% ($\pm 10\%$).

Apparatus:

The apparatus comprised eight identical chambers (20 x 24 x 30 cm), each situated in a ventilated, noise-attenuating box (74 x 38 x 60 cm, MED Associates). Each chamber was equipped with a houselight, a food cup, and a speaker, located on the right side of the wall opposite to the food cup, which could deliver a 75-dB 10-Hz clicker, a 75-dB 4-kHz tone and a 75-dB white noise. A pellet dispenser (Model ENV-203) delivered 45-mg Testdiet pellets (MLab Rodent Tablets) into the food cup. Each head entry into the food cup was detected by an LED-photocell, and recorded as a single response. Med-PC for Windows (Tatham & Zurn, 1989) controlled experimental events; trials of the same duration were delivered at the same

time across experimental chambers. The time of occurrence of each stimulus onset, stimulus termination, food delivery, and head entry response was recorded with a resolution of 10 ms.

Procedure:

In each replication subjects were semi-randomly assigned to one of two different groups ($n = 8$). The rats received four training sessions of 30 trials, each comprising the presentation of the clicker followed immediately by a single food pellet. In addition there were 30 nonreinforced trials with a tone. The trials were arranged in three, 20-trial blocks each comprising 10 trials with the clicker and 10 with the tone. The click was of a fixed 30s duration in Group F, and a variable duration with a mean of 30s in Group V; of the tone trials, five were fixed 30s and 5 of variable duration with a mean of 30s. The ITI comprised a fixed 45s plus a variable 45 seconds. Following Kirkpatrick and Church (1998) these variable CS durations were drawn from an exponential distribution with the same arithmetic mean as that of the CS in the fixed condition (Evans, Hastings & Peacock, 1993). No upper or lower limit was set for the range of potential distributions drawn by Med-PC. Animals were run in squads of eight subjects. Thus subjects in the variable group, although receiving a stimulus that varied in duration from trial to trial, had the same overall mean exposure to the CS per session as those in the fixed condition (cf. Kirkpatrick & Church 2000). In addition each trial was preceded by a 30s pre-CS period, giving a session length of approximately three hours. Both the sequence of the two types of trial (reinforced and nonreinforced) and the duration of each successive trial was identical in the two replications.

Data analysis:

Conditioned responding. Mean response rates during each type of trial were obtained by computing the total number of responses made during each CS type in each session, and

dividing by the total CS duration. Conditioned responding in each session was indexed by a *difference* score - the mean response rate during the reinforced CS after subtraction of the corresponding rate during the nonreinforced stimulus. All response rates are reported in responses per minute (rpm).

Acquisition speed. The rate of responding (number of responses per trial/trial duration) was calculated for each individual training trial for each rat, and corrected for background responding by subtracting the response rate from the corresponding pre-CS period. These data were then smoothed by calculating response rates over a 5-trial moving average (i.e. trials 1-5, 2-6, 3-7 etc), to avoid trial-by-trial variability in responding obscuring meaningful differences (Harris, 2011). A scaled Weibull cumulative function (Equation 1) was then fitted to these data. Weibull functions have been employed to study acquisition patterns in various conditioning preparations because the large trial-by-trial variability in animals' responding makes estimation of CR acquisition from the actual data difficult; thus the optimal-fit parameters of the Weibull function are instead used to interpret different aspects of the acquisition function (Gallistel et al., 2004; Harris, 2011).

$$R = \lambda (1 - e^{-(t/\beta)^s}) \quad (\text{Eq. 1})$$

Interpreting these parameters in the context of an acquisition function, responding (R) is expressed as a function of the number of trials (t) that have occurred, using three parameters: *slope* (s), a measure of the change in response rate; *latency* (β), the mean number of the trial at which R reaches 63% of its asymptotic value; and asymptote (λ), the asymptotic level of R (Gallistel et al., 2004). Fitting was performed using the L-BFGS-B matrix method (Byrd, Lu, Nocedal & Zhu, 1995) in the scipy library (<http://www.scipy.org>) for Python (<http://www.python.org>). In order to understand the range of reasonable values of these

parameters, a bootstrap procedure was applied to generate 5000 resampled datasets for each individual. For each of these, the curve was fitted again, using an identical procedure, and the parameter values stored. The standard deviation of the parameter values fitted to the resamples approximates the standard error of the mean for the respective parameter for each individual (Efron & Tibshirani, 1994). The maximum meaningful value for slope was set at 200. Relationships between the three Weibull parameters were evaluated with partial correlations that controlled for whether the data were derived from fixed or variable trials. Model fit was given by R^2 .

We computed two measures of acquisition - *onset latency*, the trial number at which animals reached 10% of asymptote, and the *dynamic interval*, the number of trials required to progress from 10% to 90% of asymptote. For one animal (in the 60s condition of Experiment 2) the asymptote parameter was negative, and the dynamic interval could not meaningfully be calculated; this animal was omitted from analyses of these two acquisition measures.

A significance level of $p < .05$ was adopted in all analyses. Significant two-way interactions were examined with simple main effects analysis, using the pooled error term for all between-subject comparisons. Mean trials to 10% of asymptote, and the number of trials between 10% and 90% of asymptote, were subjected to log transformation to correct violations from normality.

Timing: To examine the degree to which timing occurred, the number of responses occurring in successive 1-s time bins during the CS in each session was determined, and the rate of responding in each bin for each rat calculated. These response rate functions were then normalized so that each rat would contribute equally to the shape of the functions regardless of its overall response rate. Thus the response rate in each time bin was divided by the total number of responses and multiplied by 100, giving the percentage of total responses in each time bin for each subject. Then a linear function was fitted to each normalized response rate

function, and the slope determined from the best-fitting linear curve for each rat (linear fits provide a good characterization of the response rate function: Jennings et al. 2007; cf. Kirkpatrick & Church 2000). We will refer to this slope parameter as *temporal slope* in the present article, to discriminate it from the slope parameter of the Weibull function. The temporal slopes were compared against a mean of zero using one-sample t-tests on each session, applying the Bonferroni correction.

Results

Conditioned responding: Figure 1 (panel a) shows that animals in Group F responded at a higher rate than those in Group V, and this difference was reliable; although ANOVA revealed no effect of group, $F(1,30)=2.83$ $p=.10$, there was a main effect of block, $F(11,30)=39.29$ $p<.001$, and a significant interaction between these two factors, $F(11,330)=2.51$ $p=.005$; simple main effects revealed that the groups differed on blocks 7 and 8, $F(1,360)=13.87$ and 4.14 , $p=.002$ and $.043$ respectively; they also marginally differed on block 10, $F(1,360)=3.22$ $p=.074$. Neither responding during the control stimulus or during the preCS periods differed between the two groups; ANOVAs revealed only main effects of block, $F(11,330)=32.27$ and 47.42 , $ps<.001$; nothing else was significant, largest $F(11,330)=1.43$ $p=.156$. The significant effects of block reflected, for the control CS, a rise from 2.2 rpm in block 1 to a peak of 6.5 rpm in block 4, and then a gradual decline to 0.7 rpm in the final block; the corresponding values for the preCS were 3.7 rpm on block 1, rising to a high of 6.4 rpm on block 3, and then to 0.7 rpm on the final block.

Figure 1 about here

Acquisition speed: The mean value of the asymptote parameter was significantly higher in

Group F, $F(1,31)=4.34$ $p=.046$ $MSE=31.3$ (see Table 1). Latency was numerically greater in the fixed group, and slope greater in the variable group, but these values did not differ, $F_s < 1$. Asymptote was positively correlated with latency, $r=.56$ $p=.001$; the relationship between asymptote and slope was negative but not significant, $r=-.33$ $p=.07$; nor were latency and slope related, $r=-.25$ $p=.18$. The mean value of R^2 was .65 and .51 for Groups F and V respectively, $F(1,31)=4.24$, $p=.048$ $MSE=.04$, indicating a better fit for the fixed group.

Table 1 about here

Both the onset latency and the dynamic interval were numerically greater in Group F (see Table 1), but these values did not differ, $F_s < 1$. The data from the rat with the best fitting function in each group, and the actual response rates per trial are shown in Figure 1 (panel b).

Timing: The temporal slopes are shown in Figure 2 (panel a). It is clear that a marked difference emerged between the two groups, with the temporal slope for Group F ending markedly higher than that for Group V. ANOVA revealed significant main effects of group and session, $F(1,30)=28.60$ and $(3,90)=15.51$, $p_s < .001$, and a significant interaction between these two factors, $F(3,90)=2.75$ $p=.047$; Group F had a higher temporal slope than Group V from session 2, $F(1,120)=2.55$, 6.85, 13.75 and 26.38, $p_s=.11$ and .01 for sessions 1-2 respectively and $p_s < .001$ for sessions 3-4. The temporal slopes were greater than zero in Group F on sessions 3 and 4, $p < .001$, and less than zero on session 2 in Group V, $p=.032$ (all after Bonferroni correction). Panel b of Figure 2 shows the pattern of responding over the course of the CS in the final training session.

Figure 2 about here

Discussion

The results of Experiment 1 showed higher levels of conditioned responding in Group F than in Group V, thus providing evidence against a subset of the hybrid-type theories (e.g. Kirkpatrick & Church, 1998; Kirkpatrick, 2002; see also Machado, 1997). There was, however, little indication of a difference in the pattern or speed of acquisition of conditioned responding between the two groups in either experiment; no differences in onset latency or dynamic interval were obtained. Nor were there any differences in the Weibull function parameters, apart from asymptote, which was significantly higher in Group F. These results are consistent with the class of hybrid theories that make predictions about speed of acquisition. Nonetheless, inconsistent with an information processing perspective was the relationship between asymptote and latency parameters; the fact that asymptote was not independent of the latency parameter determining how fast it was attained does not accord with the type of decision-making process envisaged by many hybrid models (Gallistel et al., 2004; Harris, 2011). This, Gallistel and colleagues argue (Gallistel et al., 2004), is because these two measures are confounded in learning curves averaged over a number of subjects (because many subjects starting to respond early with a lower asymptote would produce a similar function to a few subjects beginning to respond early with a high asymptote).

Our measure of conditioning in these studies was based on the overall response rate, in line with that routinely used in the associative literature - despite the fact that this measure necessarily ignores the possibility that the distribution of responding might differ across the course of the CS. We justify our choice on the basis of precedent and practical considerations. First, the overall response rate measure is not only used throughout the associative learning literature, but has also been used in timing studies by those attempting to evaluate the ability of hybrid models to explain conditioning and timing within the same theoretical framework (e.g. Kirkpatrick & Church, 1998; Kirkpatrick, 2002; Ward et al., 2012). Second, there is no *a priori* principle for determining at which point of the CS the most accurate index of associative

strength may be obtained; using the entire CS avoids adopting an arbitrary criterion. Finally, the most complete account of the principles governing acquisition of associative strength has been developed in the conditioning literature in which, as noted above, such measures are commonplace. Thus it seems the natural choice to make contact with the conditioning literature, and enable theoretically meaningful conclusions to be drawn about levels of associative strength.

However, before we may draw any firm theoretical conclusions from these results, there are alternative explanations appealing to differences in performance engendered by the different schedules employed that must be ruled out. For example, the unpredictability of food reinforcement on variable trials might have a nonspecific and detrimental effect on performance in Group V, which could in principle result in lower levels of the CR. Such an account would predict *no* difference in responding to fixed and variable cues if both fixed and variable trials were delivered unpredictably to the same animal in the same session, as any hypothetical effect on performance produced by variable trials would affect both types of trial equally. The second experiment aimed to explore this prediction.

Experiment 2

Experiment 2 employed a within-subject design, in which all animals received training with three cues, one nonreinforced stimulus (C) and two further reinforced cues, one (F) that was always of the same fixed duration and the other (V) of a variable duration. If the difference in response rate observed in Experiment 1 were due to a nonspecific effect of the reinforcement schedule on conditioned responding, then no difference between responding on fixed and variable trials should be observed here. The experiment was conducted in two parts, one employing CSs of a mean 30s duration, and the other CSs of a mean 60s duration, which also allowed us to explore the generality of any effects observed.

Method

Subjects and Apparatus:

Subjects were 32 male Lister hooded rats (Charles River UK) with a mean free-feeding weight of 322g (range: 295-360 g; 60s CS range: 295-360 g; 30s CS range: 295-340 g). The rats were maintained exactly as in the previous experiment. The apparatus was identical to that of the previous experiment.

Method

Procedure:

Training: For both sub-experiments there were five sessions of training, each comprising 54 trials, 18 with the clicker, 18 with the noise, and 18 with the tone which served as a nonreinforced control stimulus; these three trial types were intermixed in a semi-random order, with the constraint that every successive 18 trials comprised six of each type. Half the animals trained with each CS duration experienced a fixed duration click and a variable duration noise, and the remainder the converse arrangement; each presentation of clicker and noise was followed by the delivery of a food pellet. Half the presentations of the tone were fixed and the remainder were variable, and tone presentations were always nonreinforced. In one study all CS presentations were on average 60s in duration, and the ITI comprised a fixed interval of 30s plus a variable interval of 60s; there was also an additional 60-s preCS period. In the other study all CS presentations were on average 30s in duration, and the ITI comprised a fixed interval of 30s plus a variable interval of 30s; there was also an additional 30-s preCS period. The inclusion of a fixed interval in the ITI was to ensure that the preCS period always occurred some time after delivery of the previous food pellet, so that responding during the preCS would not be contaminated by responding to food delivery itself. All other aspects of the procedure

were identical to those of the previous experiment.

Data analysis:

The data analysis procedure was identical to that of the previous experiment except that the measure of conditioning during the various kinds of trial was a difference score derived from the rate of responding during the CS minus the response rate in the corresponding preCS period.

Results

Conditioned Responding: The group mean (CS-preCS) response rates during the different types of trial are shown in Figure 3 (panel a). Animals in both CS duration conditions responded more on reinforced trials than on nonreinforced control trials; there was also more responding to the 30s than to the 60s CS. More importantly, responding also appeared to be slightly but consistently higher on F than on V trials. This impression was confirmed by ANOVA with CS duration (30/60s), trial type (F/V) and sessions as factors, which revealed a significant main effect of trial type $F(1,30)=4.98$, $p=.03$ $MSE=14.13$; none of the interactions involving trial type were significant, largest $F(4,120)=1.37$ $p=.25$ $MSE=4.56$. Thus animals responded significantly more on fixed than on variable trials, and this difference was present at both CS durations. There was also a significant main effect of CS duration, $F(1,30)=20.32$ $p<.001$ $MSE=76.43$, of session, $F(4,120)=64.51$ $p<.001$ $MSE=31.59$, and a significant interaction between CS duration and session, $F(4,120)=11.66$ $p<.001$ $MSE=31.59$; simple main effects indicated that responding was higher in the animals trained with a 30s CS on sessions 3, 4 and 5, smallest $F(1,150)=11.45$ $p<.001$ $MSE=40.56$. The mean rates of preCS responding were also higher in the 30s condition; the mean rates for the 60s group were 3.6, 40.7, 2.4, 1.3 and 1.0 rpm for sessions 1-5 respectively, and the corresponding rates for the 30s group were 5.5,

9.2, 5.5, 3.9 and 2.9 rpm. ANOVA revealed a significant interaction between CS duration and session, $F(4,120)=6.49$ $p<.001$ $MSE=2.17$; simple main effects revealed that the groups trained at the two CS durations differed on every session, smallest $F(1,150)=8.60$ $p=.004$ $MSE=3.1$.

Figure 3 about here

Acquisition speed: All ANOVAs had trial type (fixed or variable) and CS duration as factors. The mean asymptote of the fitted functions did not differ between the two trial types, $F<1$; however asymptotes were significantly higher in the 30s group, $F(1,30)=10.03$ $p<.004$ $MSE=442.7$; there was no interaction between these two factors, $F<1$ (see Table 1). Slope was numerically lower on fixed trials than on variable trials at both CS durations; latency was numerically higher for fixed trials at the 60s CS, but around the same for the 30s CS; however ANOVAs revealed no significant effects or interactions, largest $F(1,30)=2.32$ $p=.14$ $MSE=1322.3$. As in the previous study, there was a positive correlation between asymptote and latency, $r=.54$ $p<.001$; the correlations between asymptote and slope, $r=-.21$ $p=.09$, and between latency and slope, $r=-.17$ $p=.18$, were not significant. The mean values of R^2 were .49 and .46 respectively for the fixed and variable conditions of the 60s group, and 0.66 and 0.59 for the 30s group. ANOVA revealed only a main effect of CS duration, $F(1,30)=6.30$ $p=.018$ $MSE=.06$, indicating better fit in the 30s group; the effect of trial type and the interaction were not significant, largest $F(1,30)=1.83$ $p=.19$ $MSE=.03$, indicating comparable fit for the two trial types regardless of condition.

The onset latencies were similar for fixed and variable CSs, but larger for the 30s CS (see Table 1); ANOVA with CS duration and trial type as factors revealed a significant main effect of CS duration, $F(1,29)=5.66$ $p=.029$ $MSE=.1$; nothing else was significant, largest

$F(1,29)=3.24$ $p=.082$ $MSE=.08$. The dynamic interval was, as in Experiment 1, numerically higher in the fixed condition in both groups, an effect which was statistically significant. ANOVA with CS duration and trial type as factors revealed a significant main effect of trial type, $F(1,29)=5.98$ $p=.021$ $MSE=.33$; nothing else was significant, largest $F(1,29)=2.42$ $p=.13$ $MSE=.48$. The data for the rat with the best fitting functions for both trial types, and the actual response rates per trial for that animal, are shown in panels b and c of Figure 3, for animals trained with a 60s and 30s CS respectively.

Timing: As Figure 4 (panel a) shows, there appeared to be evidence for the development of timing in this experiment; at both CS durations temporal slopes for fixed trials increased more than those for the variable trials. In addition, by the end of training slopes were higher for the 30s CS, which is consistent with the sharper timing function that would be expected for the shorter duration. These impressions were confirmed by ANOVA with CS duration (30/60), trial type (F/V) and sessions as factors, which revealed a significant main effect of session, $F(4,120)=8.16$ $p<.001$ $MSE=.002$. The interactions between CS duration and session, and trial type and session, were also significant, $F(4,120)=6.26$ $p<.001$ $MSE=.002$ and $F(4,120)=2.88$ $p<.026$ $MSE=.002$ respectively; nothing else was significant, largest $F(1,30)=2.33$ $p=.14$ $MSE=.003$. The temporal slopes were significantly higher in the 30s group than in the 60s group on session 5, $F(1,150)=9.38$ $p=.003$ $MSE=.003$, on which there was also a significant effect of trial type, $F(1,15)=6.84$ $p<.01$ $MSE=.002$; the temporal slopes differed from zero for the fixed 60s CS on sessions 4 and 5, and for the fixed 30s CS on session 5 $ps<.002$ (after Bonferroni correction). The distribution of responding over the course of the two types of CS in the two CS duration conditions during the last training session is shown in panel b of Figure 4.

Figure 4 about here

Discussion

The results of Experiment 2 replicate those of Experiment 1: animals displayed higher levels of conditioned responding to a cue trained with a fixed duration than to one whose duration was drawn from an exponential distribution. Moreover, in the present study this difference was evident in a within-subject procedure. These observations do not support the view that the difference between conditioning to fixed and variable stimuli observed in Experiment 1 was due to a nonspecific effect of the different distributions on conditioned responding. Nonetheless, there is a second, related possibility that could provide an alternative explanation of these results - namely that it might be more difficult to respond at a high rate on variable than on fixed duration trials. For example, on some variable trials the CS might be too short for the animal to arrive at the foodcup before CS offset, while on others it might be so long that the animals are physically unable to sustain a high rate of conditioned responding³. Such factors could selectively reduce response rate on variable trials - even without any difference in underlying learning. This interpretation predicts that if, after training on either a fixed or a variable CS, the animals were then tested under *identical* conditions, any difference in response rate would be eliminated. Experiment 3 addressed this possibility.

In Experiment 2 there was also an indication that the pattern of acquisition of the CR differed: the dynamic interval was significantly higher for the fixed duration stimulus - suggesting slower acquisition for the fixed cue. This result is intriguing, as it is the opposite of what one would expect in terms of the information given by CS onset about the point of US delivery - which if anything would predict quicker acquisition in the fixed condition. Thus an important purpose of Experiment 3 was to replicate this finding. Finally, as in the previous experiment there was a clear positive correlation between the Weibull parameters

³ It should be noted, however, that the proportion of very short variable trials (less than one second) was low, being no more than 5% of the total number of variable trials).

corresponding to the latency and asymptote of the CR, suggesting that higher asymptotic levels of CR tend to be accompanied by later emergence of the CR.

Experiment 3

Animals in Experiment 3 received identical training to those in the 60s group of Experiment 2, with a stimulus F trained with a constant 60s duration, and a further stimulus V trained with a variable duration with a mean of 60s. Training was then followed by a test in which F and V were tested under identical conditions: half the presentations of F were fixed and the remainder variable, and the same was true of V presentations. We anticipated that, as in the previous two experiments, we would see higher levels of conditioned responding to F during training. If this were due to an effect on performance, because fixed stimuli are somehow easier to respond to, then at test the difference in responding to F and V should disappear, and be replaced by an advantage of responding on fixed duration trials *regardless* of whether F or V is being presented. But if the difference in responding to F and V reflects a difference in learning, it should be maintained at test. The experiment also aimed to replicate the difference in the dynamic intervals associated with F and V that was observed in Experiment 2.

Method

Subjects and Apparatus:

Subjects were 16 male Lister hooded rats (Charles River UK) with a mean free-feeding weight of 308g (range: 290-325 g). The rats were maintained exactly as in the previous experiment.

Procedure:

Training: Training was identical to that of the 60s group of Experiment 2; there were three

sessions of training.

Test: Training was followed by the test session, which was identical to the training phase except that half of the trials in each block with stimulus F were, as before, of a fixed 60s duration (*F* - trained fixed and tested fixed), and the remainder variable (with a mean duration of 60s; *Fv* - trained fixed and tested variable); likewise, half of the V trials in each block were presented with variable duration as before (*V* - trained variable and tested variable) and the remainder with a fixed 60s duration (*Vf* - trained variable and tested fixed); in each block there were three of each of these four trial types in an 18-trial block; one of each type was nonreinforced. In this way responding to F and V could be compared under identical conditions.

Data analysis: The data analysis procedure was identical to that in the previous experiment.

Results

Conditioned Responding: The group mean response rates (CS-preCS) during the different types of trial are shown in Figure 5 (panel a); higher responding to F than to V was again clearly evident. ANOVA revealed a significant main effect of trial type $F(1,15)=19.12$, $p<.001$ $MSE=7.18$, and also of sessions, $F(2,30)=80.07$ $p<.001$ $MSE=10.49$, as well as a significant interaction between these two factors, $F(2,30)=8.88$ $p<.001$ $MSE=2.28$; responding was significantly higher to F than to V on sessions 2 and 3, $F(1,15)=6.36$ $p=.021$ and $F(1,15)=17.46$ $p<.001$ respectively, $MSE=7.18$.

Conditioning Test Phase: Figure 5 (panel b) shows responding in the test session; here it may be seen that responding on trials with stimulus F (trials F and Fv) tended to be higher than that

on trials with stimulus V (V and Vf); in addition there was a clear effect of the mode of presentation at test - trials with a variable duration at test (V and Fv) commanded lower response rates than those with a fixed duration (F and Vf). These impressions were confirmed by ANOVA with training condition (F and Fv versus V and Vf) and testing condition (F and Vf versus V and Fv); this revealed main effects of both training and testing condition, $F(1,15)=8.09$ $p=.01$ $MSE=5.99$ and 19.43 $p<.001$ $MSE= 1.38$ respectively; there was no interaction between these two factors, $F<1$. Thus there was a highly significant effect of the temporal properties of the CS that were extant during training. In addition the testing distribution had an independent effect on responding at test, with higher rates of responding on fixed duration trials. However, it should be noted that the effect of test stimulus distribution on responding will have been exaggerated over that present under normal training conditions. This is because before the test the animals had only experienced F presented at the training duration - so that presenting F with variable durations at test was thus likely to produce a detrimental effect on performance. For V, in contrast, the animals are unlikely to have noticed the difference between the training and testing conditions, so far less difference in responding between training and test would be expected.

Figure 5 about here

Acquisition speed: The mean asymptote of the fitted functions was numerically higher for stimulus F, but this difference only approached conventional levels of significance, $F(1,15)=4.19$ $p=.059$ $MSE=194.34$. Latency was again numerically greater on fixed trials, and slope greater in the variable condition; the latter difference was significant, $F(1,15)=1.88$ $p=.19$ $MSE=305.92$ and $F(1,15)=5.58$ $p=.032$ $MSE=2590$ respectively (see Table 1). As in the previous studies, asymptote was strongly positively correlated with latency, $r=.67$ $p<.001$; the correlations between slope and both asymptote and latency were not significant, $r=-.15$ $p=.43$

and $r=-.01$ $p=.83$. The mean value of R^2 was .82 and .62 for the fixed and variable conditions respectively, and these values differed, $F(1,15)=34.47$ $p<.001$ $MSE=.009$.

The onset latencies were similar in the two conditions, $F<1$; but, just as in the previous experiment, the dynamic interval was significantly greater in the fixed condition, $F(1,15)=6.40$ $p=.023$ $MSE=.293$ (see Table 1). The data for the rat with the best fitting functions for both trial types, and the actual response rates per trial for that animal, are shown in Figure 6 (panel a).

Timing: As Figure 6 (panel b) shows, there was not strong evidence of timing in this experiment. ANOVA on the temporal slopes revealed no significant effects or interactions, largest $F(1,15)=2.43$ $p=.14$ $MSE=.00$. The temporal slopes did not differ from zero on any session, smallest $p=.19$ (after Bonferroni correction). The distribution of responding over the course of the two types of CS during the last training session is shown in Figure 6 (panel c).

Figure 6 about here

Discussion

The results of Experiment 3 confirmed the results of the previous experiment: there were higher levels of conditioned responding during stimulus F, trained with a fixed duration than during V trained with a variable duration. Moreover, the difference in responding to F and V was maintained during the test phase when both were both tested under identical conditions. Although the stimulus being fixed or variable at test had an effect on responding, there was a clear, independent effect of greater responding to stimulus F that had been conditioned with a fixed duration - a difference which cannot be explained in terms of differential ease of responding on fixed and variable duration trials.

One feature of the data from Experiment 3 that deserves mention was that no timing seemed to be manifest - there was no difference in the distribution of conditioned responding over the course of the CS, or significant difference in temporal slopes despite strong evidence of conditioned responding. This observation is consistent with the results of Experiment 2 in which, for the 30s, CS timing only emerged on session 5 - but does not fit well with the underlying principle of hybrid theories, that the development of timing essentially underlies that of conditioned responding. Nonetheless, the true extent of timing could conceivably be obscured by levels of background responding. For example to the extent that the animals learn to predict the occurrence of the next trial, responding might gradually increase with time since the end of the previous trial, even in the absence of CS presentation. As timing differences were not the core reason for this work we did not examine such possibilities - but simply note that the timing differences that we did observe could conceivably be obscured by such factors.

There was again evidence of a strong positive correlation between the optimal fit parameters of the Weibull function corresponding to asymptote and latency, just as in Experiments 1 and 2. Interestingly, the experiment also confirmed the observation made in Experiment 2, that the dynamic interval was lower for the variable CS, suggesting faster acquisition to this stimulus; conversely, there was no difference in onset latency for the two stimuli. The suggestion that there might be differences in speed of acquisition to fixed and variable CSs is ostensibly inconsistent with results recently reported by Ward et al. (2012), in which they examined a similar question in mice. As noted above, they used change point analysis to determine acquisition speed (e.g. Gallistel, Fairhurst & Balsam, 2004), a technique which detects the rapid increases in the rate of responding that are said to occur as the CR is acquired. Using this method they established that the higher the I/T ratio, the earlier the change point; however, they found no effect of whether the CS was fixed or variable on this change point measure (Experiments 2 and 3). One possible reason for the apparent discrepancy between these sets of findings could lie in the techniques employed to assess

acquisition, and the definition of what constitutes rapid acquisition. Change point analysis is designed to detect the trial on which an abrupt increase in responding occurs. If acquisition is abrupt and occurs in a handful of trials, then such a measure will be the only relevant measure of acquisition, and presumably closely related to onset latency; the dynamic interval is effectively redundant as acquisition is assumed to be uniformly rapid. But if acquisition is not uniformly abrupt, but gradual in some animals, then there could be a dissociation between onset latency or dynamic interval. Thus, if onset latency may be taken as equivalent to the change point, then our results are perfectly consistent with those reported by Ward et al. (2012); but if acquisition is not abrupt, then differences in the dynamic interval measure could still be present - as they appeared to be in our studies. This interpretation would be greatly strengthened if we could demonstrate that the curve-fitting measure of acquisition used in the present studies can yield a difference in acquisition corresponding to that reported by Ward et al. (2012) in animals trained with different *IT* ratios. To produce such a demonstration was the purpose of the final experiment.

Experiment 4

Two groups of animals were conditioned with a single CS of a fixed, 10-s duration; for Group 60 the *ITI* was on average of 60s duration, and for Group 480 it was on average 480s in duration. This *ITI* manipulation results in greatly differing *IT* ratios that should, according to the results of Ward et al. (2012), produce significantly faster acquisition in the latter group. The aim was to replicate this effect of *IT* ratio with our measure of acquisition speed.

Method

Subjects and Apparatus:

Subjects were 16 male Lister hooded rats (Charles River UK) with a mean free-feeding weight of 309g (range: 285-330 g). The rats were maintained exactly as in the previous experiment. A set of eight standard Skinner Boxes (supplied by Campden Instruments Ltd.) was used; these were fitted with Med Associates food cups and pellet dispensers identical to those used in the previous experiments. Each box had three walls of sheet aluminum, a transparent plastic door as the fourth wall, a grid floor and an aluminum ceiling. Each box was housed in a sound- and light-attenuating shell. The auditory stimulus was a white noise identical to that used in the previous experiments. The boxes were controlled by a MED Associates operating system identical to that of the previous experiment.

Procedure

Training: All animals received six sessions of training, each comprising 24 presentations of a 10-s white noise, each of which was followed by delivery of a food pellet. Each CS presentation was preceded by a 10-s preCS period. In both groups the ITI was of a variable duration drawn from an exponential distribution. For Group 60 the ITI had a mean of 50s (with the preCS giving a mean of 60s) and for Group 480 the ITI had a mean of 470s (with the preCS giving a mean of 480s), yielding *IT* ratios of 6 and 48 respectively.

Data analysis: The data analysis procedure was identical to that of the previous experiment.

Results

Conditioned responding: The course of conditioning can be seen in Figure 7 (panel a); it is clear that there was a substantial effect of ITI on conditioning, with faster acquisition of the CR in Group 480. ANOVA with group and sessions as factors revealed a significant interaction,

$F(5,70)=8.87$ $p<.001$; the groups differed on sessions 2, 3, 4 and 5, smallest $F(1,84)=15.08$ $p<.001$; by the last session, however, the groups were responding at similar rates, $F<1$. The mean rates of preCS responding were 3.9, 2.3, 1.2, 0.4, 0.1 and 0.4 rpm for sessions 1-6 respectively in Group 480, and 7.5, 11.4, 11.2, 6.6, 4.5 and 3.1 for Group 60. Again ANOVA revealed a significant interaction between these two factors, $F(5,70)=9.15$ $p<.001$, and the groups differed on every session, smallest $F(1,84)=6.78$ $p=.01$ for session 6. Nonetheless, although it is likely to have had an influence, it is unlikely that this modest difference in preCS response rate could wholly account for the substantial difference between the two groups in their difference scores (which exceeded 20 rpm on the second session).

Figure 7 about here

Acquisition speed: The mean asymptote of the fitted functions did not differ between the two groups, being 21.1 for Group 480 and 24.2 for Group 60, $F<1$. Latency was 78.1 for Group 60 and 31.2 for Group 480, and these values differed, $F(1,15)=35.03$ $p<.01$ $MSE=251.81$; slope was numerically higher for Group 60, the means being 54.5 and 6.2 for Groups 60 and 480 respectively, but these values did not quite differ significantly, $F(1,15)=4.48$ $p=.053$ $MSE=2083.89$. As in the previous studies, asymptote was strongly positively correlated with latency, $r=.67$ $p<.006$; in this experiment it was also negatively correlated with slope, $r=-.551$ $p=.033$; finally latency and slope were negatively correlated, $r=-.637$ $p<.011$. The mean values of R^2 were .71 and .70 for Groups 60 and 480 respectively, $F<1$, indicating comparable fit in the two groups.

The onset latency was 62.3 in Group 60 and 16.3 in Group 480, and these values differed significantly, $F(1,15)=191.24$ $p<.001$ $MSE=.007$; however the dynamic interval was very similar in the two groups, at 24.5 and 25.5 for Groups 60 and 480 respectively, and these values did not differ, $F<1$. The data for the rat with the best fitting functions in each group, and

the actual response rates per trial for that animal, are shown in Figure 7 (panel b).

Discussion

Animals in Group 480 were trained with a longer ITI than those in Group 60, and required fewer trials to reach 10% of asymptotic responding, suggesting faster acquisition in the former group. Thus, if this onset latency measure may be regarded as equivalent to the change point, the results of this experiment confirm the findings of Ward et al. (2012) using an alternative means of measuring acquisition. However, in terms of dynamic interval, no difference was observed between the two groups. These findings may resolve the apparent inconsistency between the results reported by Ward et al. (2012) and those from our experiments, as they suggest that the two measures of acquisition may be dissociable. We found that differences in ITI produced significant effects on onset latency, but no effect on the dynamic interval, while differences in the temporal distribution of the CS produced the opposite pattern.

Figure 8 about here

General Discussion

In each of Experiments 1-3 there was evidence that the rate of conditioned responding was higher to CSs trained with a fixed duration than to those whose duration varied from trial to trial. Furthermore, these effects were demonstrated in both between and within-subjects procedures, and were maintained when animals were tested under identical conditions. These observations suggest that the higher responding on fixed trials reflects a true difference in associative strength, rather than some idiosyncratic effect of the different schedules on

performance. Such a difference in the level of conditioned responding to these two types of stimulus is not consistent with those hybrid or decision theories that make predictions about rates of conditioned responding (e.g. Gibbon & Balsam, 1981; Kirkpatrick & Church, 1998; Kirkpatrick, 2002; see also Machado, 1997). As outlined in the introduction, although some studies had previously examined differences in learning about fixed and variable duration CSs, their results were generally inconsistent. Nonetheless, some of these studies showed higher asymptotic rates to fixed duration stimuli (Jennings et al, 2006; 2011; Kirkpatrick & Church, 1998; but see Kamin, 1960); the present results confirm these findings, while at the same time ruling out several potentially artefactual explanations.

For example, one could argue that if higher responding on fixed duration trials were an *unconditioned* effect on performance, and increased the likelihood that the US was collected promptly at the end of the stimulus, then this could indirectly boost the level of conditioning to the fixed duration CS. Although there was no sign of such a difference in responding at the *start* of training in any of the experiments, as such an account would predict, we examined the latency between the end of the CS and the first response as a function of CS type. In none of the experiments was there evidence that the animals collected pellets more quickly at the end of the fixed CS; in fact the difference in latency to respond after fixed and variable duration CSs was not significant in any of the experiments ($F_s < 1$) - although in Experiment 1 there was a tendency for animals to be *slower* to respond after the offset of the fixed duration CS which approached statistical significance, $F(1,30)=4.12$ $p=.051$ (the mean latencies being 9.6s and 6.6s for Groups F and V respectively) - the opposite of what such an alternative account would need to assume. We, therefore, found no evidence that our results are a by-product of an unconditioned effect of stimulus distribution on performance.

But although finding a difference in conditioned responding to fixed and variable duration stimuli is inconsistent with some hybrid theories, many others make predictions not about the asymptotic rate of conditioned responding, but about the speed with which it is

acquired (e.g. RET; Gallistel & Gibbon, 2000; Balsam & Gallistel, 2009). Our data are not consistent with this class of model either, as in Experiments 2 and 3 we found a significantly shorter dynamic interval for variable than for fixed trials. (It has been argued that within-subject procedures are more sensitive for detecting acquisition differences (cf. Gottlieb & Rescorla, 2010), which may explain why we did not observe the same effect in our between-subject Experiment 1). As noted above, the majority of hybrid theories do not predict a difference in speed of acquisition for fixed and variable CSs; the only potential exception is the recent model of Balsam and Gallistel (2009), where it is argued that informativeness forms the basis of acquisition of the CR. Given the additional assumption that the extra information given by a fixed duration CS can contribute to the speed of CS acquisition, then it is in the spirit of their model to predict a difference (cf. Ward et al., 2012) - but the opposite to that reported here.

At face value these effects contradict the results reported by Ward et al. (2012), who found no difference in acquisition speed between fixed and variable duration CSs, although a substantial difference when I/T ratio was manipulated. We argued that the root of this apparent discrepancy might lie in the indices used to measure acquisition - we used dynamic interval, whereas Ward et al., used change point, which we argued was equivalent to our onset latency measure. We suggested that if acquisition is abrupt, then dynamic interval will be minimal, and onset latency the primary measure of acquisition speed - but if acquisition is gradual, then dynamic interval could be considered an additional and independent measure of acquisition speed. Consistent with this analysis, in Experiment 4 we demonstrated that I/T ratio had a profound effect on onset latency, but none on dynamic interval - the opposite effect to that seen in the preceding experiments.

This argument relies on the assumption that acquisition is gradual - which is contrary to what is supposed by many information processing models incorporating a decision process (e.g., Church & Broadbent, 1990; Gallistel & Gibbon 2000; Gallistel, Fairhurst & Balsam 2004). Such models argue that acquisition often *appears* to be gradual because of an averaging

artefact - the ability of individual animals to determine reinforcement rate varies, and so responding starts on different trials (Gallistel et al., 2004 Figure 2), and pooling over subjects will yield an apparently gradual increase in the CR. In contrast, associative models employ an error correction term to describe learning, such that conditioned responding is proportional to the difference between the maximum associative strength that a US will support, and any strength that has been acquired to that point. Thus acquisition of the CR will be gradual *within* an individual, occurring most rapidly at the start of training and at a negatively accelerating rate thereafter, as the difference between the current associative strength and the asymptote declines (e.g. Hull 1943; Rescorla & Wagner 1972). Moreover, although these models are constrained to predict that learning is gradual, the expression of learning can depend on other factors that lie outside the scope of the theory. Thus, despite the rules governing the way in which associative strength is translated into performance not being specified, these models could accommodate the variability in acquisition speed that was actually observed. Finally Harris (2011), using a simulated dataset, has shown that an abrupt increase in responding could be generated by an underlying sensitivity in the data to random variation in responding, and therefore that an abrupt pattern of responding is possible even when there is in fact an incremental increase in learning.

Recent attempts to explore whether acquisition is abrupt or not have been inconclusive, some reporting abrupt acquisition, more consistent with decision-type models (Gallistel et al., 2004; Morris & Bouton 2006) and others gradual learning, consistent with associative accounts of acquisition (e.g. Kehoe, Ludvig, Dudeney, Neufeld & Sutton 2008; Harris, 2011). To establish the abruptness of acquisition in the present experiments we examined responding by subjects from all four experiments, only considering the data from fixed CSs for consistency with other studies. In addition, in order to avoid the possibility that the smoothing technique used in our acquisition measures might mask a tendency to abrupt acquisition, no smoothing was employed in the abruptness evaluation. After deriving Weibull functions as before, the

dynamic interval was calculated. A scatter plot of the resulting data is shown in Figure 8, for a total of 77 subjects; three subjects were omitted because their functions had a negative asymptote. It is clear that there is considerable variation, but that while some subjects did show clearly abrupt acquisition, with dynamic intervals of 10 trials or less (the criterion for rapid acquisition adopted by Gallistel, Fairhurst & Balsam, 2004), the majority (65%) did not; the mean dynamic interval was 39.9 (the median 17). These results are rather different from those reported by Gallistel, Fairhurst & Balsam (2004) - but one possible reason for the discrepancy might lie in the conditioning parameters employed. For example, Gallistel et al. (2004) analysed some comparable head entry data from rat subjects, and found considerably more abrupt acquisition than that observed here; their general conclusion was that subjects acquired the CR in 10 trials or less. However, these animals were being presented with three, 10-second reinforced CSs in 90 minutes yielding an *IT* ratio of 180 -- considerably higher than those used here (or indeed in most standard conditioning experiments). But whether or not this speculation is correct, the critical point is that acquisition was not uniformly abrupt in the studies reported here, but was gradual in the majority of animals (see also results reported by Harris, 2011). This lack of abruptness lends some credence to the suggestion that onset latency and dynamic interval may in some cases be regarded as independent and dissociable measures of acquisition speed - and it may be that different factors affecting acquisition speed have differential effects on these two measures.

Other aspects of our results are problematic for decision theories. For example, these accounts assume that the emergence of the CR is based on the speed with which evidence may be gathered, and a decision to respond reached. Assuming that the evidence is gleaned rapidly, then such theories predict that the final level of responding should not be systematically related to the speed with which it develops - in other words there should be no relationship between the asymptote of responding and the parameters of latency and slope, which indicate the speed at which the CR emerges (Gallistel et al., 2004). This was not the

case here. In all four experiments there was a highly significant positive relationship between asymptote and latency; such a relationship is consistent with the idea that higher response rates are accompanied by slower acquisition – and inconsistent with the prediction that asymptote should be independent of the speed with which it is acquired.

In summary, the present results add to a body of evidence that casts doubt on the ability of hybrid information processing theories to provide an adequate account of conditioning effects. Indeed one could argue that, by their very nature, models such as that proposed by Balsam and Gallistel are constrained to provide only an impoverished account of conditioning. However detailed the temporal information they may provide about the occurrence of the US, it is limited in the sense that it says nothing about *what* the US might be, or what its motivational valence is – or the degree to which information about the occurrence of one US to the next might be generalized to another. But if these models are not well adapted to explain conditioning, are conditioning models any better equipped to explain timing? The most popular accounts of conditioning are those general models of learning proposed by Rescorla & Wagner (1972), Mackintosh (1975), and Pearce and Hall (1980). However, theories of this type do not incorporate time in an explicit manner, and so cannot address the effects of temporal manipulations. One possible exception is the temporal difference model (Sutton & Barto 1987, 1990; see also e.g. Wagner, 1981; Vogel, Brandon & Wagner, 2000). The temporal difference model is a real-time extension of the Rescorla-Wagner learning rule (Rescorla & Wagner, 1972), which assumes that a stimulus consists of a series of temporally ordered components that can acquire associative strength independently of each other (cf. Moore et al, 1998⁴). The final component, CS_n , conditions directly to the US, but the strength of the component immediately preceding it, CS_{n-1} , will change according to the mismatch between its own associative strength and the associative strength of the final component - essentially second

⁴ The authors are aware that an alternative representation based on microstimulus is under investigation (Ludvig, Sutton & Kehoe, 2012)

order conditioning; CS_{n-2} then conditions to CS_{n-1} , and so on. The amount of associative strength accruing to successive units is determined by a parameter gamma (γ), so that if CS_n acquires an associative strength of 1 unit, CS_{n-1} will acquire this strength discounted by $\gamma - 0.9$ units - and CS_{n-2} will acquire CS_{n-1} 's strength also discounted by $\gamma - 0.81$ units - and so on. Thus higher values of gamma result in more conditioning to CS components earlier in the CS. Moreover, the amount of learning produced by this temporal difference (TD) learning rule is modulated by the magnitude of the *eligibility trace* which grows and declines for each CS component according to a parameter *delta* which is constant for each component, such that high delta means decay is rapid and conditioning curtailed. To allow that as many units as possible "inherit" strength (i.e. even when the limitation imposed by the eligibility trace is high), the value of γ is set to a value close to 1 by default. In combination these considerations ensure that the later portions of the CS will condition more effectively than earlier ones, and yield a timing function when the stimulus is of a fixed duration (Moore & Choi, 1997; Sutton & Barto, 1990).

The temporal difference model can account for our finding of higher response rates on fixed trials, because it can allow that a variable stimulus will acquire less associative strength than a fixed, even though the mean duration of the two stimuli is the same (Gray, Alonso, Mondragón & Fernández, 2012). Although the variable CS will comprise the same total number of time steps as the fixed, on some trials the variable CS will be either shorter or longer than the fixed stimulus. Consequently, because many elements of the variable CS will be contiguous with the US on some trials, and distant on others, they will gain strength on some trials and lose it on others, thus never reaching a stable value; in contrast elements of the fixed stimulus will be able to reach a stable asymptotic value. It is less clear, however, that it can consistent predictions about the rate of acquisition; moreover it should also be noted that the TD model has no systematic explanations of some core results from the conditioning literature, such as the quantitative effect of *IT* ratio on the speed of CR acquisition.

One further issue concerns the relationship between the emergence of conditioning and timing. According to hybrid theories, the emergence of conditioning depends on the assessment of temporal information about US occurrence - with the implication that timing should occur before conditioning. Associative models that incorporate temporal factors, in contrast, would explain timing as a difference in conditioning to different components of the CS according to their proximity to US delivery - so that it would be possible for conditioning to emerge first, and only later for the difference in conditioning to the start and end elements of the CS to develop, so that timing becomes manifest. In Experiment 2 for the animals trained with a 30s CS and Experiment 3, a profound conditioning effect was present well before any discernible timing effect - which could be taken to support the second of these two possibilities, and thus add to an existing body of findings showing similar effects (e.g. Delamater & Holland, 2008). Nonetheless, as we have already noted above, the measures of timing in these studies were potentially compromised by noise in baseline response levels, so these arguments can be only suggestive.

In summary, it is parsimonious that the same theory should be able to account for both conditioning itself, and also the effects of temporal factors on the conditioning process. Until recently the most advanced accounts for such a unified theory were hybrid models developed from a timing perspective, and extended to incorporate an account of conditioning. We have tested the predictions of two classes of such model, and the results cast doubt on their ability to provide an integrated account of acquisition, conditioning and timing effects. In contrast, an adaptation of an associative model was able to provide an explanation of the most reliable aspect of our findings, higher rates of responding to a fixed duration CS. Our results suggest that continued effort should be devoted into developing current associative theories to allow them to explain a greater variety of time-based effects. The results also question the idea that there can be a single measure of acquisition speed, and suggest that a broader approach is required to capture the full subtlety of the effects of temporal parameters on the speed of

conditioning.

Author Notes

The authors would like to thank Eric Tam for helpful discussion and helping us with the Weibull functions, BMSU staff for excellent technical support, and Lorenzo More for collecting the data from Experiment 1. John Peirce programmed the Weibull function and its implementation in Python. This work was funded by the BBSRC.

References

- Balsam, P.D., & Gallistel, C.R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neurosciences*, 32, 73-78.
- Bouton, M.E., & Sunsay, C. (2003). Importance of trials versus accumulating time across trials in partially reinforced appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 29, 62-77.
- Byrd, R.H., Lu, P.H., Nocedal, J. & Zhu, C.Y. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific and Statistical Computing*, 16, 1190-1208.
- Church, R.M., & Broadbent, H.A. (1990). Alternative representations of time, number and rate. *Cognition*, 37, 55-81.
- Efron, B. & Tibshirani, R.J. (1994). *An Introduction to the Bootstrap*. Boca Raton: Chapman & Hall/CRC.
- Delamater, A.R., & Holland, P.C. (2008). The influence of CS-US interval on several different indices of learning in appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 32, 202-222.
- Evans, M., Hastings, N., & Peacock, B. (1993). *Statistical distributions*. New York: Wiley.
- Gallistel, C.R., Fairhurst, S., & Balsam, P.D. (2004). The learning curve: Implications of a quantitative analysis. *Proceedings of the National Academy of Sciences*, 101, 13124-13131.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. *Psychological Review*, 107, 289-344.
- Gallistel, C. R., & Gibbon, J. (2002). *The symbolic foundations of conditioned behavior*. Mahwah, NJ: Erlbaum Associates.
- Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, 84, 279-325.

- Gibbon, J., Baldock, M.D., Locurto, C., Gold, L., & Terrace, H.S. (1977). Trial and intertrial intervals in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, 3, 264-284.
- Gibbon, J., & Balsam, P. (1981). Spreading association in time. In L. C Locurto, H.S. Terrace and J. Gibbon (Eds.) *Autoshaping and conditioning theory* (pp.219-253). New York: Academic Press.
- Gottlieb, D.A., & Rescorla, R.A. (2010). Within-subject effects of number of trials in rat conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 36, 217-231.
- Gray, J., Alonso, E., Mondragón, E. & Fernández, A (2012). Temporal Difference Simulator © V.1 [Computer software]. London: CAL-R. <http://www.cal-r.org/index.php?id=software>
- Harris, J.A. (2011). The acquisition of conditioned responding. *Journal of Experimental Psychology: Animal Behavior Processes*, 37, 151-164.
- Holland, P.C. (2000). Trial and intertrial interval durations in appetitive conditioning in rats. *Animal Learning and Behavior*, 28, 121-135.
- Hull, C.L. (1943). *Principles of Behavior*. Appleton Century Crofts.
- Holland, P.C. (2000). Trial and intertrial interval durations in appetitive conditioning in rats. *Animal Learning and Behavior*, 28, 121-135.
- Jenkins, H.M., Barnes, R.A., & Barrera, F.J. (1981). Why autoshaping depends on trial spacing. In L. C Locurto, H.S. Terrace and J. Gibbon (Eds.) *Autoshaping and conditioning theory* (pp.255-284). New York: Academic Press.
- Jennings, D., Alonso, E., Mondragón, E., & Bonardi, C. (2006). Temporal Uncertainty during overshadowing. *Proceedings of the Society for the Study of Artificial Intelligence and the Simulation of Behaviour: Adaptation in Artificial and Biological Systems*. (T. Kovacs & A.R. Marshall, Eds.): University of Bristol.
- Jennings, D.J., Alonso, E., Mondragón, E. & Bonardi, C. (2011). Temporal uncertainty during

overshadowing: A temporal difference approach. In E. Alonso and E. Mondragón (Eds.). *Computational Neuroscience for Advancing Artificial Intelligence: Models, Methods and Applications*. Hershey, PA: IGI Global.

Jennings, D., Bonardi, C., & Kirkpatrick, K. (2007). Overshadowing and stimulus duration. *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 464-475.

Kamin, L.J. (1960). Acquisition of avoidance with a variable CS-US interval. *Canadian Journal of Psychology*, 14, 1-6.

Kehoe, E., Ludvig, E.A., Dudeney, J.E., Neufield, J., & Sutton, R.S. (2008). Magnitude and timing of nictitating membrane movements during classical conditioning of the rabbit. *Behavioral Neuroscience*, 122, 471-476.

Kirkpatrick, K., & Church, R. M. (1998). Are separate theories of conditioning and timing necessary? *Behavioral Processes*, 44, 163-182.

Kirkpatrick, K., & Church, R. M. (2000). Stimulus and temporal cues in classical conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 26, 206-219.

Kirkpatrick, K. (2002). Packet theory of conditioning and timing. *Behavioral Processes*, 57, 89-106.

Lattal, K.M. (1999). Trial and intertrial durations in Pavlovian conditioning: Issues of learning and performance. *Journal of Experimental Psychology: Animal Behavior Processes*, 25, 433-450.

Libby, M.E., & Church, R.M. (1975). Fear gradients as a function of the temporal interval between signal and aversive event in the rat. *Journal of Comparative and Physiological Psychology*, 88, 911-916.

Low, L.A., & Low, H.I. (1962). Effects of variable versus fixed CS-US interval schedules upon avoidance responding. *Journal of Comparative and Physiological Psychology*, 55, 1054-1058.

Ludvig, E. A., Sutton, R. S., & Kehoe, E. J. (2008). Stimulus representation and the timing of

- reward-prediction errors in models of the dopamine system. *Neural Computation*, 20, 3034-3054.
- Machado, A., (1997). Learning the temporal dynamics of behavior. *Psychological Review*, 104, 241-265.
- Mackintosh, N.J. (1975). A theory of attention: variation in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298.
- Moore, J. W., & Choi, J.-S. (1997). The TD model of classical conditioning: Response topography and brain implementation. In J. W. Donahoe & V. P. Dorsel (Eds.), *Neural-networks models of cognition* (pp. 387-405). New York: Elsevier Science.
- Moore, J.W., Choi, J., & Brunzell, D.H. (1998). Predictive timing under temporal uncertainty: the TD model of the conditioned response, in: D. Rosenbaum, A.C.E. Collyer (Eds.), *Timing of Behavior: Neural, Computational, and Psychological Perspectives* (pp.3-34). The MIT Press, Cambridge, MA.
- Morris, R.W., & Bouton, M.E. (2006). Effect of unconditioned stimulus magnitude on the emergence of conditioned responding. *Journal of Experimental Psychology: Animal Behavior Processes*, 32, 371-385.
- Patterson, M.M. (1970). Classical conditioning of the rabbit's (*Oryctolagus cuniculus*) nictitating membrane response with fluctuating ISI and intracranial CSs. *Journal of Comparative and Physiological Psychology*, 72, 193-202.
- Pavlov, I. (1927). *Conditioned reflexes*. Oxford University Press.
- Pearce, J.M., Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532-552.
- Perkins, C.C., Beavers, W.O., Hancock, R.A., Hemmendinger, P.C., Hemmendinger, D., & Ricci, J.A. (1975). Some variables affecting rate of key pecking during response-independent procedures (autoshaping). *Journal of the Experimental Analysis of Behavior*, 24, 59-72.

- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Theory and research* (pp. 64-99). New York: Appleton-Century-Crofts.
- Shannon, C.E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27, 370-423.
- Sutton, R. S., & Barto, A. G. (1987). A temporal difference model of classical conditioning. Technical report TR 87-509.2. GTE Lab, Waltham, MA.
- Sutton, R. S., & Barto, A. G. (1990). Time derivative models of Pavlovian reinforcement. In M. R. Gabriel & J. W. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497-537). Cambridge, MA: MIT Press.
- Tatham, T. A., & Zurn, K. R. (1989). The Med-PC experimental apparatus programming system. *Behavioral Research Methods, Instruments, and Computers*, 21, 294-302.
- Terrace, H.S, Gibbon, J., Farrell, L., & Baldock, M.D. (1975). Temporal factors influencing the acquisition and maintenance of an autoshaped keypeck. *Animal Learning and Behavior Processes*, 3, 53-62.
- Vogel, E.H., Brandon, S.E., & Wagner, A.R. (2000). Stimulus representation in SOP II: An application to inhibition of delay. *Behavioral Processes*, 110, 67-72.
- Ward, R.D., Gallistel, C.R., Jensen, G., Richards, V.L., Fairhurst, S., & Balsam, P.D. (2012). CS informativeness governs CS-US associability. *Journal of Experimental Psychology: Animal Behavior Processes*, 38, 217-232.
- Wagner, A.R. (1981). SOP: A model of automatic memory processing in animals. In N.E. Miller & R.R. Spear, (Eds.). *Information processes in animals: Memory Mechanisms*. pp.95-128. Hillsdale: Erlbaum.

Figure Legends

Figure 1: Panel a. Group mean difference scores (reinforced - nonreinforced C) (\pm SE) for the 4 training sessions of Experiment 1; the data are presented in 20-trial blocks. *Panel b:* Data from rat in each group with best fitting response functions, and corresponding corrected response rates per trial, for the fixed and variable group of Experiment 1.

Figure 2. Panel a: Group mean temporal slope of responding in the four training sessions of Experiment 1. *Panel b:* Group mean responses per minute (\pm SE) over the course of the CS in the final training session of Experiment 1.

Figure 3. Panel a: Group mean response rates (CS-preCS) during fixed, variable and control trials during each of the five training sessions of Experiment 2, for the groups trained with a 30s and with a 60s CS. *Panels b & c:* Data from rat with best fitting response functions, and corresponding corrected response rates per trial, for fixed and variable trials in animals trained with a 60s (panel b) and 30s CS (panel c) in Experiment 2.

Figure 4. Panel a: Group mean temporal slope of responding for fixed and variable trials for the groups trained with a 30s and a 60s CS in the five training sessions of Experiment 2. *Panel b:* Group mean responses per minute over the course of the CSs in the final training session of Experiment 2.

Figure 5. Panel a: Mean response rates (CS-preCS) during fixed, variable and control trials during each of the three training sessions of Experiment 3. *Panel b:* responding on the test trials with stimulus F (trained with a fixed duration), during fixed and variable duration test trials

(F and Fv respectively), and with stimulus V (trained with a variable duration) during fixed and variable variable duration test trials (V and Vf respectively).

Figure 6. Panel a: Data from rat with best fitting response functions, and corresponding corrected response rates per trial, for the fixed and variable conditions of Experiment 3. *Panel b:* Mean temporal slope of responding for fixed and variable trials in the three training sessions of Experiment 3. *Panel c:* Group mean responses per minute over the course of the CS in the final training session of Experiment 3.

Figure 7. Panel a: Group mean response rates (CS-preCS) in Groups 60 and 480 during each of the six training sessions of Experiment 4. *Panel b:* Data from rat in each group with best fitting response functions, and corresponding corrected response rates per trial, for Groups 60 and 480 of Experiment 4.

Figure 8. A scatter plot of the dynamic interval (unsmoothed data) for the fixed CS for all subjects.

Figure 1

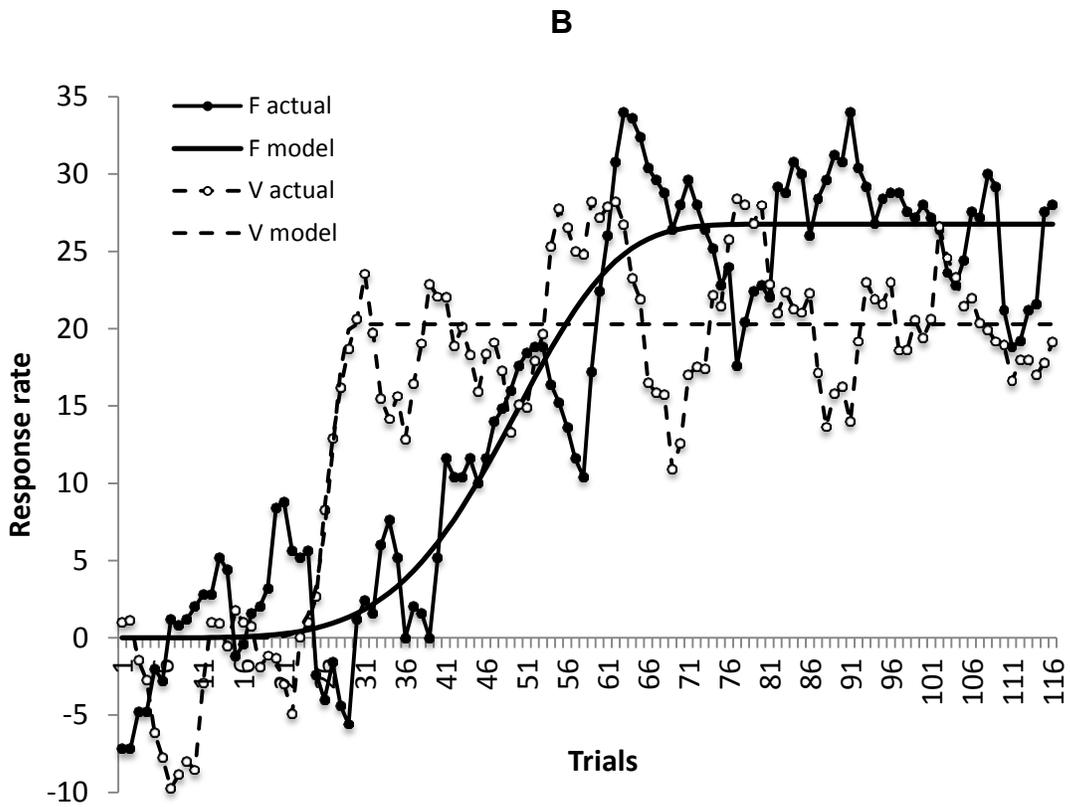
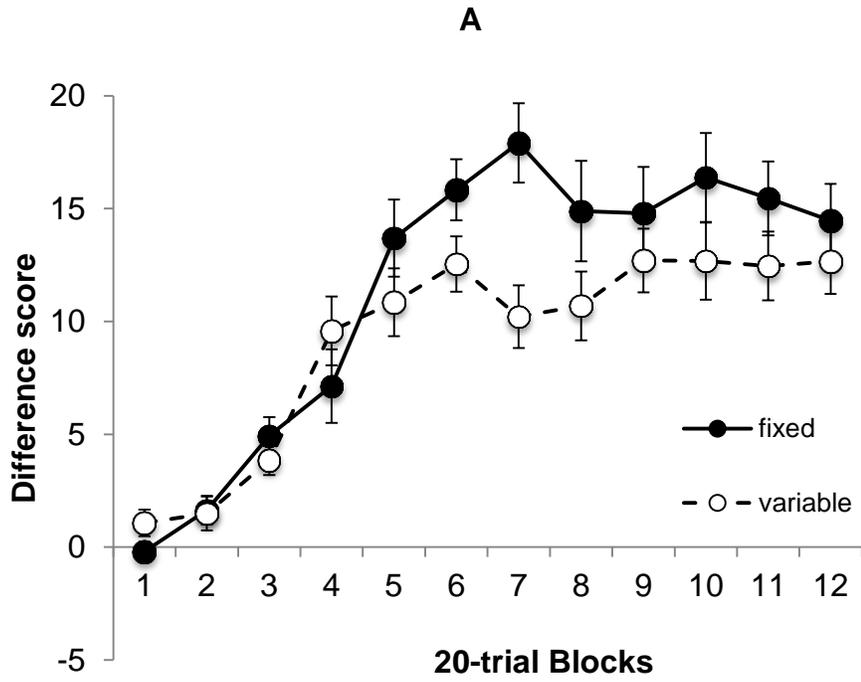


Figure 2

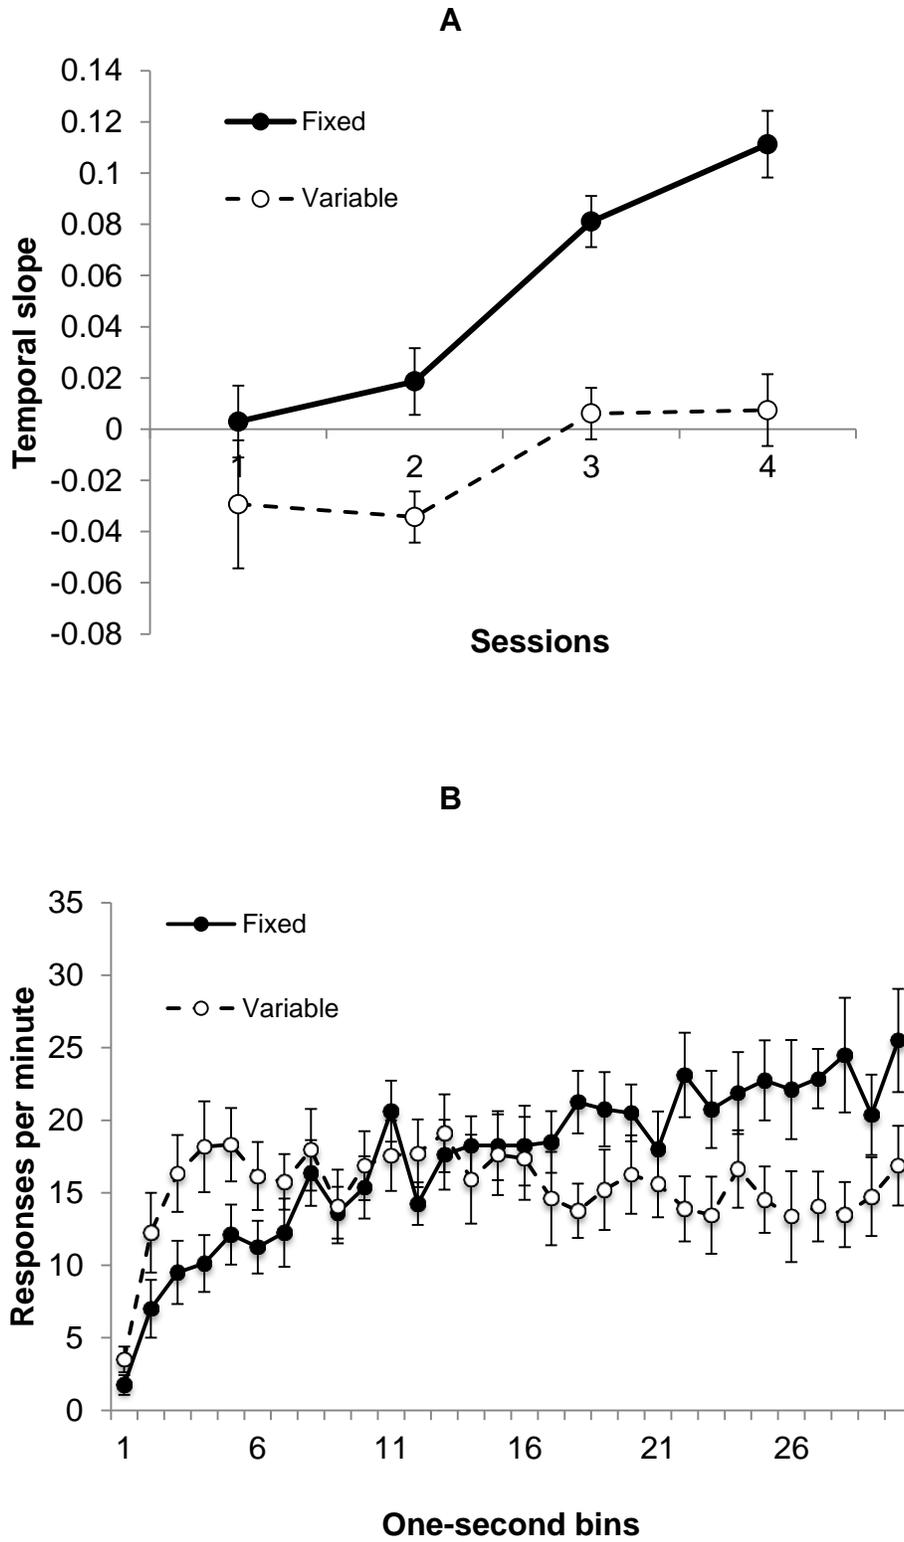
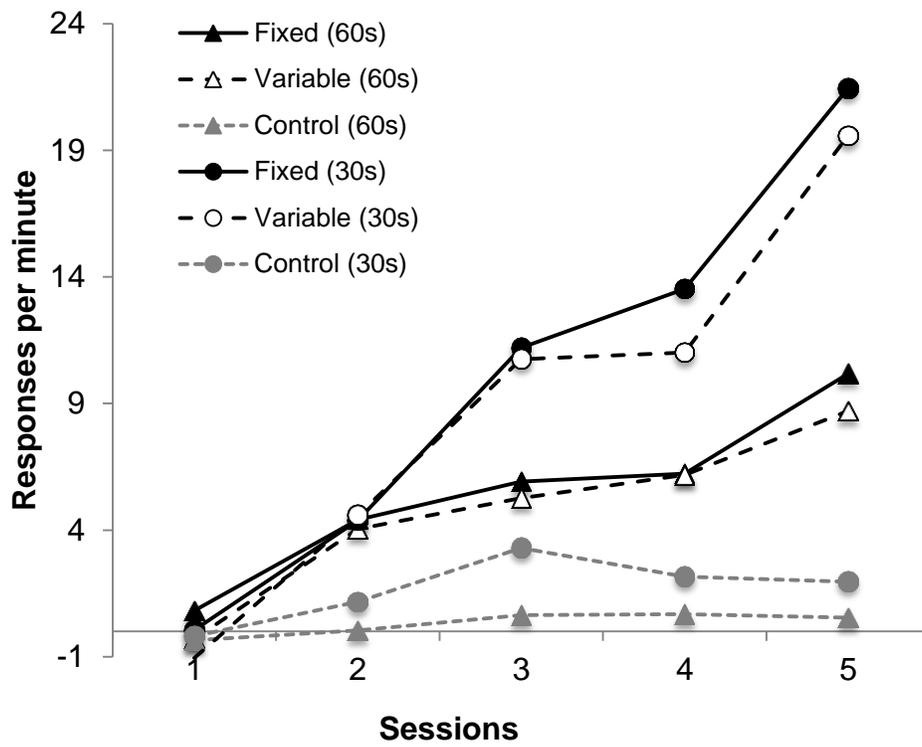
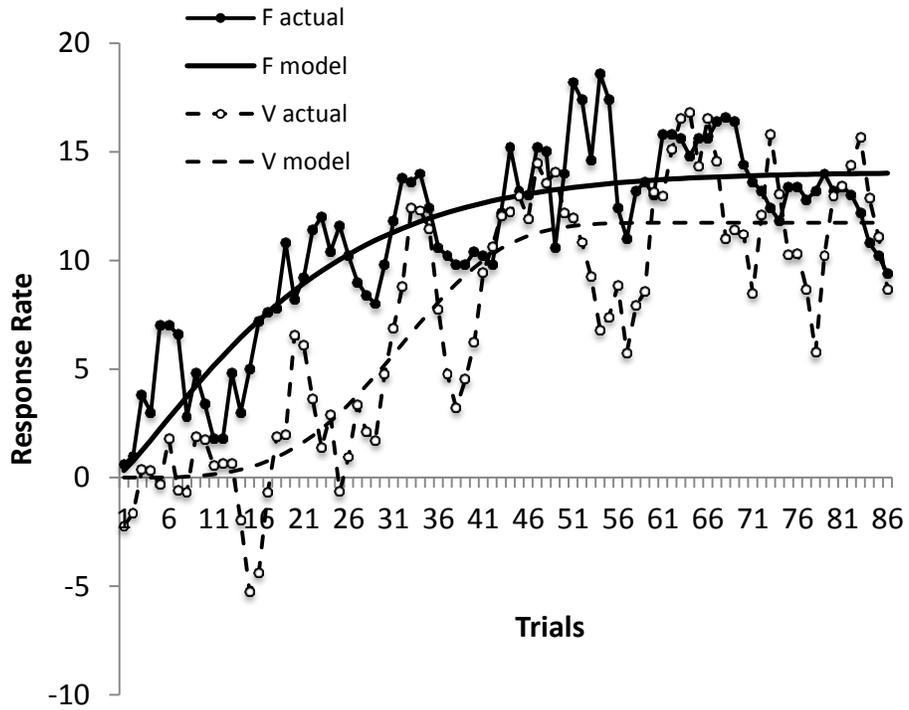


Figure 3

A



B



C

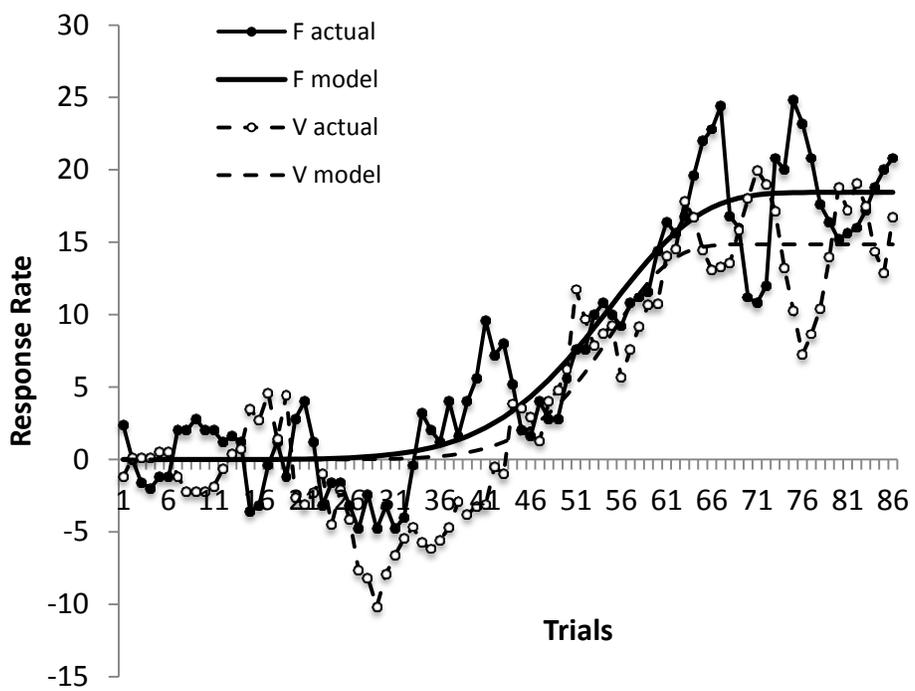
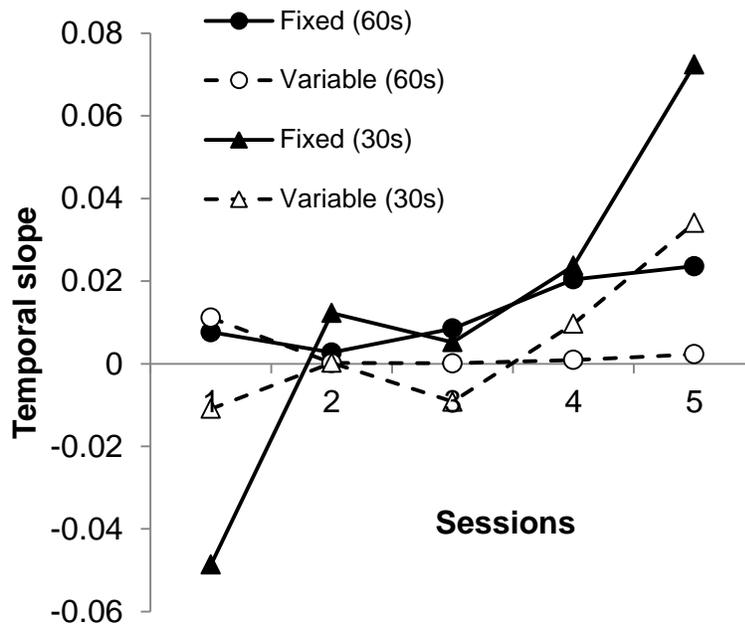


Figure 4

A



B

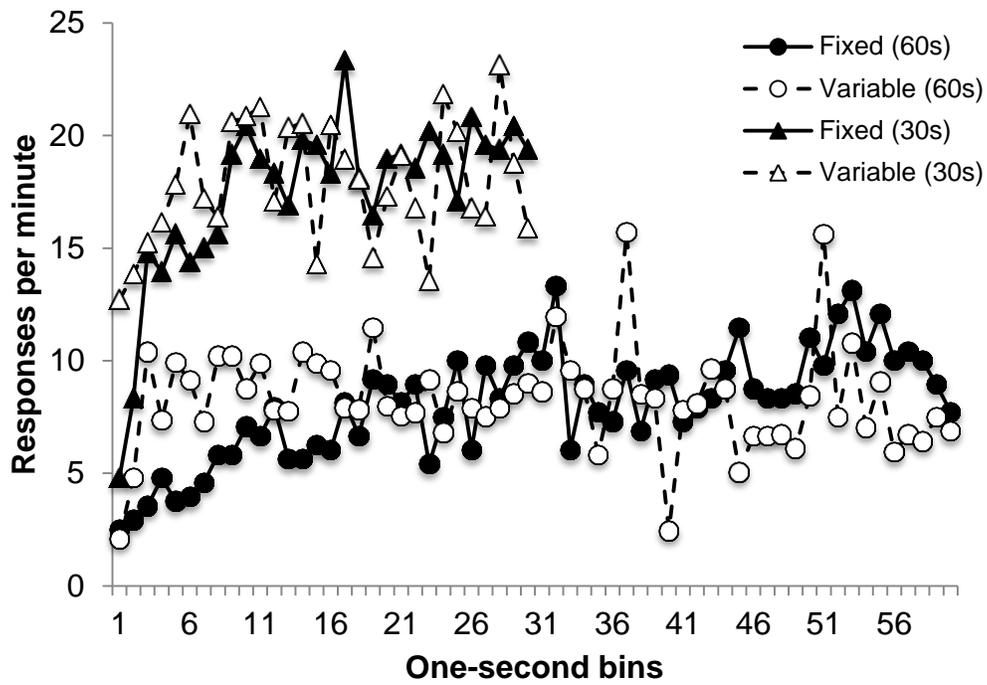
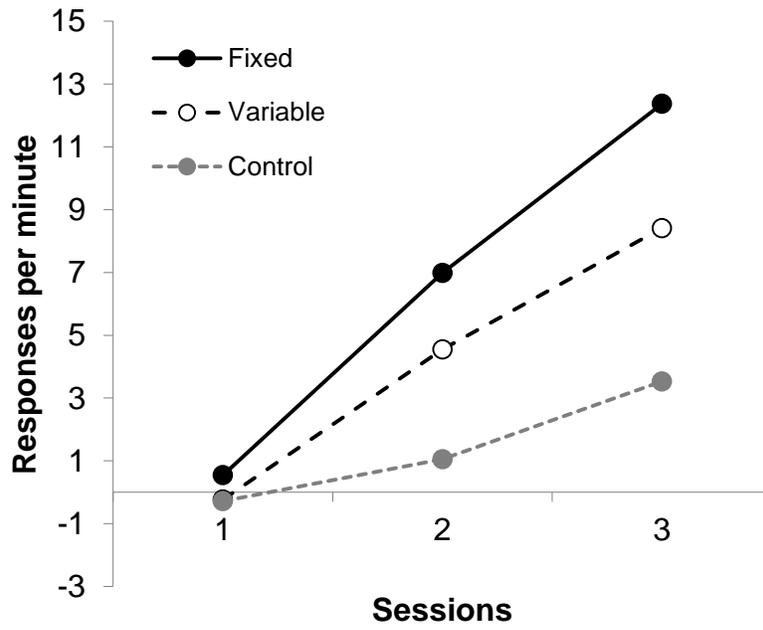


Figure 5

A



B

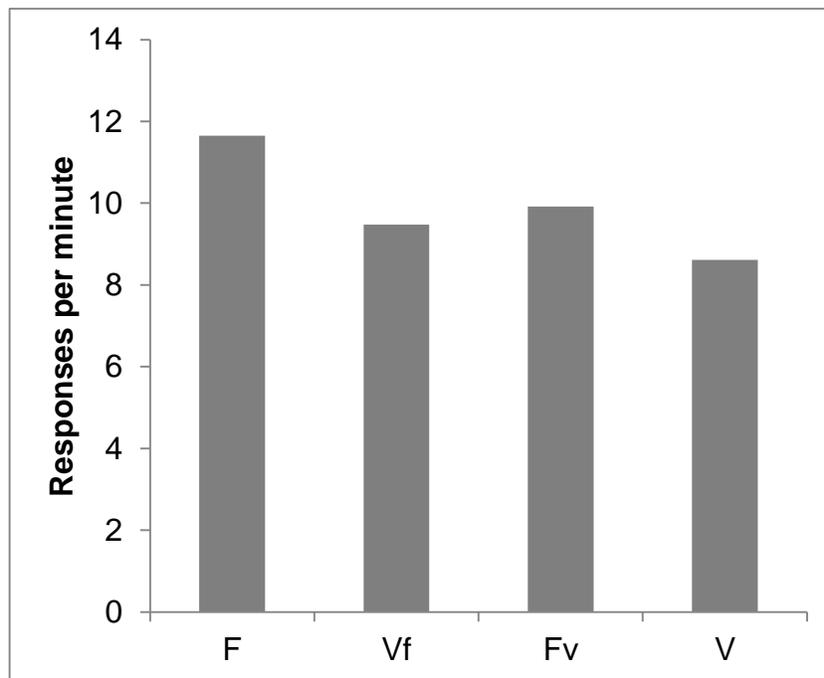
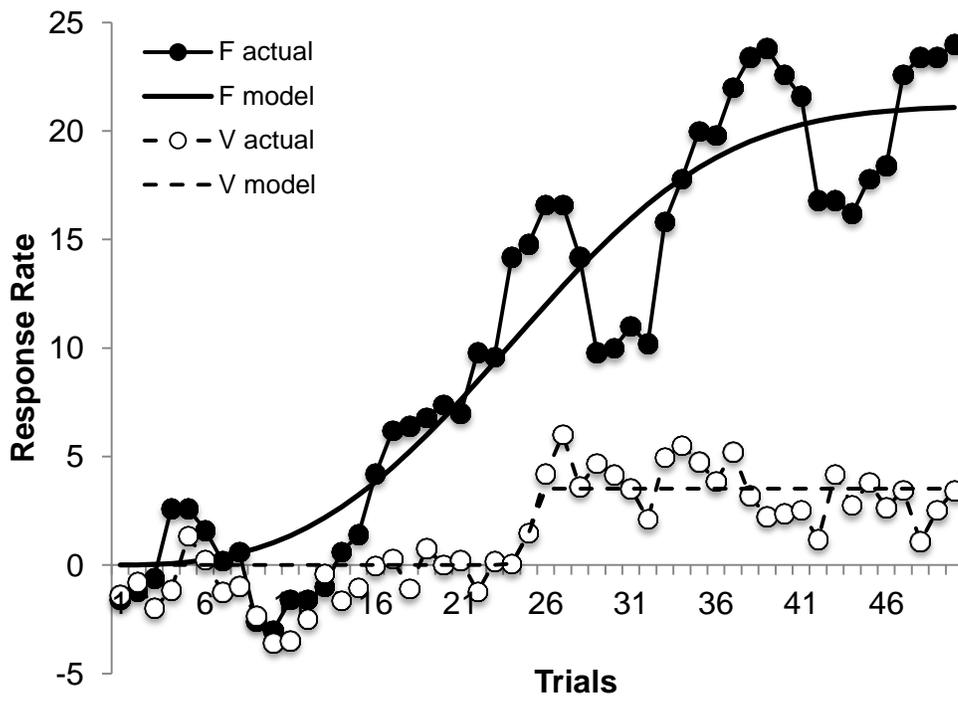
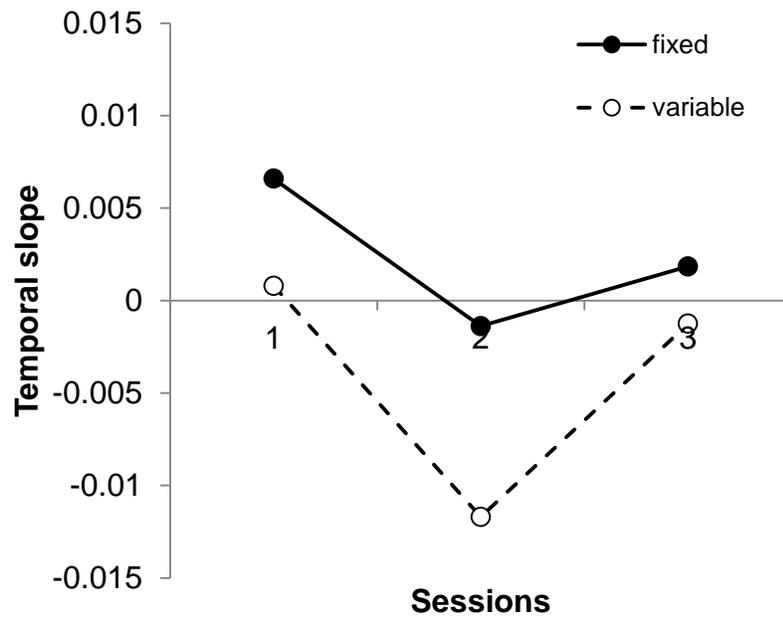


Figure 6

A



B



C

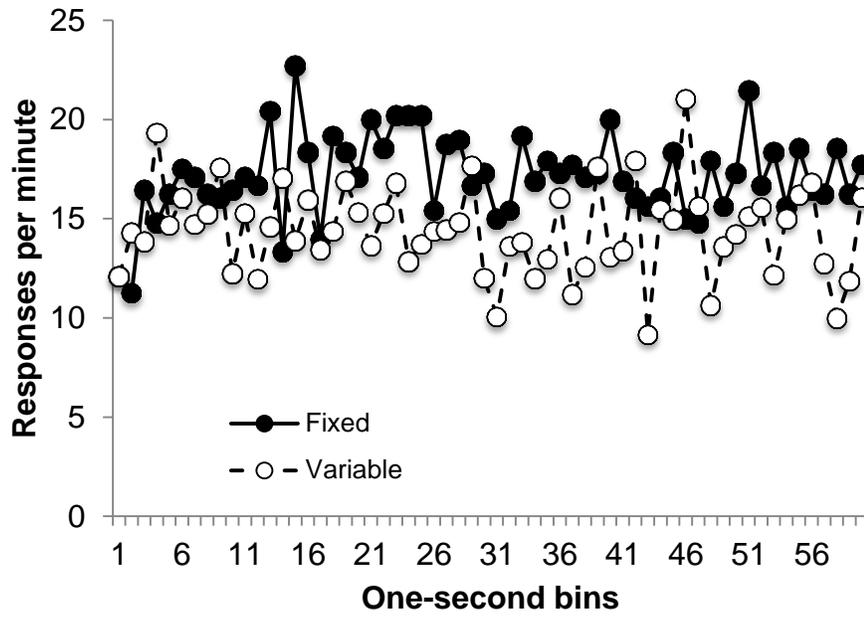
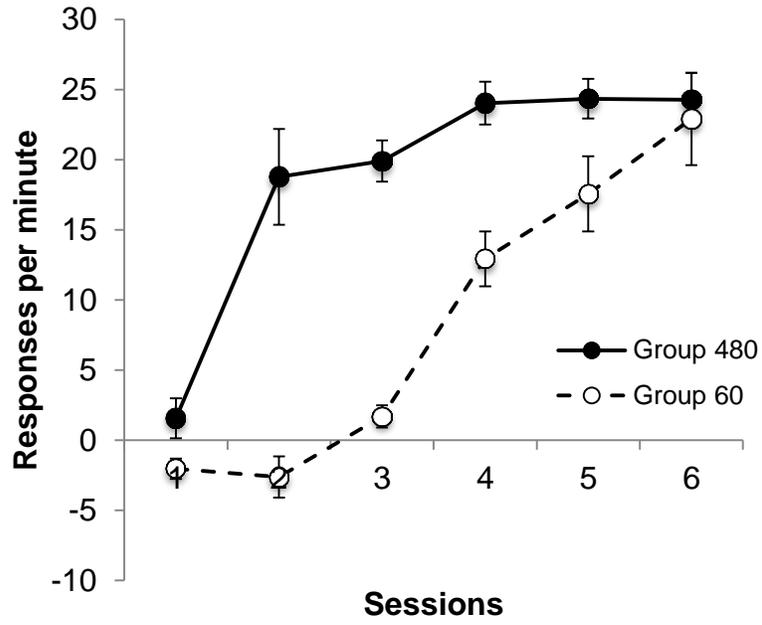


Figure 7

A



B

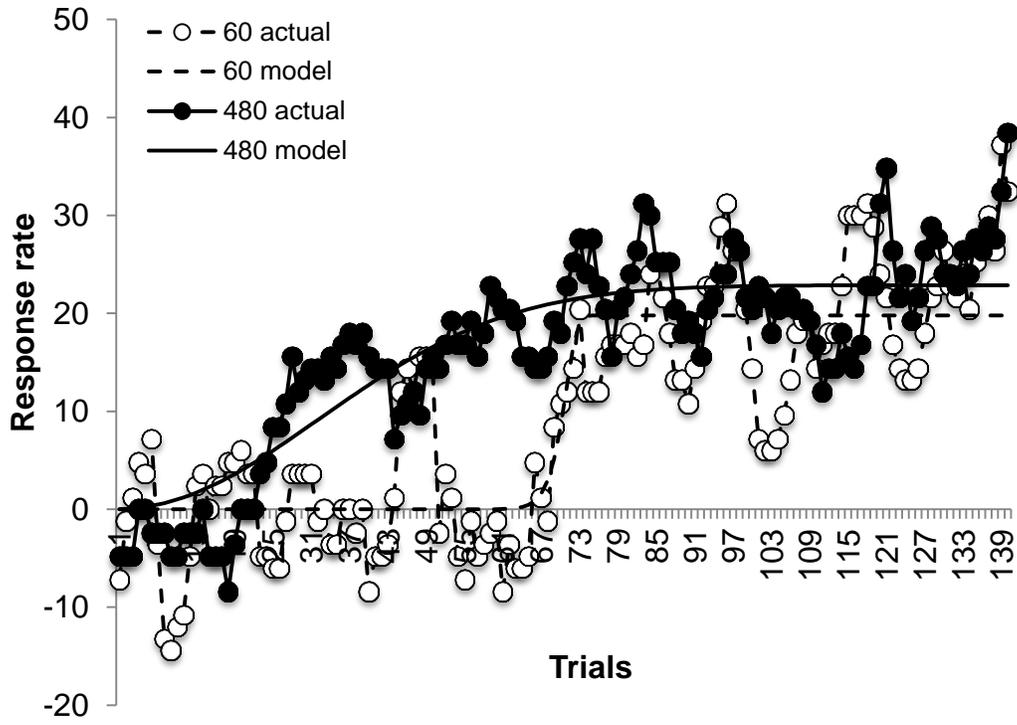


Figure 8

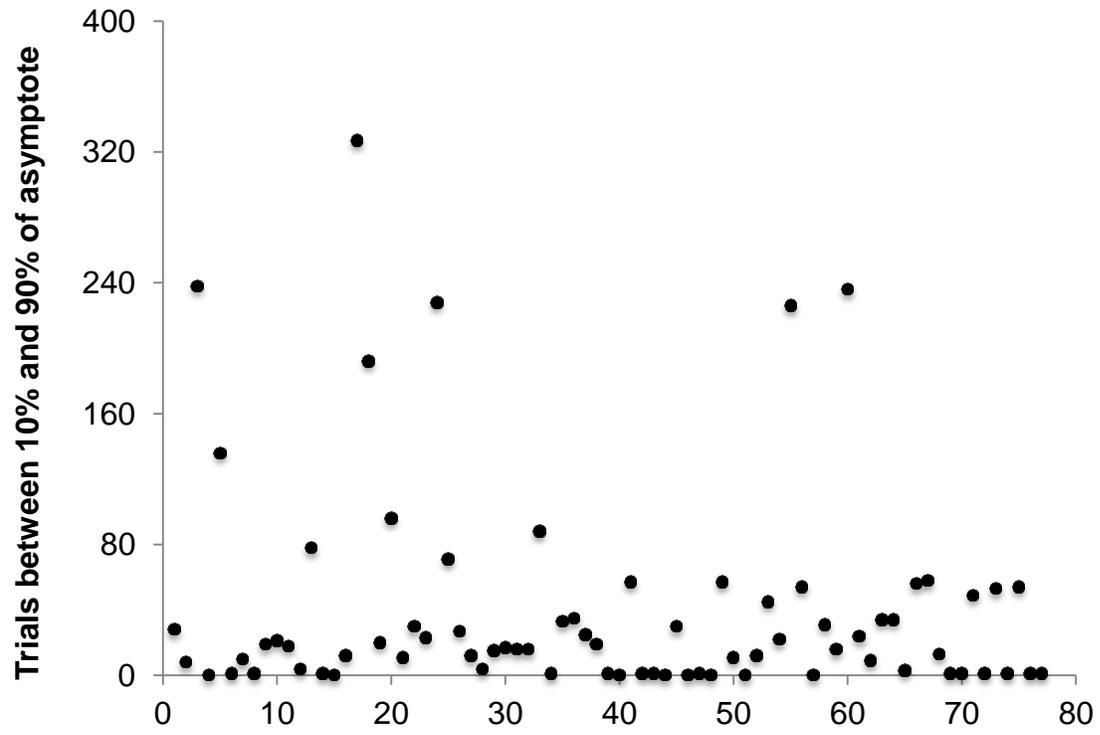


Table 1. Weibull parameters, mean fit and trials to 10%, and between 10% and 90% of asymptote, for fixed (F) and variable (V) conditions in the five experiments. * denotes statistically significant difference, bold indicates a numerical difference consistent with faster acquisition in the variable condition and higher asymptotic rates in the fixed condition.

	<i>Exp. 1</i>		<i>Exp. 2a 60s</i>		<i>Exp. 2a 30s</i>		<i>Exp. 3</i>	
	F	V	F	V	F	V	F	V
Asymptote	17.8	13.7 *	7.0	6.2	18.9	27.5	21.7	11.6
Latency	49.6	41.3	45.2	26.4	48.7	50.7	37.9	29.4
Slope	13.8	23.8	22.3	40.6	30.4	34.6	6.7	49.2
Onset Latency	24.8	21.0	15.6	19.8	20.6	29.8	19.5	20.6
Dynamic Interval	68.6	42.9	54.6	10.9 *	62.9	37.1*	32.6	15.9 *