

Accepted Manuscript

Title: Acquisition of Peak Responding: What is learned

Authors: Fuat Balci, Charles R. Gallistel, Brian D. Allen, Krystal M. Frank


PII: S0376-6357(08)00222-2
DOI: doi:10.1016/j.beproc.2008.09.010
Reference: BEPROC 1935

To appear in: *Behavioural Processes*

Received date: 16-5-2008
Revised date: 3-9-2008
Accepted date: 21-9-2008

Please cite this article as: Balci, F., Gallistel, C.R., Allen, B.D., Frank, K.M., Acquisition of Peak Responding: What is learned, *Behavioural Processes* (2008), doi:10.1016/j.beproc.2008.09.010

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Acquisition of Peak Responding: What is learned?

Fuat Balci^{1,2}, Charles R. Gallistel³, Brian D. Allen¹, Krystal M. Frank¹,

Jacqueline M. Gibson², Daniela Brunner^{1,4}

¹PsychoGenics Inc., Tarrytown, NY

²Princeton Neuroscience Institute, Princeton University

³Department of Psychology, Rutgers University, New Brunswick

⁴Department of Psychiatry, Columbia University

Running Head: Acquisition of Peak Responding

Correspondence Address:

Fuat Balci
Green Hall, 3-N-12,
Princeton University,
Princeton, NJ, 08544
Phone: 609-258-7511
Email: fbalci@princeton.edu

Acquisition of Peak Responding: What is learned?

Fuat Balci, Charles R. Gallistel, Brian D. Allen, Krystal Frank, Jackie Gibson,
Daniela Brunner

ABSTRACT

We investigated how the common measures of timing performance behaved in the course of training on the peak procedure in C3H mice. Following fixed interval (FI) pre-training, mice received 16 days of training in the peak procedure. The peak time and spread were derived from the average response rates while the start and stop times and their relative variability were derived from a single-trial analysis. Temporal precision (response spread) appeared to improve in the course of training. This apparent improvement in precision was, however, an averaging artifact; it was mediated by the staggered appearance of timed stops, rather than by the delayed occurrence of start times. Trial-by-trial analysis of the stop times for individual subjects revealed that stops appeared abruptly after three to five sessions and their timing did not change as training was prolonged. Start times and the precision of start and stop times were generally stable throughout training. Our results show that subjects do not gradually learn to time their start or stop of responding. Instead, they learn the duration of the FI, with robust temporal control over the start of the response; the control over the stop of response appears abruptly later.

Keywords: Mice, Peak Procedure, Interval Timing, Acquisition, Animal Cognition

INTRODUCTION

The extraction of temporal information from the conditioning experience is an essential aspect of the process that leads to conditioned behavior (Gallistel & Gibbon, 2000; Gibbon & Balsam, 1981; Honig, 1981; Savastano & Miller, 1998). Temporal intervals have been shown to inform conditioned behavior at the very beginning of the training when the conditioned response first appears (Balsam et al., 2002; Drew et al., 2005; Guilhardi and Church, 2005; Kahoe et al., 2008; Kirkpatrick and Church, 2000; Wynne and Staddon, 1988). In some cases, it has been possible to show that the temporal intervals were learned prior to the appearance of the conditioned response (Ohshima and Mauk, 2001). Temporal control over responding is observed after a single CS-US pairing in aversive conditioning paradigms (Bevins and Ayres, 1995; Davis et al., 1989).

In the peak procedure (Catania 1970), a widely used method of studying interval timing, subjects are typically trained on a fixed-interval (FI) schedule prior to the introduction of peak trials (e.g. Abner et al., 2001; Cheng et al., 2007a,b; Kaiser, 2008; Kirkpatrick-Steger et al., 1996; Matell et al., 2006; Matell and Portugal, 2007; Taylor et al., 2007). FI pre-training ensures a well-established memory representation for the delay of reinforcement by the time the probe (peak) trials are introduced. Despite an established temporal memory, post-peak responding changes during peak interval training (Abner et al., 2001; Kaiser, 2008; Kirkpatrick-Steger et al., 1996; Matell and Portugal, 2007). This has been attributed to extinction of post-peak responses (Machado, 1997). In this report, we present the molar and molecular (trial-by-trial) investigation of the behavioral transition that occurs with the introduction of peak trials following FI pre-training.

Responding in the peak procedure generates multiple measures, which are assumed to relate to different components of the interval timing apparatus. For instance, the trial time of the maximal response rate has been proposed to reflect the target temporal interval or maximal expectancy of reinforcement (Gibbon, 1977). On the other hand, variability in the time of maximal response rate is assumed to reflect noise in the subjective representation of reinforcement time (Gibbon, 1977). These measures are typically gathered from the average response rates. When averaged across trials, response rates increase and decrease smoothly around the target time, but when analyzed at the level of individual trials abrupt onsets (starts) and terminations (stops) are seen, which occur at different times on different trials. These abrupt starts and stops are proposed to represent the onset and offset of the behavioral translation of reinforcement expectancy (see Church et al., 1994). The normalized variability signatures (coefficient of variation) of these derivations have also been used as secondary molecular measures of temporal precision (Abner et al., 2001; Gallistel et al., 2004).

Few studies have characterized the changes in a subclass of these measures that appeared in the course of training. For instance, although lacking a molecular characterization, the presentation of response curves over the course of training in Abner et al. (2001 – in mice) and Kirkpartick-Steger et al. (1996 – in pigeons) suggests gradual changes, particularly in post peak responding. Kaiser (2008 – in rats) took a closer look at such changes in response curves in the course of training. He found that post-peak responding decreased in the course of training while pre-peak responding was stable. Consistent with this finding, Matell and Portugal (2007) found decreasing post-peak but stable pre-peak responding with extended training in rats.

To test the assumption that the gradual looking decrease in post-peak responding is mediated by extinction, Kaiser (2008) looked at the acquisition of peak responding with low vs. high proportions of peak trials. He assumed that if extinction underlay this behavioral transition, then a higher proportion of probe trials would result in earlier steady-state response curves. On the contrary, he found that a higher proportion of peak trials (50%) resulted in delayed emergence of steady-state performance when compared to a lower proportion of peak trials (10%), which led him to conclude that extinction did not govern decreasing post-peak responding.

In all of these studies, the behavioral transition observed in the course of peak interval training was studied at the level of group averages rather than individual subjects, which might be responsible for the gradual looking change (see Papachristos and Gallistel, 2006). Moreover, almost all of the investigations of the acquisition of peak responding lacked start-stop analysis, which would provide a more in-depth trial-by-trial assessment of the temporally controlled responses. Start-stop analysis was used in a study by Matell and Portugal (2007), however it only captured the effect of additional 27 training sessions (extended training) on peak responding that had already been acquired within six to eight training sessions (brief training).

In our study, we investigated how these critical measures of timing performance behaved in the course of the acquisition of timed responding by mice in the peak procedure using both average response and single-trial analysis. Our characterization of the behavioral transition was conducted both at the level of group averages and individual subjects.

METHOD

Subjects

The subjects were 45 experimentally-naïve male C3H mice (Taconic, Germantown, NY), aged eight weeks at the start of the study. They were maintained at 85% of their ad libitum body weights by supplementary feedings following sessions during weekdays and at the session time during weekends. Water was continuously available in the home cages. The mice were kept on a 12/12 light/dark cycle and all tests were conducted during the light period of the cycle. Mice were treated in accordance with the Guide for the Animal Care and Use of Laboratory Animals (NRC 1996), and all procedures were approved by the Institutional Animal Care and Use Committee at PsychoGenics, Inc.

Materials

The mice were tested in standard mouse operant chambers (MED Associates, Burlington, VT; Model: ENV-307A) that contained two retractable levers on either side of a food trough along one wall. A small cue light was mounted above each lever. Only one of the levers was used. Reinforcement was 0.01 cc of evaporated milk (CarnationTM) delivered via dipper through the receptacle. The operant chambers were located in sound-attenuating chambers and controlled by the Med-PC IV software package. The fan was on throughout the session.

Procedure

Behavioral testing. Behavioral testing was conducted daily Monday to Friday. The mice were first trained to press the levers, using an alternative FR-1 FT-60s schedules, in which reinforcement was delivered upon the completion of either schedule, whichever was met first. In this phase, the session continued until a mouse received 64

reinforcers. Almost all mice readily learned to press the lever under these conditions within five sessions. They were then trained for up to ten days with a fixed-interval (FI) schedule, wherein only the first response after the criterion duration was reinforced. During the first five training sessions, the mice were trained on a FI-10s schedule, while during the second five training sessions, they were trained on a FI-30s schedule.

They were then exposed to the peak procedure. The peak procedure consisted of 45 trials of two types: fixed-interval (30 trials) and peak trials (15 trials). Each trial started with the extension of the retractable lever and the illumination of the small cue light above the lever. On fixed-interval trials, the first response after 30 s was reinforced with two seconds of access to evaporated milk. After this period, the lever was retracted and the cue light was extinguished. On peak trials, which were randomly interspersed, the lever remained extended and small cue light stayed on for 90 s, and there was no reinforcement. Trials were separated by a fixed ITI of 20 ± 10 s (uniformly distributed). Each session lasted approximately 80 minutes and training lasted for 16 sessions.

DATA ANALYSIS

We used two relatively novel methods of data analysis. One was based on average response rates (the average response analysis), whereas the other involved trial-by-trial estimates of start and stop times (the change-point analysis; also see Church et al, 1994). These different techniques were used to eliminate the analytical interdependence of two different sets of measures; one set composed of the peak time and spread, and the other of start and stop times. In the first instance, we used smoothed cross-trial average response rates to estimate peak and spread. In the second instance, when peak and spread are derived from trial-by-trial analysis, they depend on the more basic start and stop values

($spread = stop - start$ and $peak = (start + stop) / 2$), which we estimated using a change-point algorithm.

Average response analysis: This analysis was conducted on individual subject's session-based data (averaged across trials), which, as expected, was noisier than the response rates averaged across subjects and sessions. Thus, we calculated a running average smoothing of the response rates as a function of time into the trial (with span of 19 s). This smoothing made the response rate data unimodal. The mode of the smoothed response rate data was our estimate of the peak time. The only constraint in this analysis was that the peak time not be allowed to occur during the first and last five seconds of the trial (which occurred in around 3% of the trials). This eliminated spurious results caused by early responding due to CS onset, and late responding due to anticipation of the next trial (or the end of current trial, see Church et al 1991). The start time was defined as the point at which the smoothed data first exceeded 70% of maximum response rate and the stop time was the time at which it first fell below 70% of maximum response rate. If the response rate decreased below 50% of maximum response rate at any point prior to the peak time, the start time was the point at which the smoothed data first exceeded 70% of maximum response rate after that point. This again eliminated the effect of early bursts of responding. Figure 1 shows the normalized response rates as a function of trial time and amount of training.

Single trial change-point analysis: Church et al (1994) showed that responding in individual peak trials was not a smooth function as suggested by average response rates; rather, it followed a "break-run-break" pattern. In order to capture the temporal limits of the "run" period, they developed an algorithm that assumed a single onset and offset.

However, response rate was reported to vary within the “run” phase; response rates were slightly lower at the beginning and end of the run period compared to its middle (Church et al 1994; Cheng and Westwood, 1993). Thus, we used a method that searched for statistically significant¹ changes in the response rates (see also Taylor et al. 2007). This method estimated the start and stop times on individual trials, using a relative-likelihood change-point algorithm (also see Taylor et al. 2007 for a variant).

The algorithm assumes that the onset and offset of responding are changes in the expectation of a random rate (Poisson) process. It proceeds through the sequence of interresponse intervals on a given trial datum by datum, computing for each successive datum the relative likelihood (Bayes Factor) of two different models for the sequence: the no-change model and the change model. The no-change model assumes that all of the intervals up to and including the current datum come from the same exponential distribution. The change model assumes that the intervals up to some earlier point in the sequence (the putative change point) come from one exponential distribution, while the intervals after that point come from an exponential distribution with a different expectation.

The putative change point is the point that maximizes the likelihood of the change model. Because it has an extra free parameter, the change model is always more likely than the no-change model. The Schwartz Criterion (in this case, half the log of the number of data entering into the likelihood computations) corrects for this bogus gain in

¹ “Significant” has a slightly different meaning from its use in traditional statistical hypothesis testing, because the change-point computation uses Bayes Factors rather than p values. Its meaning here is that the Bayes Factor (odds in favor of a change) exceeded our decision criterion, which is analogous to saying that the p value was less than the alpha level chosen by the researcher.

relative likelihood (Schwartz, 1978). The algorithm is recursive: when the Bayes Factor (odds) in favor of the change model exceeds a user-specified decision criterion, the data are truncated at the putative change point, and the algorithm begins anew with the first datum after the point of truncation.

In our analysis of individual trials, we used decision criteria ranging from; 10 (liberal), to 100 (conservative), to 1000 (extremely conservative). For a criterion of 100, the odds had to favor the change hypothesis by 100:1 in order for the algorithm to decide that there has been a change in the expectation of the distribution from which the measures were drawn. The first reliable positive and first negative change points were defined as start and stop points. The only constraint used in this calculation was that stop points were not allowed to occur earlier than the initial 15 s of the trial. This constraint, as in the previous analysis, was used to eliminate the effect of early bursts of responding due to CS onset (non-temporally controlled), a case observed in some trials. These two methods, namely the average response analysis and single trial change point analysis, were used to determine different aspects of the peak responding (i.e., peak time, response spread, start times, and stop times).

Analysis of Acquisition:

We then checked how the output of these methods changed over the course of training. We initially assessed the relation between timing performance and the number of training sessions by log-log regression. The regression was run for each individual subject's data (averaged across the trials in a session) and for the data averaged across subjects (per session). We present the analysis of average values (see Figure 2A) for only

those start and stop times that were detected by criterion 100. The results gathered from start and stop times detected by criteria 10 and 1000 were identical to these.

We then asked when the temporal control of stopping appeared in the course of training. We first classified trials in two classes, those without a stop (failed trial) and those with a well-defined stop (success trial). The failed trials were those in which the algorithm failed to find a stop time or the stop time was longer than 80 s. The success trials were those in which the algorithm found a stop time, which was shorter than 80 s. We then proceeded to analyze the sequence of failures and successes as a function of trial number. To determine if there was a change in the frequency of trials with a controlled stop (successes), we applied the change-point algorithm that was applied to the individual trials. Whether in the single trial analysis, the algorithm did or did not find a stop, depended to some extent on the decision criterion. Thus, the sequence of failures and successes to which we now applied the algorithm (the second phase use of the algorithm) depended on the criterion used in the single trial analysis (the first phase use). For that reason, we repeated the first-phase analysis using three very different decision criteria, therefore obtaining somewhat different sequences of failures and successes, and we repeated the second-phase analysis on each of those sequences.

When the change-point algorithm is applied to a binary sequence, the sequence is assumed to have been drawn from a Bernoulli distribution, the “coin-flip” distribution in which there are only two possibilities, 0 and 1, and the only parameter is p , the probability of drawing a 1. The algorithm again proceeds datum-by-binary-datum, asking at each point in its progress, whether the binary sequence up to and including that point is better represented by a model in which there has been no change in p or by a model in

which there has been a change in p (at some earlier point in the sequence). Again, whether the algorithm does or does not find a change in the underlying probability of a well-defined stop depends on the decision criterion adopted. Here, we also used three very different criteria. In both the first and second phases, our three decision criteria were odds of 10:1 (liberal), 100:1 (conservative), and 1,000:1 (extremely conservative). Crossing the three criteria used in the first phase (single-trial analysis) with the same three used again in the second phase gives nine repetitions of the analysis for each subject. What varies from analysis to analysis is the combination of decision criteria used in the first and second phases (see Table 1).

Finally, linear regression was used to test whether the timing of the well-defined stops did or did not depend on the trial number for each subject. An alpha level of .05 was selected for all null hypothesis testing. Before presenting the results, we stress again that all peak times and response spread measures were gathered from the average response analysis and all start and stop times (and their relative variability) were gathered from single-trial change-point analysis.

RESULTS

Figure 1 shows the average normalized response rate as a function of elapsed time in a trial on day 1, 4, 7, 10, 13, & 16 (every third session). The average normalized response rate data suggest that the spread got narrower as training progressed, which could be interpreted as an increase in the subject's temporal precision over the course of training. Figure 1 further suggests that the gradual narrowing of the spread was mediated by the gradual decrease in the stop times over the course of training, with little change in the start responses.

On the other hand, the apparent gradual increase in the temporal precision suggested in Figure 1, that is, the decrease in the spread, might be an artifact of averaging across subjects. Different subjects may have abruptly learned to stop responding at different points during the training. When averaged across subjects, this would produce the appearance of a gradual change in the stop times and narrowing of the spread.

Consistent with the gradual change interpretation, when the temporal performance was represented as the average across subjects (per session), the spread and stop times appeared to change gradually over the course of training (see Figure 2A). The number of training sessions was a significant predictor of the spread, $R^2 = .93$, $F(1,14) = 188.78$, $p < .0001$, the spread normalized by the peak time, $R^2 = .51$, $F(1,14) = 14.55$, $p < .01$ and the stop times, $R^2 = .95$, $F(1,14) = 251.9$, $p < .0001$. These results were further complemented by the results of the same analysis applied to individual subjects (see Figure 2B). For 75% of the subjects there was also a decrease in the normalized spread but this relation reached significance in only 33% of those cases.

On the other hand, scatter plots of the stop times (see Figure 3 left panel for 4 representative subjects) revealed a seemingly abrupt transition from a pre-acquisition phase, during which trials with a well-defined stop (a stop detectable by the change point algorithm) occurred infrequently or not at all, to a post-acquisition phase, in which most trials had a well-defined stop. In the scatter plots in Figure 3, the trials when there was no well-defined stop time are represented by the points at the upper limit (90 s). In each plot, there is a clear point, somewhere around trial 50, at which these points become much less dense and the majority of points then fall below this upper limit. In the scatter of points below the upper limit (that is, in the well-defined stops), there is no clear trend toward

earlier stops as training progresses. The abrupt transition from a pre-acquisition phase to a post-acquisition phase is also reflected in the individual subject's average response rates from these phases (see Figure 3 right panel).

We confirmed these impressions using the change-point analysis and linear regression on well-defined stops. Table 1 shows the median training sessions and trials to acquisition, defined as the session in which well-defined stops appeared and the trial after which they became frequent, respectively. These values essentially show how many training sessions/trials were required prior to the abrupt appearance of frequent well-defined stops. Table 1 also shows the median frequency of well-defined stop trials immediately following their abrupt appearance. This value gives the frequency of well-defined stop times in the interval immediately after the change in their frequency.

Finally, Table 1 shows the median percentage of the frequency of stops immediately following their abrupt appearance to the mean frequency of trials with a well-defined stop during the last half of the trials. This value indicates whether there was a change in the frequency of well-defined stop-times after their abrupt appearance. A value of one for this index means that the frequency of well-defined stops during the last half of the trials was the same as in the interval immediately following their abrupt appearance. These summary statistics depend to some extent on the combination of decision criteria used in the two phases of change-point analysis, which is why each statistic appears nine times in Table 1.

Change-point analysis revealed the abrupt appearance of frequent well-defined stops between the third (30th trial) and the fifth training session (60th trial), depending on how conservative the decision criteria employed were. The frequency of trials with a

well-defined stop was fixed after the first change point: the change in the number of such trials was abrupt in that their median frequency, immediately following the first change point, was usually 86 to 99 % of their mean frequency during the last half of the trials that came after the change point.

Subject-by-subject linear regressions of the well-defined stop times on training trials showed no systematic dependence of stop time on the number of training trials: The slopes of these regression lines were distributed about a mean of $-.01$, (standard deviation (std) = $.05$), $-.02$ (std = $.1$), and $-.02$ (std = $.16$) - depending on the how conservative the decision criterion for single trial analysis was. Following the first change point the median frequency of trials with a well-defined stop ranged between 47 & 81 % of the trials (see Table 1). Usually a large portion of the trials following the first change-point contained a well-defined stop time. What Table 1 shows is that the dependence of these summary statistics on the combination of decision criteria used in the two phases of change-point analysis is weak; no matter which combination one focuses on the conclusion is the same.

Thus, we conclude that the apparent gradual narrowing of the spread via apparent gradual decrease in stop times was due to the abrupt appearance of timed response terminations by individual subjects at different points during training, a conclusion that would not follow from a consideration of the average response rates. In order to have a better characterization of changes in peak responding, we also looked at other timing measures.

Temporal Accuracy: The number of training sessions was a significant predictor of the absolute temporal distance between the peak time and scheduled fixed interval averaged

across subjects per session, ($R^2 = .86$, $F(1,14) = 83.33$, $p < .0001$; see Figure 2A).

However, when the correlation was done on individual data, no significant trend was found (see Figure 2B). It is possible that decreasing temporal distance was mediated by reinforcement times that were experienced closer to FI as training progressed; however, the number of training sessions was not a significant predictor of the average experienced reinforcement time, $R^2 = .22$, $F(1,14) = 3.96$, $p > .07$ (see Figure 2).

Start Times: There was no apparent change in the timing of the beginning of responding (starts) over the course of training. Supporting this statement, the number of training sessions was not a significant predictor of either the average start times, $R^2 = .39$, $F(1,14) = 1.57$, $p < .23$ (see Figure 2A) or the individual starts times (see Figure 2B).

Start and Stop Time CV (Coefficient of Variation): Another set of temporal precision measures that cannot be directly observed from the average response rates in Figure 1 is the CV of start and stop times. In all of the analyses conducted on the Stop CVs, we only looked at the data gathered from the last 10 sessions of training. This was done because a meaningful estimate of variability of stop times could only be gathered from well-defined stops. We found that the CV of well-defined start and stop times did not decrease as the training progressed. Session was not a significant predictor of the average start CVs ($R^2 = .05$, $F(1,14) = .66$, $p < .43$), or the average stop CVs ($R^2 = .15$, $F(1,8) = 1.38$, $p < .28$; see Figure 2A). This was confirmed with the analysis of individual data (see Figure 2B). If anything, the average CV of stop times increased as the training progressed. Figure 2B summarizes the output of individual log-log regressions for all measures of temporal performance in terms of the percentage of subjects.

In all test days, the CV of the stop times was significantly lower than CV of the start times (all p s < .0001). There was a significant positive correlation between start and stop CVs when they were averaged across subjects for each session ($R^2 = .53, p < .05$, comparison done for the last 10 sessions of training). There was no significant correlation between start and stop times when they were averaged across subjects for each day ($R^2 = .02, p < .59$).

FI-30s pretraining: Finally, we checked if the start times changed during FI training. The number of FI-pretraining sessions was not a significant predictor of the average start times, $R^2 = .10, F(1,3) = .35, p < .6$ (see Figure 2A). This was confirmed with the analysis of individual data.

DISCUSSION

In this study, we investigated the acquisition of timed responding in the peak procedure by mice, following FI pretraining. As previously reported (i.e. Kaiser, 2008) temporal precision appeared to increase gradually in the course of training, that is, response spread decreased. The decrease in spread was mainly mediated by the earlier termination of responding (stop times) rather than by delayed start times (also see Matell and Portugal, 2007). The earlier stop times, however, appeared to be a result of a threshold-like abrupt acquisition process. Once they emerged, these response terminations (stops) were under temporal control, the parameters of which were stable throughout the remaining trials. Based on these findings, we conclude that the subjects simply learned to stop responding in the peak trials in the course of training rather than learning to stop earlier in the trial as would be suggested by average response rates. When generalized to the underlying processes, this suggests that what would be acquired and

manifested throughout the peak-interval training (following FI pretraining) would be the decision to stop responding when the target time passed without reinforcement.

Abrupt behavioral transitions of timed responses have also been demonstrated in the context of experimental manipulations of temporal parameters (e.g. Guilhardi and Church, 2005; Meck et al., 1984; Wynne and Staddon, 1988). For instance, Meck et al. (1984) demonstrated the abrupt nature of re-acquisition of peak times when the scheduled fixed interval in the peak procedure was changed from one value to another (10 s \rightarrow 20 s and 20 s \rightarrow 10 s). Specifically, they found that the rats adjusted to the new temporal criterion (peak time) in two abrupt steps, a finding not manifested in the average peak times, but rather in the individual data. During the intermediate state, subjects aimed at the geometric mean of the initial and novel criteria, while in the final state they aimed for the novel criterion, itself. However, it is possible that the intermediate state might be an artifact of smoothing function and the same individual data sets can be explained also by a single abrupt change in the peak value (no smoothing function was applied on our data sets that were evaluated for change points).

Our finding on the abrupt emergence of stops is consistent with the abrupt acquisition of anticipatory autoshaped head poking in individual mice (see Papachristos and Gallistel, 2004). Our finding constitutes a generalization of this phenomenon to the acquisition of more complex responses such as the temporal controlled ones. These findings give an important message, that is, group-averaged learning curves do not represent acquisition in individual subjects and thus should be interpreted with care in trying to get at the neurobiological underpinnings of learning in the context of both simple and complex behaviors.

Given the abrupt emergence of response terminations in our data, the gradual looking asymmetrical decrease in response spread throughout peak interval training (prior to steady state peak responding) should not be treated as the behavioral manifestation of increasing temporal precision. Indeed, given earlier reports that show the acquisition of timing very early in the training (e.g. Balsam et al., 2002) one would rather expect the temporal parameters to be established during FI pre-training. In cases where peak interval training follows this FI pre-training, the complete behavioral manifestation of this temporal memory might be blocked by other processes such as frustration-induced perseverative responding or perseveration due to the lack of change detection in the state of the world. Whatever these processes are, they appear to be facilitated by the sequential nature of FI and peak-interval training.

This interpretation is supported by Drew et al. (2005) where they investigated the acquisition of timed responding in the peak procedure with goldfish in an aversive conditioning task. In this experiment, peak trials were intermixed with the FI trials from the beginning of the training. They found that temporal precision and accuracy measures did not change after the subject acquired the conditioned response, that is, peak responding was under temporal control from the start of training. The introduction of peak trials from the very start in that study resulted in the behavioral manifestation of stop thresholds early in the training, which is directly supported by stable post-acquisition stop times.

An alternative explanation for the gradual looking decrease in latent portions of average response rates with peak training is extinction, where behavioral states that follow scheduled FI in peak trials are assumed to get inhibited and thus expressed less

with peak interval training (see Kaiser, 2008; Machado, 1997). However, the abrupt occurrence of well-timed response terminations in our study combined with the earlier occurrence of steady-state peak responding with relatively lower proportion of peak trials (see Kaiser, 2008) suggest that extinction or a habituation like process is not likely to mediate these behavioral transitions.

The strong transition in the response terminations was not accompanied by changes in the start times and CV of starts and stops (also see Matell and Portugal, 2007), which suggests that in general temporal control over many aspects of peak responding is robust to changes in the experimental parameters that are irrelevant to the temporal properties of the critical event (but see Ludvig et al., 2007). Supporting this interpretation, in the overshadowing paradigm, timing of stimuli (both target and overshadowing) were also found to be insensitive to overshadowing, itself (Jennings et al., 2007).

Interestingly however, temporal accuracy appeared to increase in the course of training, that is, the difference (Peak Time – FI) decreased. Since the scale of change in the experienced reinforcement time was very small in relation to the scale of the observed change in “Peak Time – FI” (see Figure 2), the changes in the experienced reinforcement time cannot explain this increase in temporal accuracy. However, the apparent decrease in “Peak Time – FI” was mainly due to the lack of response terminations that delayed the time of maximum response rate at the very beginning of peak training. Basically, at the start of training the veridical peak time appeared to be masked by the asymptotic post-FI responding. Overall, these results show that starting from the very beginning of peak training, anticipatory responses are in general under cognitive control of well-established

temporal memory and what appears to be acquisition of timing is essentially due to lack of temporal control rather than lack of well-established temporal memory.

Consistent with previous reports, the CV of start times was consistently larger than the CV of stop times (also see Cheng et al., 1993; Church et al., 1994; Gallistel et al., 2004; Matell and Portugal, 2007). Matell & Portugal showed that at least in rats the higher CV of start times was mainly due to the impulsive responding that resulted in a bimodal distribution of start times (but see Brunner et al., 1996). The distribution of start times however did not reveal any form of bimodality in our data set, even when it was evaluated separately for the first and second half of training sessions.

The lack of correlation between start and stop times (also see Matell and Portugal, 2007 for the lack of such relation with extended training but see Cheng et al., 1993; Church et al., 1994; Gibbon and Church, 1990) and their differential variability signatures further suggest that these behaviors might be mediated by different decision processes (Cheng and Westwood, 1993; Church et al., 1994; Gallistel et al., 2004; Matell and Portugal, 2007), while the trend for a positive correlation in their variability suggests that they might also be affected by a common variability source.

Finally, with this work we used relatively novel tools for an analytical approach to timed responses observed in peak procedure. The use of change point single trial analysis as a statistical approach to detecting start and stop times (also see Taylor et al., 2007) has an advantage over the exhaustive search algorithms (e.g. Church et al., 1994) since “break-run-break” pattern of responding is not necessarily binary as it is assumed by those algorithms. Change-point analysis makes no assumptions about the form of the function that best described the evolution of the response rate throughout a trial. Change

point analysis also allows the multiple user-defined threshold values with different levels of conservativeness. The use of non-model based average response analysis has a particular advantage over Gaussian fits in our data set, since early in training the response curves have clearly different shapes compared to their steady state Gaussian looking nature. Also, this analysis is robust to bursts of responding due to CS onset and the anticipation of the next trial or the end of the current trial (Church et al., 1991).

Considered together, our results suggest that the subjects established a temporal memory during FI pre-training and relied on this memory throughout training; the accuracy and precision in the memory representation did not change. After roughly 50 probe or peak trials subjects abruptly began to inhibit responding when the expected time of reinforcement had elapsed without reinforcement. This inhibition is under the control of the previously established memory for the time of reinforcement. The apparent gradual decrease in temporal precision and accuracy was a result of lack temporal control on the responses following scheduled FI rather than lack of a well-established temporal memory.

Our results further suggest that the investigation of attentional and/or memory enhancement/disruption through temporal information processing, might require at least 15 daily sessions of training prior to behavioral or pharmacological challenge; in order to reach at the steady state timing performance in peak procedure (e.g. spread). Any decrease in the spread measure through “earlier stop times” (as a result of manipulation) prior to the steady state performance might on the other hand reflect enhancement in processes other than memory and/or attention such as disappearance of perseveration.

In our experiment we used a fixed peak trial time, which was three times longer than the FI trial. Although, earlier reports exhibit identical response patterns with longer relative peak trial durations, testing the acquisition of peak responding with variable peak trial times that are longer than $3 \times \text{FI}$, with different FI values and other experimental parameters would add to the generality of our findings.

Accepted Manuscript

REFERENCES

- Abner, R.T., Edwards, T., Douglas, A., Brunner, D., 2001. Pharmacology of temporal cognition in two mouse strains. *Int. J. Comp. Psychol.*, 14: 189–210.
- Balsam, P.D., Drew, M.R., Yang, C., 2002. Timing at the start of associative learning. *Learn. Motiv.*, 33: 141–155.
- Bevins, R.A., Ayres, J.J.B., 1995. One-trial context fear conditioning as a function of the interstimulus interval. *Anim. Learn. Behav.*, 23: 400–410.
- Brunner, D., Kacelnik, A., Gibbon, J., 1996. Memory for inter-reinforcement interval variability and patch departure decisions in the starling, *Sturnus vulgaris*. *Anim. Behav.*, 51: 1025–1045.
- Catania, A.C., 1970. Reinforcement schedules and psychophysical judgments: a study of some temporal properties of behavior. In: Schoenfeld, WN (ed) *The theory of reinforcement schedules*. Appleton-Century-Crofts, New York, pp. 1–42.
- Cheng, R.K., Ali, Y.M., Meck, W.H., 2007a. Ketamine “unlocks” the reduced clock-speed effect of cocaine following extended training: evidence for dopamine–glutamine interactions in timing and time perception. *Neurobiol. Learn. Mem.*, 88: 149–159.
- Cheng, R.K., Hahak, O., Meck, W.H., 2007b. Habit formation and the loss of control of an internal clock: inverse relationship between the level of baseline training and the clock-speed enhancing effects of methamphetamine. *Psychopharmacology*, 193: 351–362.
- Cheng, K., Westwood, R., 1993. Analysis of single trials in pigeons timing performance. *J. Exp. Psychol. Anim. Behav. Process.*, 19: 56–67.
- Cheng, K., Westwood, R., Crystal, J., 1993. Memory variance in the peak procedure of timing in pigeons. *J. Exp. Psychol. Anim. Behav. Process.*, 19: 68–76.
- Church, R.M., Meck, W.H., Gibbon, J., 1994. Application of scalar timing theory to individual trials. *J. Exp. Psychol. Anim. Behav. Process.*, 20: 135–155.
- Davis, M., Schlesinger, L.S., Sorenson, C.A., 1989. Temporal specificity of fear conditioning: Effects of different conditioned stimulus/unconditioned stimulus intervals on the fear-potentiated startle effect. *J. Exp. Psychol. Anim. Behav. Process.*, 15: 295–310.
- Drew, M.R., Zupan, B., Cooke, A., Couvillon, P.A., Balsam, P.D., 2005. Temporal Control of Conditioned Responding in Goldfish. *J. Exp. Psychol. Anim. Behav. Process.*, 31: 31–39.

- Gallistel, C.R., Gibbon, J.G., 2000. Time, rate, and conditioning. *Psychol. Rev.*, 107: 289–344.
- Gallistel, C.R., King, A., McDonald, R., 2004. Sources of variability and systematic error in mouse timing behavior. *J. Exp. Psychol. Anim. Behav. Process.*, 30: 3–16.
- Gibbon, J., 1977. Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.*, 84: 279–325.
- Gibbon, J., Balsam, P.D., 1981. The spread of association in time. In C. Locurto, H. S. Terrace and J. Gibbon (Editors), *Autoshaping and conditioning theory*, Academic Press, New York, pp. 219–253.
- Gibbon, J., Church, R.M., 1990. Representation of time. *Cognition*. 37: 23–54.
- Guilhardi, P., Church, R.M., 2005. Dynamics of temporal discrimination. *Learn. Behav.*, 33: 399–416.
- Honig, W. K. 1981. Working memory and the temporal map. In N. E. Spear and R. R. Miller (Editors), *Information processing in animals: Memory mechanisms*. Erlbaum, Hillsdale, NJ, pp. 167-197.
- Jennings, D.J., Bonardi, C., Kirkpatrick, K., 2007. Overshadowing and stimulus duration. *J. Exp. Psychol. Anim. Behav. Process.*, 33: 464–475.
- Kaiser, D.H., 2008. The proportion of fixed interval trials to probe trials affects acquisition of the peak procedure fixed interval timing task. *Behav. Process.*, 77: 100–8.
- Kehoe, E.J., Ludvig, E.A., Dudeney, J.E., Neufeld, J., Sutton, R.S., 2008. Magnitude and timing of nictitating membrane movements during classical conditioning of the rabbit (*oryctolagus cuniculus*). *Behav. Neurosci.*, 122: 471–76.
- Kirkpatrick, K., Church, R.M., 2000. Independent effects of stimulus and cycle duration in conditioning: The role of timing processes. *Anim. Learn. Behav.*, 28: 373–388.
- Kirkpatrick-Steger, K., Miller, S.S., Betti, C.A., Wasserman, E.A., 1996. Cyclic responding by pigeons on the peak timing procedure. *J. Exp. Psychol. Anim. Behav. Process.*, 22: 447–460.
- Ludvig, E., Conover, K., Shizgal, P., 2007. The effects of reward magnitude on timing in rats. *J. Exp. Anal. Behav.*, 87: 201-218.
- Machado, A., 1997. Learning the temporal dynamics of behavior. *Psychol. Rev.*, 104: 243–262.

- Malapani, C., Rakitin, B., Meck, W.H., Deweer, B., Dubois, B., Gibbon, J., 1998. Coupled temporal memories in Parkinson's disease: a dopamine-related dysfunction. *J. Cogn. Neurosci.*, 10: 316–31.
- Matell, M.S., Bateson, M., Meck, W.H., 2006. Single-trials analyses demonstrate that increases in clock speed contribute to the methamphetamine-induced horizontal shifts in peak-interval timing functions. *Psychopharmacology*, 188: 201–212.
- Matell, M.S., Portugal, G.S., 2007. Impulsive responding on the peak-interval procedure. *Behav. Process.*, 74: 198–208.
- Meck, W.H., Komeily-Zadeh, R.N., Church, R.M., 1984. Two-step acquisition: Modification of an internal clock's criterion. *J. Exp. Psychol. Anim. Behav. Process.*, 10: 297–306.
- NRC, 1996. Guide for the care and use of laboratory animals. National Academy, Washington, DC
- Ohyama, T., Mauk, M.D., 2001. Latent acquisition of timed responses in cerebellar cortex. *J. Neurosci.*, 21, 682–690.
- Papachristos, E.B., Gallistel, C.R., 2006. Autosshaped Head Poking in the Mouse: A Quantitative Analysis of the Learning Curve. *J. Exp. Anal. Behav.*, 85: 293–308.
- Savastano, H. I., Miller, R. R., 1998. Time as content in Pavlovian conditioning. *Behav. Proc.*, 44: 147-162.
- Schwartz G, 1978. Estimating the dimension of a model. *Ann. Stat.*, 6: 461–464.
- Taylor, K.M., Horvitz, J.C., Balsam, P.D., 2007. Amphetamine affects the start of responding in the peak interval timing task. *Behav. Process.*, 74:168-75.
- Wynne, C., Staddon, J., 1988. Typical delay determines waiting time on periodic-food schedules: static and dynamic tests. *J. Exp. Anal. Behav.*, 50: 197-210.

FIGURE CAPTIONS

Figure 1. Average normalized response distribution on different days.

Figure 2. A) Log-log regression conducted on data averaged across subjects B) Summary of individual log-log regression results in terms of percentage of subjects. Sig: Significant.

Figure 3. Left Panel: Scatter plot of stop times as a function of trial number for 4 representative individual subjects. The stop times presented in this plot were detected with criterion 10. Right Panel: Average response rates for the corresponding subjects on the right panel. Solid lines represent the average response rates from two sessions that immediately preceded the session during which a change was detected in the stop times (using criterion 100) presented on the corresponding left panel. Dashed lines represent the average response rates from two sessions that immediately followed that session.

Figure 1

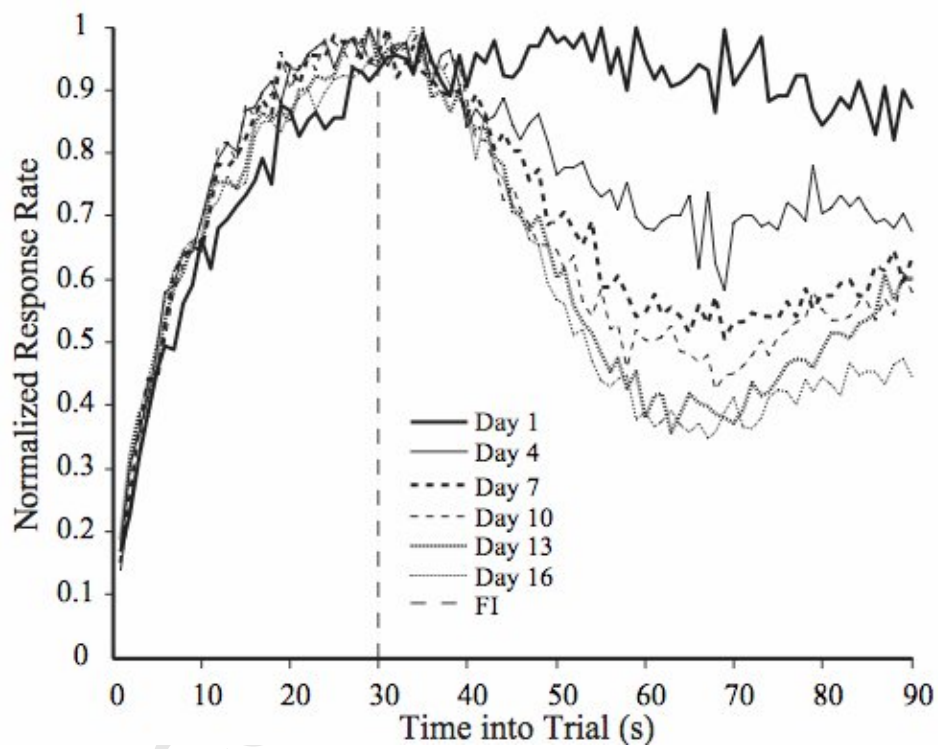
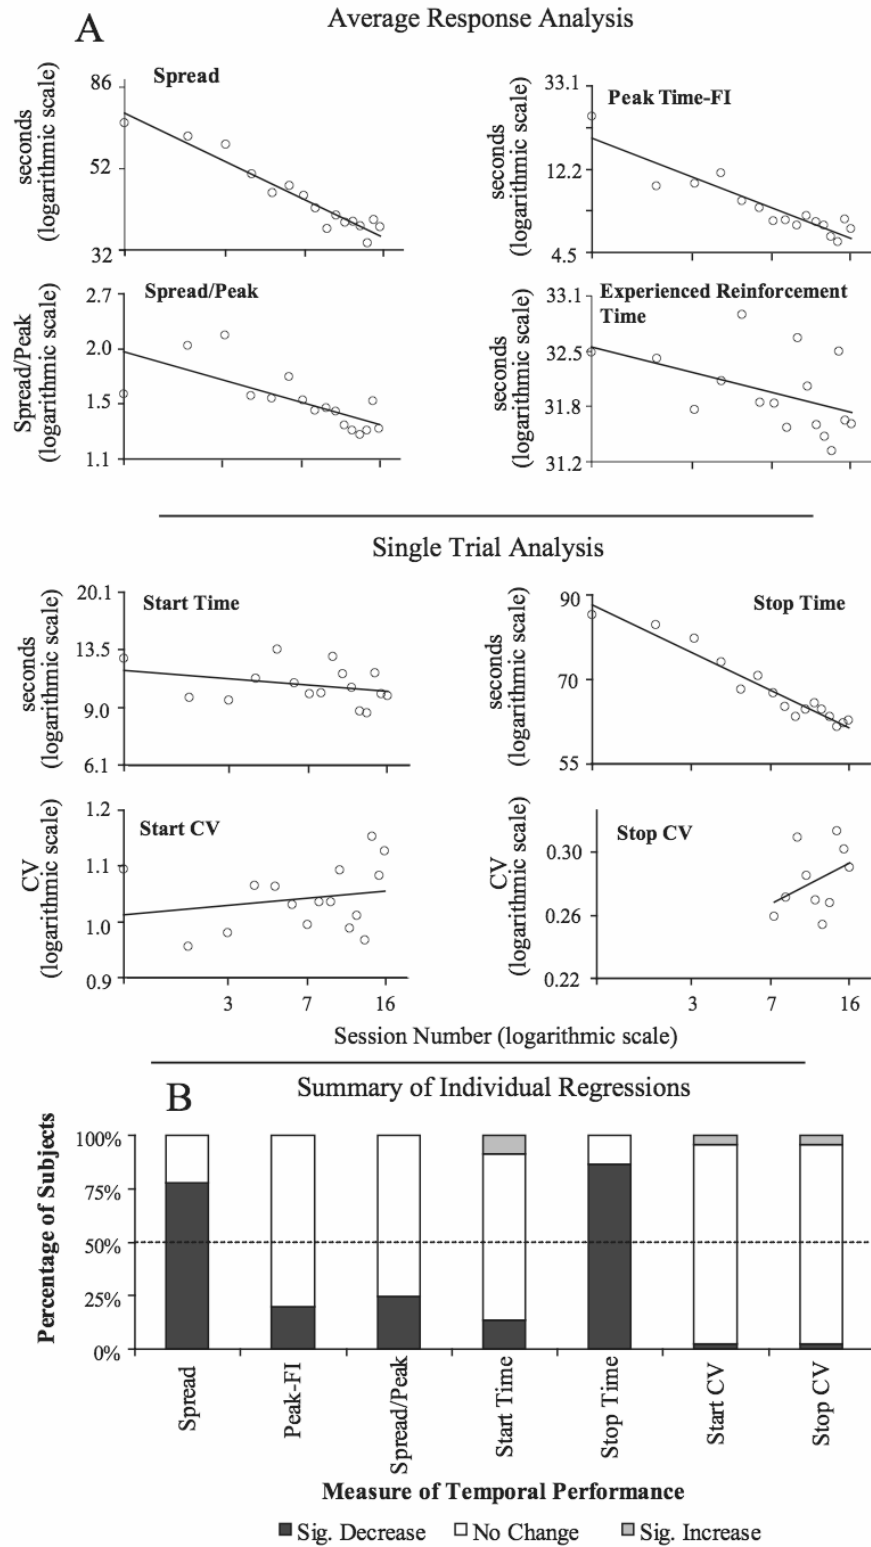


Figure 2



Acquisition of Peak Responding 30

Figure 3

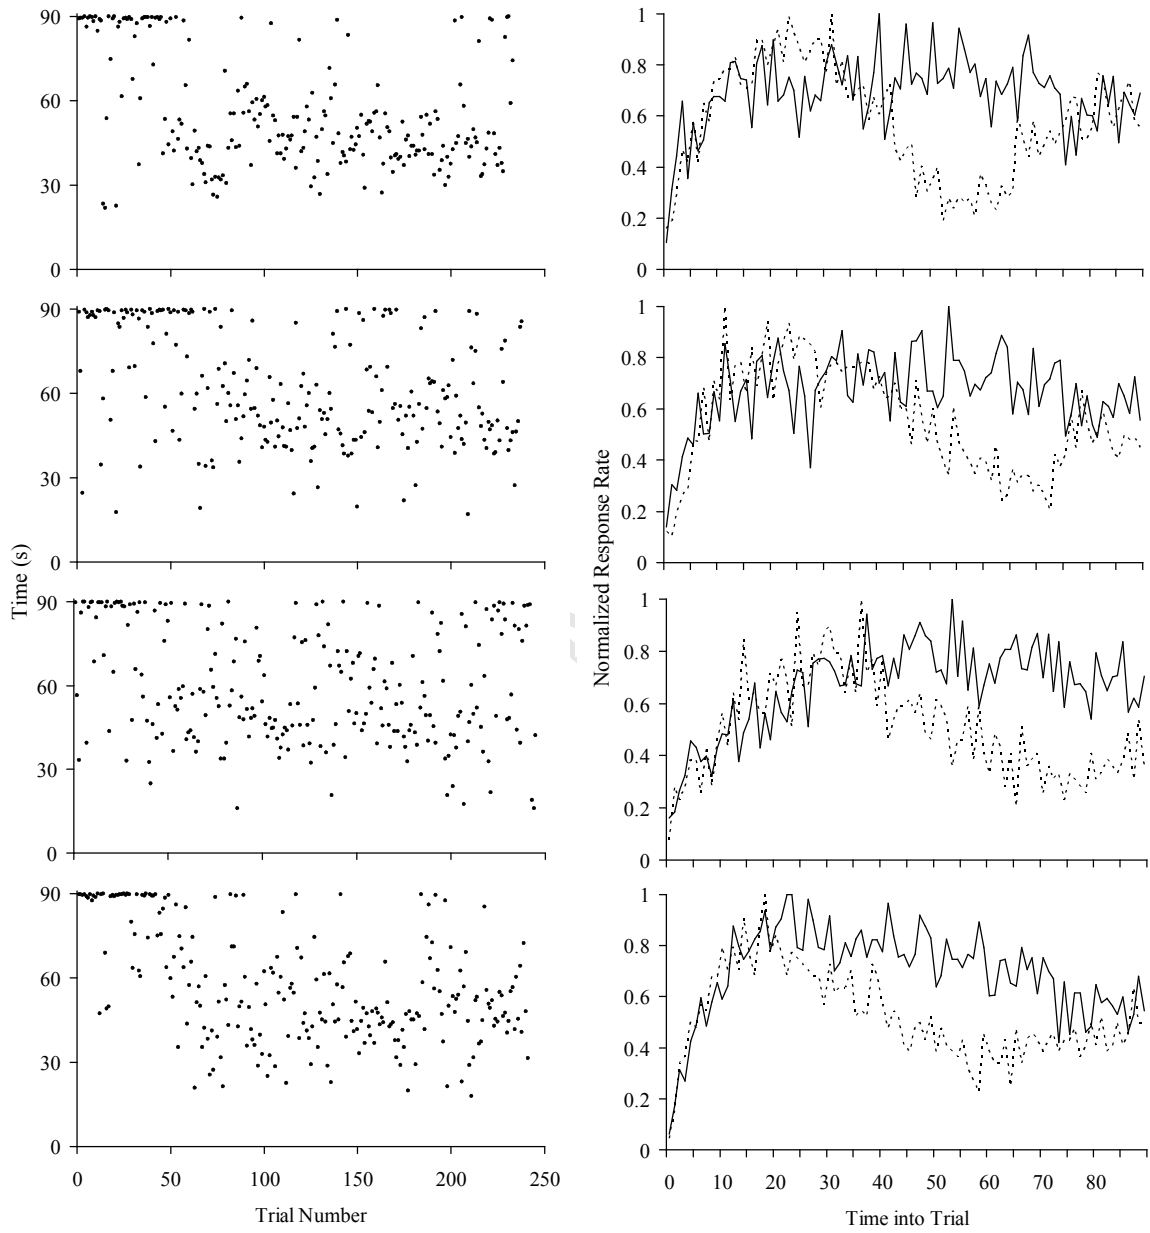


Table 1. Summary of change point analysis with different criteria

		Change Detection to Acquisition											
		10				100				1000			
p/(1-p)		Session	Trial	Frequency	Fraction	Session	Trial	Frequency	Fraction	Session	Trial	Frequency	Fraction
Change Detection for start and stop	10	<u>2</u> 1-4	<u>30</u> 14-51	<u>.67</u> .48-.88	<u>.92</u> .66-1	<u>4</u> 2-5	<u>49</u> 29-73	<u>.76</u> .59-.88	<u>.97</u> .77-1	<u>4</u> 3-7	<u>52</u> 37-106	<u>.81</u> .64-.88	<u>.97</u> .81-1
	100	<u>3</u> 2-4	<u>42</u> 28-55	<u>.57</u> .41-.86	<u>.98</u> .63-1.2	<u>4</u> 3-8	<u>54</u> 41-110	<u>.69</u> .49-.82	<u>.99</u> .83-1	<u>5</u> 4-9	<u>61</u> 45-129	<u>.74</u> .5-.82	<u>.99</u> .89-1
	1000	<u>4</u> 2-5	<u>44</u> 28-61	<u>.56</u> .38-.78	<u>.92</u> .65-1.2	<u>4</u> 3-7	<u>58</u> 41-106	<u>.60</u> .37-.76	<u>.96</u> .61-1.1	<u>5</u> 4-13	<u>66</u> 45-182	<u>.6</u> .42-.71	<u>.96</u> .84-1.1

Note: Each cell gives the median (underlined) and interquartile interval (below median) of a summary statistic for a given combination of decision criteria. The criterion used in the first phase (the search for a stop time in the sequence of interresponse intervals on each trial) varies by row; the criterion used in the second phase (the search for a change in the frequency of well-defined stops) varies by hypercolumn. Session is the session in which well-defined stops appeared; Trial is the trial after which they become frequent; Frequency is the frequency of well-defined stops in the interval immediately following the onset of frequent stops; Fraction is this frequency as a proportion of the frequency during the second half of the trials.