Journal of Experimental Psychology:
Animal Behavior Processes
2007, Vol. 33, No. 3, 339–348

0097-7403/07/$12.00 DOI: 10.1037/0097-7403.33.3.339

Occasion Setting of Timing Behavior

Charlotte Bonardi and Dómhnall Jennings
University of Nottingham

Rats were trained on a temporal switching discrimination, with 2 features, *A* and *B*, signaling 2 target conditioned stimuli, *x* and *y*; feature offset and target onset were separated by a 5-s feature–target interval, and all target stimulus presentations terminated in a food pellet. The target conditioned stimuli were either short or long (6 or 30 s): Specifically, when signaled by *A*, *x* was 6 s and *y* 30 s, but when *x* and *y* were signaled by *B*, *x* was 30 s and *y* 6 s. Trials with 6-s and 30-s targets were termed *short* and *long* trials, respectively. Probe tests indicated that the animals correctly anticipated when food was to be delivered on these 2 types of trials. In further testing, the interval between feature offset and target onset was lengthened, to investigate the precise mechanism underlying this behavior. This manipulation did not have a substantial effect on discrimination performance. These results are discussed with reference to theories of occasion setting, timing, and configural learning.

Keywords: occasion setting, timing, rats

An occasion setter is a stimulus that signals whether or not some other cue will be reinforced. For example, in a feature-positive occasion-setting discrimination (*F – t+, t–*) the animal will learn to respond to the target stimulus, *t*, more when it is signaled by the feature, *F*, than when it is presented alone. Such discriminations can be solved by means of Pavlovian conditioning: The feature, being the best available predictor of the unconditioned stimulus (US), acquires all of the associative strength, so that more responding is seen on trials when the feature is present than when it is absent. But in some cases discrimination performance is independent of the feature's associative properties, and in these instances the feature is termed an *occasion setter* (see Swartzentruber, 1995, for a review).

Several theories have been proposed to explain this behavior. Rescorla (e.g., 1985) argued that the occasion setter lowers the activation threshold of the US representation, making it easier to activate, whereas Holland argued that it acts on the CS–US association like an AND gate, facilitating the flow of activation between CS (conditioned stimulus) and US (Holland, 1983, 1985; see also Bonardi, 1989, 1998; Bouton, 1990; cf. Skinner, 1938). An increasing body of evidence suggests that occasion setters are, at least in part, specific in their action to the CS with which they were trained, supporting the AND gate suggestion (e.g., Bonardi, 1996; Rescorla, 1991a, 1991b). Evidence indicating that occasion setters are specific in their action to the particular CS–US association with which they were trained also supports this view (Bonardi, 2007; Bonardi & Ward-Robinson, 2001; cf. Jenkins, 1985). The implication of this account is that the feature allows access to the associative properties of the target CS. However, the associative properties of a CS may do more than simply tell the animal that the US will happen—they may also tell the animal when the US will occur. Given the considerable literature demonstrating animals' capability to time in simple conditioning procedures (e.g., see Gibbon & Allan, 1984; Killeen & Fetterman, 1988; Miller & Barnet, 1993; Staddon & Higa, 1993), it seems possible that occasion setters might also give the animal temporal information about the relationship between the CS and the US. The aim of the present experiments was to evaluate this suggestion.

Some prior work by Holland and colleagues investigated this question. Holland, Hamlin, and Parsons (1997) reported an experiment in which two groups of animals were trained on two feature-positive occasion-setting discriminations. The groups differed in the interval between feature offset and target onset (FTI)—for one group the FTI was 5 s, and for the other it was 25 s. In subsequent testing with a range of FTIs, the animals' performance was best with the FTI with which they had been trained. In a second experiment it was demonstrated that this temporal specificity of optimal performance was mediated by the feature, not by the target. A single group of animals was trained on two discriminations, one with an FTI of 15 s and one with an FTI of 35 s. Animals were then tested at both FTIs to examine whether performance was better when the test FTI matched that associated with the feature during initial training or that associated with the target. The authors found that regardless of the target's identity, the feature was always more effective at elevating responding when the test FTI matched the feature's training FTI than when it did not.

These results suggest that temporal information is being encoded with the feature stimulus. But there are a number of ways in which this information might be mediated, not all of them consistent with the idea that the feature is acting as an occasion setter that gives conditional information about when the CS is to be reinforced. For example, the animals might simply time the feature (i.e., learn exactly how long after feature onset food will occur)

Charlotte Bonardi and Dómhnall Jennings, School of Psychology, University of Nottingham, Nottingham, England.

Dómhnall Jennings is now at the Division of Psychology, School of Biology and Psychology, University of Newcastle upon Tyne, England.

We thank Jo Temperton, Stuart Morley, and Richard Wood for technical assistance and Kim Kirkpatrick for helpful advice.

Correspondence concerning this article should be addressed to Charlotte Bonardi, School of Psychology, University of Nottingham, University Park, Nottingham NG7 2RD, England. E-mail: cmb@psychology.nottingham.ac.uk

and use this information to time their responding at the appropriate time after feature offset. In this way, responding to the feature would coincide with that elicited by target CS presentation, and summation of these two response tendencies would produce maximal responding during the target CS. Changing the FTI would remove this coincidence of feature and target responding, and performance would be attenuated. However, Holland and colleagues (1997) argued that this was not the correct explanation of their results. In their first experiment, they showed that discrimination performance survived extinction of the feature in both groups (and was actually unchanged by extinction in the long FTI condition). Because feature extinction should eliminate the feature's Pavlovian properties, this in turn should abolish any summation effect. But although this argument is logical, it is not conclusive. First, feature extinction did have a substantial and significant effect on the more robust discrimination trained with a 5-s FTI, suggesting (as the authors acknowledged) a contribution of Pavlovian conditioning in this case. Moreover, any residual discrimination in this case could be explained if extinction were for some reason incomplete; to the extent that the feature retained any Pavlovian properties, then maintained discrimination would be expected. Second, although the researchers failed to find any effect of extinction on the discrimination trained with the 35-s FTI, this could be attributed to the fact that the levels of responding on this discrimination were already relatively low, so that differences between the extinction and control conditions were difficult to detect.

A second piece of evidence that was proposed to rule out this simple Pavlovian interpretation came from a study reported by Holland (1998). Animals were trained on two feature-positive discriminations, one with an FTI of 10 s and the other with an FTI of 30 s. In a subsequent test, the animals were tested at various FTIs with a simultaneous compound of the two features. The occasion-setting ability of this compound to control responding to the various targets was examined, as was its ability to elicit a Pavlovian conditioned response (assessed by recording food cup entry in the period preceding target presentation). Holland found numerical (but not statistically significant) evidence for "averaging" of the compound's occasion-setting ability: Animals responded most at a 15-s FTI, despite the fact that the compound's component features had been trained with 10-s and 30-s FTIs. In contrast, the peak of Pavlovian responding to the compound feature was indiscriminable from that shown by the feature trained with a 10-s FTI. Holland argued that this dissociation contradicted the idea that the feature's temporally specific control over target responding was due to its Pavlovian properties. But although indicative, again these data are inconclusive: Not only was one of the key effects, the averaging of occasion-setting power with the compound feature, not significant, but the two effects were assessed by different measures and never explicitly compared. In conclusion, it seems premature to reject the idea that in an occasion-setting discrimination the animal learns only to time the feature; moreover, whether or not one accepts Holland's interpretation of his data, it is still the case that different training procedures might produce different results. Accordingly, the present experiment aimed to examine whether an occasion setter conveys temporal information about the reinforcement of a CS, independent of its own Pavlovian properties.

A switching design was employed. Animals were trained with two auditory occasion setters, *A* and *B*, both of 10-s duration. These were followed, after a trace interval, by one of two visual target CSs, *x* and *y*. This yielded four stimulus combinations, *A–x*, *A–y*, *B–x*, and *B–y*, and all four trial types were reinforced: On *A–x* and *B–y* trials, the targets were of 6 s duration (short trials), and in the remaining trial types, *A–y* and *B–x*, the targets were of 30 s duration (long trials). In all cases, the targets terminated with US delivery. With this arrangement, neither the targets nor the features were unique predictors of whether the trial would be long or short, and so accurate performance could not be attributed to simple conditioning or timing to either type of stimulus. If animals can learn this discrimination, then it would suggest that the feature is able to provide conditional information about when the target CS is to be reinforced.

Experiment 1

Method

Subjects. The subjects were 16 male hooded Lister rats with a mean ad lib weight of 500 g (range = 425–510 g). They had previously participated in an experiment on flavor aversion learning but were naïve to the stimuli and procedures employed here. They were deprived to 80% of their ad lib weight before the start of the experiment and were maintained at this level for the rest of the experiment 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. The colony rooms were lit from 8 a.m. to 8 p.m.; the subjects were tested during the light portion of the cycle.

Apparatus. A set of four Skinner Boxes (Paul Fray, Cambridge, United Kingdom) were employed. Each was housed in a sound- and light-attenuating shell equipped with an exhaust fan that both ventilated the chamber and generated a background noise level of 65dB. The boxes contained a recessed food tray to which 45-mg food pellets (Noyes Formula A/I) could be delivered. A transparent plastic flap (4 cm high × 4.5 cm wide) that was hinged at the top covered the tray. Pushing this flap actuated a microswitch, closure of which was recorded as a response. The standard response levers were retracted throughout the course of the experiment. The floor was made from stainless steel rods. A loudspeaker mounted on the roof opposite the food tray was used to present a 79-dB, 10-Hz clicker and an 80-dB white noise (produced by Campden Instruments, Loughborough, United Kingdom), tone and noise generator, respectively. Illumination of the chambers was by means of an 8-W house light (operated at 12 V), situated in the center of the chamber ceiling. Two visual stimuli, dark and light, were produced, respectively, by turning off the house light and by illuminating three 8-W jewel lights around the food tray (one placed centrally 14.5 cm above the food tray, one 8 cm to the left of the food tray, and the other 8 cm to the right, the last two being 10 cm above the top of the food tray). The boxes were controlled by a BBC microcomputer programmed in a version of BASIC.

Procedure

Subjects first received a single 40-min session of magazine training in which reinforcers were delivered according to a variable time 60-s schedule.

Training. All animals were then trained on a switching discrimination in which each of the two auditory occasion setters, the noise and the click, signaled the reinforcement of each of the visual stimuli, light and dark, after either 6 or 30 s. Thus, for half of the animals the noise signaled that the light would be reinforced after 6 s and the dark after 30 s, and the click signaled the opposite, that the light would be reinforced after 30 s and the dark after 6 s. The remaining animals experienced the converse arrangement. Each session consisted of 40 trials, 10 of each type delivered in a semirandom order. Each trial consisted of a 10-s presentation of one of the auditory feature stimuli, followed by a 5-s trace period during which no stimuli were presented. This in turn was followed by presentation of one of the visual target cues for the appropriate duration. The reinforcement was a single food pellet, and the intertrial interval a fixed 50 s plus a variable interval with a mean of 60 s. There were 40 trials per session, giving a mean session duration of 101 min; there were 56 sessions in this stage.

Probe Test I. There followed 40 sessions of probe testing. These were identical to those of the previous stage, except that 4 probe trials were added. There were two types of probe session, one for each type of visual target stimulus. In each, the visual stimulus was signaled twice by one occasion setter and twice by the other. The probe trials were identical to the corresponding training trials except that (a) the visual stimulus duration was 90 s and (b) no reinforcement was delivered. The different types of probe session were delivered in a double alternation sequence (A, B, B, A, B, A, A, B).

Data treatment. One of the animals became ill over the course of the experiment and so was excluded from all analyses after the initial training phase. A significance level of $p < .05$ was adopted in all of the analyses that follow. Data across the course of the various types of trial were collected in 2-s bins. As all of the comparisons were within-subject, the standard errors are not shown in the figures.

Results

Training. The animals appeared able to learn the discrimination over the course of training. Figure 1 shows the number of responses, presented in 2-s bins, averaged over short trials and long trials and pooled over the last four sessions. It may be seen that on short trials responding increased rapidly at the start of the trial, whereas on long trials it rose more gradually. The groups appeared to differ during the initial 6-s portion of the trial, an impression confirmed by the results of an analysis of variance (ANOVA) performed on responding during the first three 2-s bins, with trial type (short, long) and bin (1–3) as factors: This revealed a significant main effect of bin, $F(2, 30) = 5.52$, and of trial type, $F(1, 15) = 6.68$; the interaction between these two factors was not significant, $F(2, 30) = 2.41$. The main effect of bin indicated that responding increased over bins in both trial types; the main effect of trial type confirmed that animals responded more on short than on long trials in this early part of the CS.

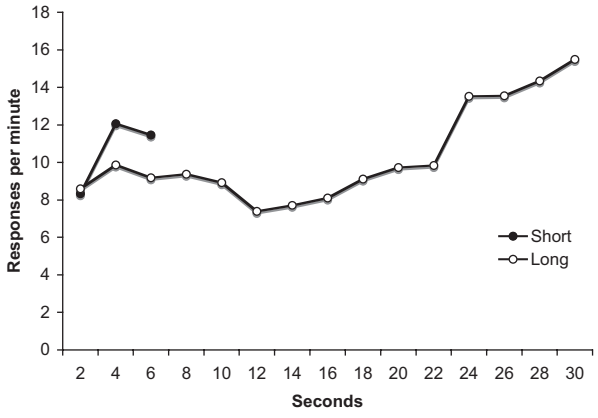

Figure 1. Response rates for short and long trials, from the last four sessions of initial training. Rates are presented in 2-s bins.

Probe Test I. In order to establish that responding during the two trial types also differed in the later part of the CS, it was necessary to examine the probe trial data. Responding in the two types of probe trial (short and long) was pooled over all the probe trials in the test phase; the resulting data are presented in Figure 2. It may be seen that in both short and long trial conditions there appeared to be a sharp peak of responding at around 6 s and a broader increase in responding at 30 s, suggesting that the animals were expecting food at both intervals on both types of trial. However, the relative size of these peaks did appear to differ between the two trial types, as predicted. On short trials the 6-s peak was substantially more pronounced than the 30-s peak, whereas on long trials the two peaks were much more similar in size, with the 30-s peak slightly more pronounced than the 6-s peak. Moreover, in the vicinity of the 6-s peak responding was greater on short trials than on long trials, whereas at the 30-s peak responding was greater on long than on short trials. This pattern of results seemed to confirm the prediction that animals would be able to learn this discrimination and respond appropriately on short and long trials. The fact that the discrimination was not perfect (in that animals showed both peaks on both trial types) probably reflects generalization between the two features and between the two targets; it could also stem from the animal having forgotten which of the two features had just occurred when the target was presented.

The statistical analysis aimed to confirm the significance of these observations by seeking statistical evidence that responding at the two peaks interacted with trial type—in other words, that responding at the 6-s peak was greater on short than on long trials, but that the reverse was true at the 30-s peak. In order to separate responding at the two peaks, we divided the data into two 24-s blocks (Block 1: 1–24 s, Bins 1–12; Block 2: 25–48 s, Bins 13–24) based on the predicted peaks of responding. This meant that the two blocks were exactly comparable in the sense that on short trials reinforcement would be expected 6 s after the start of Block 1, whereas on long trials it would be expected 6 s after the start of Block 2. We predicted that we would obtain a significant Block × Trial Type interaction, with post hoc tests showing greater responding on short trials in Block 1 and on long trials in Block 2. We also predicted that these differences in responding would be more pronounced at the actual times when reinforcement was normally delivered—6 and 30 s after target onset. Accordingly, we included 2-s bin as a factor to examine the peak of responding with more accuracy. If the difference in responding between the two trial types showed a different pattern across bins in the two blocks, then the critical Block × Trial Type interaction should also inter-

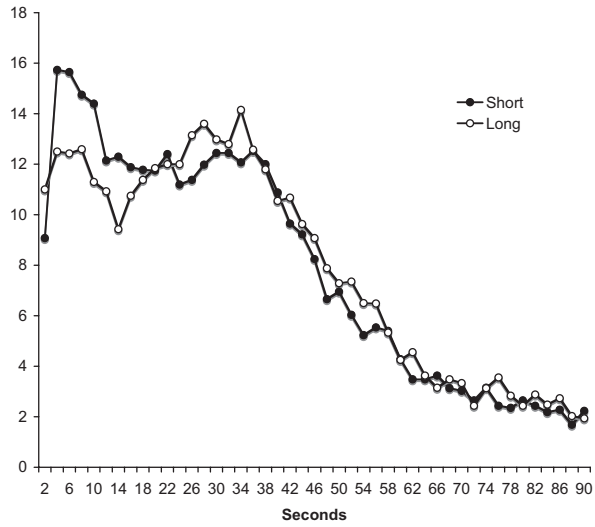

Figure 2. Response rates for short and long probe trials, pooled over probe testing in Experiment 1. Rates are presented in 2-s bins.

act with bins. The scalar property of timing led us to expect such an interaction, simply because peaks early in the trial (Block 1) would tend to be sharp and confined to a few bins, whereas peaks later in the trial (Block 2) would be more likely to be broad and less localized to particular bins.

An ANOVA with block, trial type (short, long), and bin (1–12) as factors confirmed most of these predictions. It revealed a main effect of bin, $F(11, 154) = 11.58$, $p < .01$, and interactions between block and trial type, $F(1, 14) = 8.66$, $p < .02$; block and bin, $F(11, 154) = 7.30$, $p < .01$; and trial type and bin, $F(11, 154) = 3.12$, $p < .01$. Critically, the three-way interaction between the three factors was significant, $F(11, 154) = 2.23$, $p < .02$. The main effects of block and trial type were not significant, $F(1, 14) = 1.68$ and 1.72, respectively. The three-way interaction allowed us to (a) investigate the prediction that responding on the two trial types differed in the predicted direction in each block and (b) isolate in which bins this difference lay. Accordingly, separate ANOVAs were conducted on the data from each block, with trial type and bin as factors. For Block 1 this revealed significant main effects of trial type and bin, $F(1, 14) = 8.94$, $p < .01$, and $F(11, 154) = 5.03$, $p < .01$, respectively, and a significant interaction between these two factors, $F(11, 154) = 4.24$, $p < .01$. The effect of trial type supported our prediction that responding would be greater on short trials, and the interaction between trial type and bin implied that this difference was not found in all bins of the block. Subsequent simple main effects analysis confirmed that responding was greater on short than on long trials on Bins 2 (Seconds 3–4) and 3 (Seconds 5–6), $F(1, 14) = 5.03$, $p < .05$, for both, and also on Bin 5 (Seconds 9–10), $F(1, 14) = 4.64$, $p < .05$. No other differences in responding between the two trial types were significant, largest $F(1, 14) = 3.99$.

A parallel analysis was conducted on the data from Block 2. This revealed a significant main effect of trial type, $F(1, 14) = 4.84$, $p < .05$, and of bin, $F(11, 154) = 16.42$, $p < .01$; the interaction between these two factors was not significant, $F < 1$. The effect of trial type confirmed that in Block 2 animals were doing the opposite of what they had been doing in Block 1—responding more on long trials than on short trials. The lack of an interaction with bin also confirmed our suspicion that, because of the scalar property, this difference would not be confined to specific bins. The analysis indicated only that responding on the two trial types differed over the whole of Block 2.

These analyses confirmed that animals could learn the discrimination and respond differentially to the two target stimuli depending on how they had been signaled: Early in the stimulus animals responded more on short trials, but later in the stimulus they responded more on long trials. Also, the accuracy with which they could track the time of reinforcement was greater at the shorter intervals, as scalar timing theory would predict: In the first block (Bins 1–12), animals responded most between 3 and 6 s after CS onset when reinforcement was expected, whereas in the second block there was no such temporal specificity but rather a general tendency to respond more on long than on short trials.

These conclusions were supported by a consideration of the times of peak responding. The data were smoothed to reduce the impact of random fluctuations by taking a 3-bin running mean, and then the bin at which the maximum number of responses was made was determined for each animal on each type of trial. Although animals clearly showed two peaks, if the 6-s peak were more prominent on short trials, and the 30-s peak on long trials, this would strongly support our conclusion that animals were responding appropriately on the two types of trial. The mean peak was 12.34 s for short trials and 20.98 s for long trials, and these values differed significantly, $F(1, 14) = 5.72$, confirming our prediction. This result was more notable in view of the fact that response rates generally tended to be higher early in the trial. This makes it particularly striking that, on long trials, responding was faster later in the trial.

Although this difference in peak times was consistent with the contention that the animals discriminated between the two types of trials, the actual values of the peak times did not encourage the view that they were responding at the appropriate times (6 and 30 s after target onset), as the average peak times were around 12 and 21 s. A scatter plot of the data, shown in Figure 3, gives more insight into why these figures were obtained. It may be seen that the individual peak response times that contributed to the two

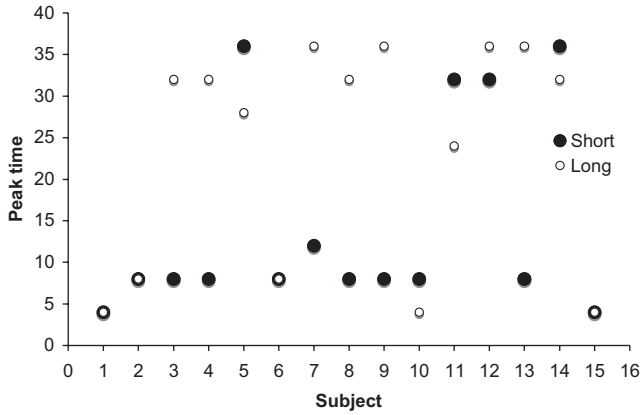

Figure 3. Scatter plot of time of peak responding (in seconds) for short and long trials, pooled over probe testing in Experiment 1.

means appear to be bimodally distributed, occurring either 2 to 10 s after CS onset, or 22 to 34 s after CS onset. On short trials, 11 of the 15 animals had peak times between 2 and 10 s, and only 4 had peaks between 22 and 34 s. Conversely, on long trials, 10 of the 15 animals had peak times between 22 and 34 s, and the remaining 5 subjects' peaks fell between 2 and 10 s.

Discussion

The results from the first experiment provide good evidence that animals can encode conditional temporal information. The animals showed selective peaks of responding depending on the combination of stimuli they received. This behavior cannot be explained in terms of simple conditioning or timing to either the features or the targets, as neither type of stimulus uniquely predicted a specific delay of reinforcement. Nevertheless, it is unclear from these results whether the temporal information is encoded by the feature or the target. The experiment looked for evidence that presentation of a particular feature could convey how long after a specific target's onset that target would be reinforced, and the results just presented are consistent with this interpretation. But they are equally consistent with the proposal that the presence of a particular target stimulus could tell the animal how long after a specific feature's offset that that feature would be reinforced. Either the feature tells the animal how to time the target, or the target tells the animal how to time the feature.

One way of discriminating between these two possibilities is to extend the FTI, as then the accounts' predictions differ. For example, if the conditional temporal information is encoded with the feature, then when the target is presented the animal must remember which feature preceded it; the identity of the feature will tell him when the target he is currently experiencing will be reinforced, and he will time his responding from the onset of that target. If the feature indicates that, with this target, the trial type is short, then he will respond 6 s after target onset. If the feature indicates that, with this target, the trial type is long, then he will respond 30 s after target onset. In this case, extending the FTI should have no effect on responding, as long as the identity of the feature can still be recalled, and the animal will show peak responding at the same times as he did in training. In contrast, if the identity of the target stimulus tells the animal how long after its offset the feature will be reinforced, then the prediction is different. For example, suppose that on a long trial the FTI is increased from 5 to 29 s. Now the animal must not only remember the identity of the feature to know what sort of trial it is when the target is presented, he must also use the time since the offset of that feature to time his response, as he knows that food occurs 35 s after feature offset. If the FTI is increased from its training value of 5 s to, say, 29 s, then 35 s after feature offset coincides with 6 s after target onset—in other words, he will now respond as though the long trial were a short trial.

Another explanation of the results is in terms of configural learning. According to configural theory, stimuli should not be viewed as independent entities. Thus, a stimulus compound is said to differ from the sum of its parts—for example, by an additional, unique cue that is only present in the compound (e.g., Rescorla, 1972). This class of theory has proved very powerful in explaining a wide variety of behavioral phenomena; however, configural theories have not usually been applied as explanations of timing effects. One exception to this generalization is a proposal made by Vogel, Brandon, and Wagner (2000). These authors assumed the basic theoretical framework of the sometimes-opponent process theory (Wagner, 1981), according to which presentation of a stimulus recruits a selection of corresponding stimulus elements into an active memory state, from which they decay into a secondary state of activation before becoming inactive again. This recruitment process persists as long as the stimulus is being presented. Vogel et al. suggested that timing could be incorporated into this framework by assuming that the order in which elements are recruited is maintained constant. This means that elements occurring at the start of the CS will acquire less associative strength than those recruited at the end, allowing timing functions to develop.

Applying these ideas to the present data, one could argue that differential responding on short and long trials arises because configures of particular feature/target combinations are reinforced at specific times. Specifically, on long trials the memory elements active when reinforcement occurs will be a configure of the feature elements 35 s after feature offset (5-s trace interval plus 30 s of target) and the target elements 30 s after target onset. This account thus predicts that the pattern of responding will alter if the FTI is increased. For example, if the FTI is now extended to 29 s, there will be a conflict as to when to respond—the feature elements will tend to elicit responding 6 s after target onset and the target elements 30 s after target onset; perhaps more importantly, the critical configure of these two stimulus patterns will no longer occur. It is difficult to see how accurate responding could be maintained under these circumstances.

In order to address these predictions, a second set of probe trials was administered. These were identical to those of Experiment 1 except for the fact that in half these trials the FTI was increased from 5 to 29 s. If the original pattern of responding is maintained under these conditions, then this would support our prediction that the feature gives conditional information about when the target is to be reinforced.

Experiment 2

Method

Subjects and apparatus. The subjects were those from the previous experiment, and the same apparatus was employed.

Retraining. Following Probe Test I from the previous experiment, all animals received 28 sessions of retraining.

Probe Test II. All animals received a series of test sessions in which they received probe trials with either the training FTI or a longer FTI of 29 s. In all other respects these probe trials were like those of the previous probe test, except that there were now four types of session rather than two: two with each type of visual stimulus and, of these two pairs, one was with the training FTI and the other was with the longer FTI. Sessions were arranged in nine 4-session blocks, with each block containing one of each session type. The different session types were programmed randomly within each block.

Data treatment. This was identical to that employed in the previous experiment.

Results

The mean rates of responding on long and short trials are shown for probe trials with normal and extended FTIs in the top and bottom panels, respectively, of Figure 4. In both FTI conditions, the discrimination was apparently maintained as before: On short trials the animals seemed to respond more around 6 s than around 30 s, whereas in the long condition the opposite was the case. The main difference between the two FTI conditions was on the short trials: Animals with the short FTI appeared to respond considerably more than those with the extended FTI, especially early in the trial.

First we compared responding on short trials with the training FTI with that on long trials with the extended FTI; the critical issue was whether long trials with the extended FTI were treated as long or short. If the conditional temporal information lies with the target, then behavior characteristic of short trials should be seen on long trials with the extended FTI—responding should peak at 6 s when, on the basis of the feature's trace, food is predicted to occur. But if, as we predicted, conditional temporal information lies with the feature, then the characteristic long trial response pattern should be retained on these trials. Once again it appeared our predictions were borne out. On short trials the animals responded more early in the trial than later on, whereas on long trials the opposite pattern was apparent. This description of the data was evaluated statistically, using the same ANOVA as in the previous experiment, with trial type, block, and bin as factors. This revealed significant main effects of trial type, $F(1, 14) = 19.68$, and bin, $F(11, 154) = 11.83$, $ps < .01$; the effect of block was not significant, $F < 1$. The critical interaction between trial type and block was significant, $F(1, 14) = 8.46$, $p < .02$, although in this instance it did not interact with bin, $F(11, 154) = 1.46$. The interactions between trial type and bin, and between block and bin, were also significant, $F(11, 154) = 2.16$, and $F(11, 154) = 3.94$, respectively. Simple main effects analysis performed on the critical Trial Type × Block interaction revealed an effect of trial type on responding in Block 1, $F(1, 14) = 27.03$, $p < .01$, although not on Block 2, $F(1, 14) = 1.16$. The fact that responding did not differ on the second block might be taken to indicate that animals were not responding appropriately on long trials. However, the analysis also showed that, although on short trials animals responded significantly more in the first block than in the second, $F(1, 14) = 12.42$, $p < .01$, on long trials they responded significantly more in the second block than in the first, $F(1, 14) = 7.81$, $p < .02$, a pattern that was consistent with appropriate long trial responding. Overall these results lent little support to the prediction that extending the FTI would stop animals responding in a manner characteristic of long trials.

Further support for this conclusion came from the mean peak response times. For the long trials with the extended FTI condition, the peak response time was 25.93 s, considerably higher than the mean of 9.93 s seen in the short trials with the unchanged FTI condition. These means differed significantly, $F(1, 14) = 8.72$, $p < .02$. Indeed, the peak time for long trials with the long FTI seemed more similar to that seen on long trials with the unchanged FTI condition, 28.00 s. This impression was confirmed with an ANOVA performed on these three sets of peak time data, which revealed that these scores differed significantly, $F(2, 28) = 8.31$, $p < .01$. Tukey's honestly significant difference test revealed that 9.93 s differed from both 28 and 25.93 s, $ps < .01$; the latter two conditions did not differ. These results appeared to indicate that animals still recognized long trials as such even with an extended FTI.

The second question was whether the difference in responding on short and long trials seen in the previous experiment was affected by extending the FTI. Configural theory would predict that the long/short discrimination would be profoundly impaired by extending the FTI. In order to evaluate this prediction, we performed an overall factorial ANOVA on the data from short and long trials, with both normal and extended FTIs, with block (Bins 1–12, 13–24) and bin (1–12) as factors. We predicted that extending the FTI would not impair the discrimination, and that the difference between responding on short and long trials would persist. The ANOVA revealed significant main effects of FTI, $F(1, 14) = 15.82$; trial type, $F(1, 14) = 16.47$; and bin, $F(11, 154) = 13.10$, $ps < .01$. The critical Block × Trial Type interaction was again significant, $F(1, 14) = 14.63$, but this interacted with FTI, $F(1, 14) = 8.29$. There were also significant two-way interactions between FTI and trial type, $F(1, 14) = 11.66$, and block and bin, $F(11, 154) = 5.12$, $ps<.01$; and between trial type and bin, $F(11, 154) = 2.29$, $p < .02$. There were significant three-way interactions between FTI, trial type, and bin, $F(11, 154) = 2.30$, $ps < .02$. Nothing else was significant, largest $F(1, 14) = 1.71$.

The critical three-way interaction between FTI, block, and trial type implied that extending the FTI might have influenced the key response pattern of high early responding on short trials and high late responding on long trials. Accordingly, the interaction was examined further by conducting two-way ANOVAs on block and trial type separately for the two FTI conditions. For the short FTI condition this revealed no effect of block, $F < 1$, but a significant main effect of trial type, $F(1, 14) = 22.86$, and also a significant interaction between block and trial type, $F(1, 14) = 15.84$, $p < .01$. Simple main effects performed on the critical Trial Type × Block interaction revealed that the effect of trial type was significant in both Block 1, $F(1, 14) = 117.40$, and Block 2, $F(1, 14) = 16.60$, $ps < .01$. Thus, in the training FTI condition the original discrimination had been maintained from the previous experiment.

A similar pattern was observed in the long FTI condition: An identical analysis revealed no effect of block, $F(1, 14) = 1.09$, or of trial type, $F < 1$, but a significant interaction between these two factors, $F(1, 14) = 6.47$, $p < .05$. Simple main effects analysis again confirmed that the effect of trial type was significant on both Blocks 1 and 2, $F(1, 14) = 7.13$, $p < .02$, and $F(1, 14) = 15.04$, $p < .01$, respectively. These results confirmed that extending the FTI did not eliminate the critical discrimination between short and long trials. Inspection of the data suggested that the Block × Trial Type × FTI interaction may have arisen from the disproportionately high responding on short trials in the short FTI condition. This could have yielded the main effect of trial type that was found in the short FTI condition but not in the long FTI condition.

Scatter plots of the smoothed peak times for short and long trials for the unchanged and extended FTIs are shown in the top and bottom panels, respectively, of Figure 5. The means for short and long trials were 9.93 and 28 s, respectively, for the training FTI and 17.97 and 25.93 s, respectively, for the extended FTI condition. An ANOVA performed on these data revealed no main effect of FTI, $F < 1$, but a main effect of trial type, $F(1, 14) = 11.13$, $p < .01$, and a marginally significant interaction between these two

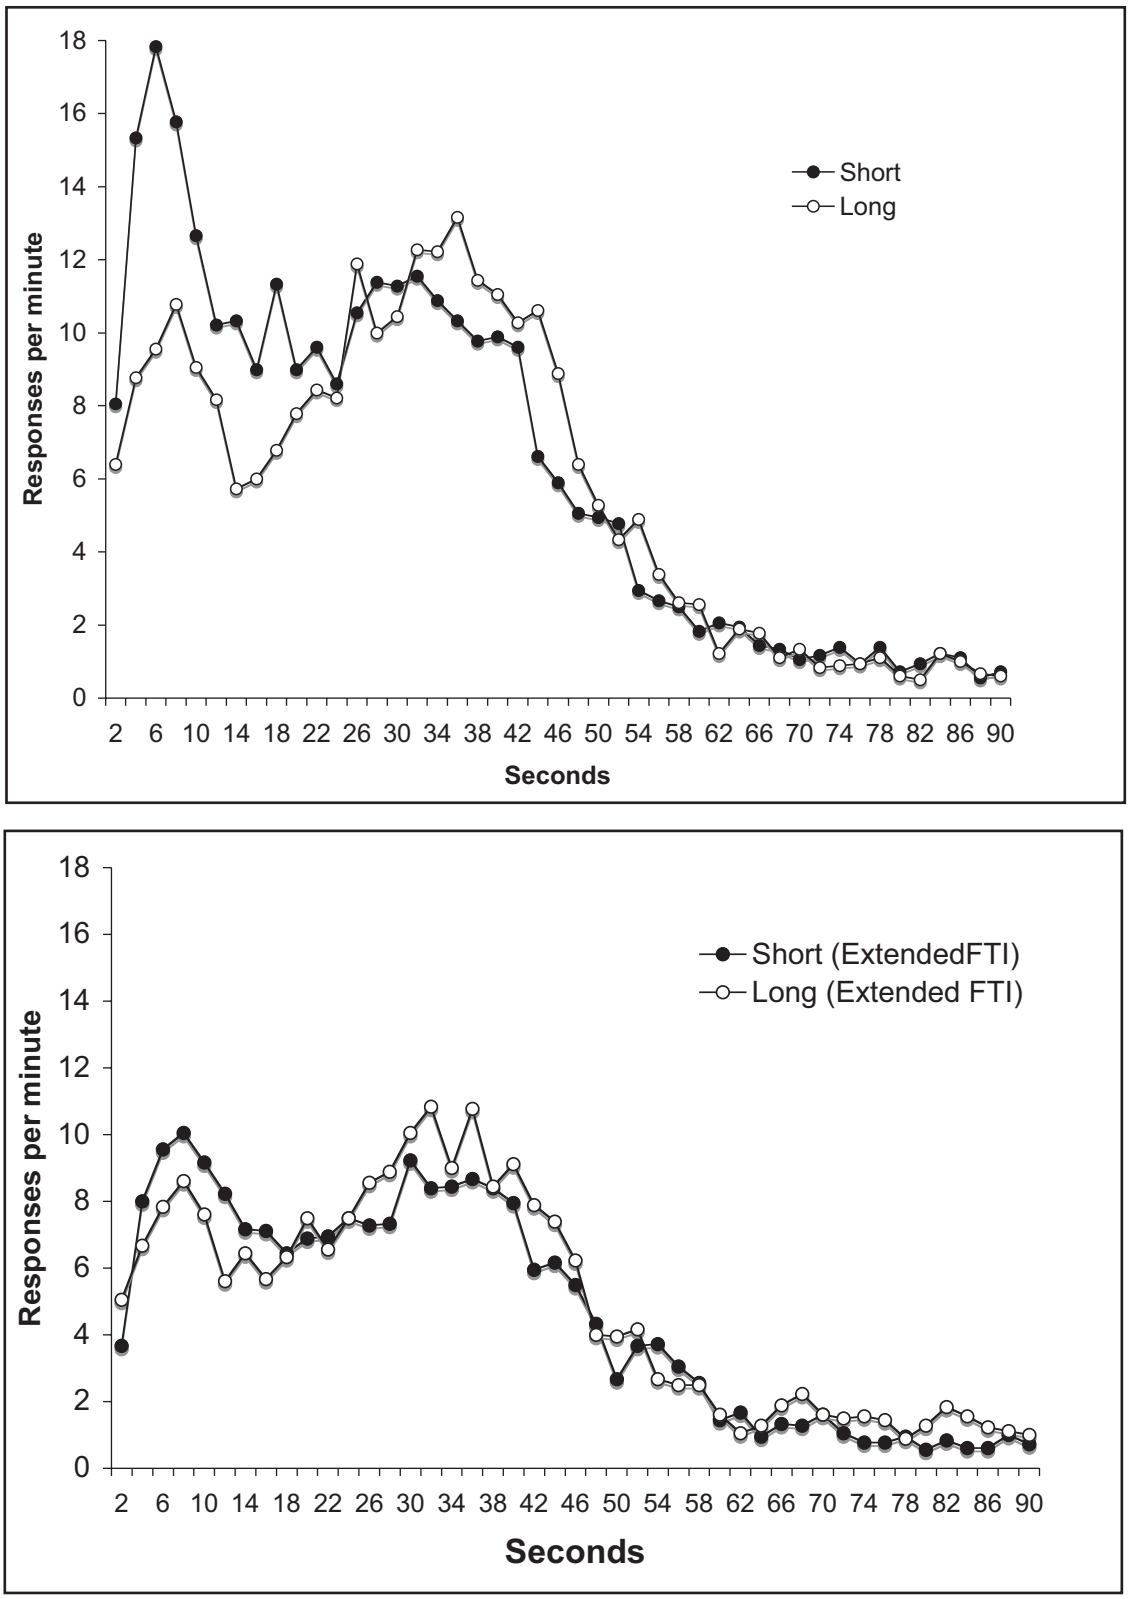

Figure 4. Top: Response rates for short and long probe trials with the original feature–target interval (FTI), pooled over probe testing in Experiment 2. Bottom: Response rates for short and long probe trials with the extended FTI, pooled over probe testing in Experiment 2. All rates are presented in 2-s bins.

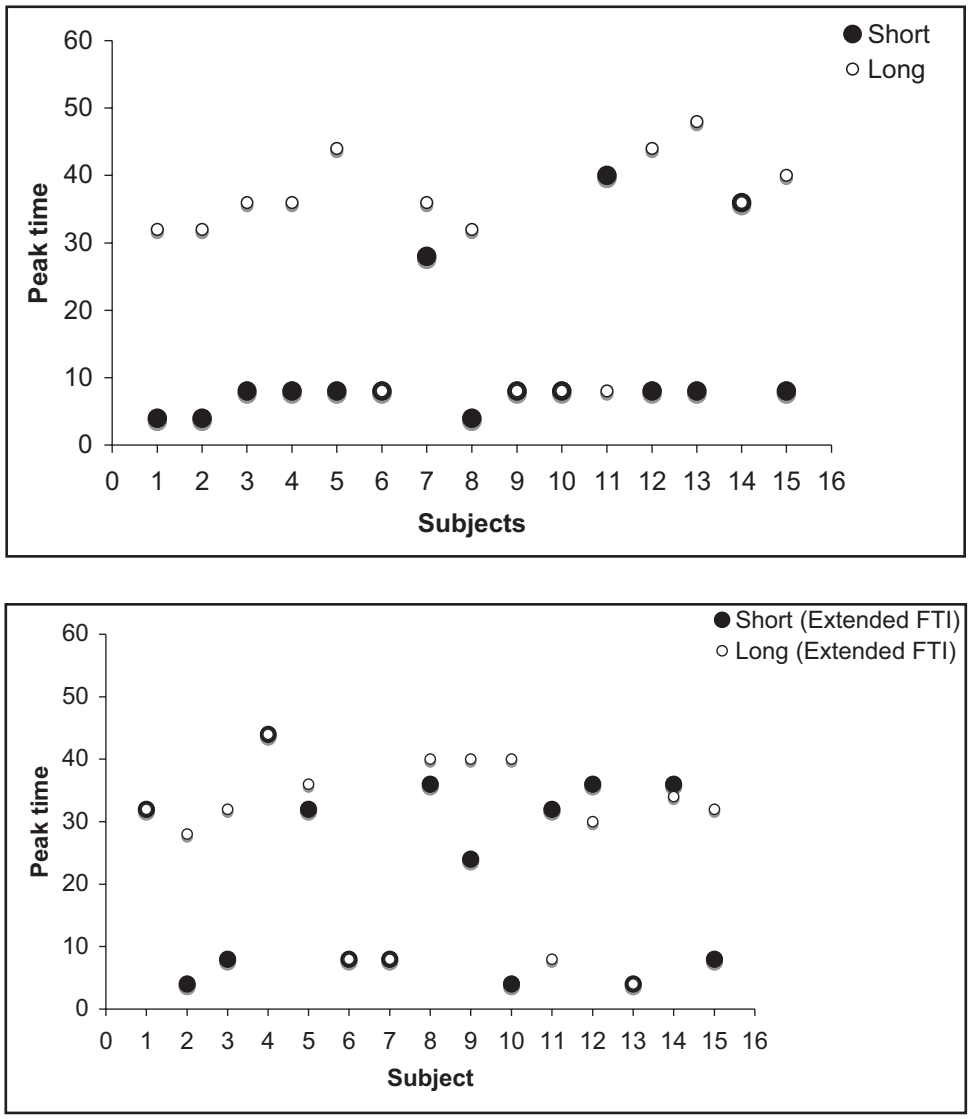

Figure 5. Top: Scatter plot of time (in seconds) of peak responding for short and long trials with the original feature–target interval (FTI), from data pooled over probe testing in Experiment 2. Bottom: Scatter plot of time (in seconds) of peak responding for short and long trials with the extended FTI, from data pooled over probe testing in Experiment 2.

factors, $F(1, 14) = 4.58$, $p = .051$. Simple main effects revealed that the difference in peak times between short and long trials was significant in the training FTI condition, $F(1, 14) = 10.72$, $p < .01$, but not in the extended FTI condition, $F(1, 14) = 2.09$. Thus, although there was some differentiation between the peak times in the extended FTI condition, it was not statistically significant.

Discussion

These results do not support the view that the target was signaling when the feature would be reinforced, rather than vice versa; nor did they confirm the predictions of configural theory. There was no convincing evidence that extending the FTI on long trials made the animals respond as though they were in a short trial, as the first of these accounts would predict. The peak response times especially reflected their apparent belief that they would still be reinforced after 30 s rather than after 6 s. Moreover, extending the FTI had little effect on the short/long trial discrimination, as both of these accounts would predict: The characteristic pattern of responding that had been observed during training was also observed with the extended FTI, with animals responding more on short trials early in the trial, and more on long trials later in the trial. The only obvious difference in discrimination performance between the two FTI conditions was that the times of peak responding no longer differed when the FTI was extended. More generally, the main difference between the two FTI conditions appeared to be on short trials: Animals with the short FTI appeared to respond considerably more than those with the extended FTI, especially early

in the trial. This may be due to summation of the feature's Pavlovian associative strength with that of the target, an effect that would be attenuated by extending the FTI.

General Discussion

The results of these two experiments allow us to draw a number of conclusions about the way in which temporal information is encoded in this switching task. Experiment 1 demonstrated that animals were capable of learning the task, and the design of the experiment meant that the animals' discrimination performance could not be explained in terms of simple conditioning or timing to either the feature or the target stimuli. We then considered various alternative explanations. The first suggestion was that, consistent with the results reported by Holland (1998) and Holland and colleagues (1997), the conditional information was encoded by the feature—specifically, the identity of the feature gave the animal information about how long after its onset the target stimulus would be reinforced. This was contrasted with the converse view, that the identity of the target gave the animal information about how long after the feature's offset reinforcement would be delivered. The results of Experiment 2 support the first of these possibilities: Extending the FTI had little effect on discrimination performance; moreover, there was no evidence that extending the FTI turned long trials into short trials, as this account would predict. These results also rule out a configural interpretation, according to which animals are conditioned to a configural cue comprising a specific trace of the target and a specific trace of the feature. Extending the FTI would, according to this account, eliminate the reinforced configure, resulting in a profound disruption of discrimination performance. No such effect was observed.

Timing theories have trouble accommodating these results; the underlying reason for this lies in the relationship between theories of conditioning and of timing. Theories of timing are not designed to explain conditioning (e.g., scalar timing theory; Gibbon, Church, & Meck, 1984), whereas theories of conditioning frequently have no way of conceptualizing temporal factors (e.g., Rescorla & Wagner, 1972). Thus, a theoretical approach such as that proposed by scalar timing theory would have no means of representing the conditionality of timing that was seen in the present experiment. Moreover, although the need for a unified theory of timing and conditioning has been recognized (e.g., Kirkpatrick & Church, 1998; Savastano & Miller, 1998), the success of the hybrid theories developed to address this need is as yet limited. For example, one type of approach is to propose that conditioning arises from a timing-like mechanism (Gibbon, 1977; Gibbon & Balsam, 1981; see also Gallistel & Gibbon, 2000), which estimates the probability that a CS will produce a US by comparing the rate of reinforcement during the CS and in its absence. To the extent that the CS produces an increase in the probability of reinforcement, then conditioned responding will occur. Theories of this type can explain conditioning in simple paradigms but fail to adequately address several fundamental properties of associative learning, such as blocking. At present they have no way of conceptualizing conditional learning phenomena such as those demonstrated in the present experiments and so are powerless to explain these results.

There are alternative approaches to explaining these data. One of these is exemplified by real-time theories such as the temporal difference model (Sutton & Barto, 1981, 1990; see also Vogel et al., 2000). This attempts to incorporate learning about time into a conditioning framework based on Rescorla–Wagner type principles. For example, the temporal difference model conceptualizes the CS as comprising a series of independently conditionable units. Timing arises because those units proximal to the US acquire more associative strength than those that are more distant, resulting in a gradient of responding that increases to the point of US delivery. Models of this type, although not at present incorporating conditional learning phenomena, would presumably be able to do so with as much ease as any other standard conditioning model. A slightly different computational approach, proposed by Buhusi and Schmajuk (1999), is one that proposes that a CS evokes multiple memory traces that peak at different moments in time, have different peak amplitudes, and are active over different time intervals. These traces can condition independently, activating the US either directly or, in combination with a second CS, via a hidden unit that is in turn associated with the US. The hidden units may be taken to correspond to configural cues, and they allow the CS to act as an occasion setter. The authors argued that these assumptions allow the CS to have different roles, as a CS or as an occasion setter, at different moments in the trial. It is a possible that a model with this flexibility would be able to explain our pattern of results. In conclusion, one of these two classes of approach will probably represent the way forward in explaining results of this type.

References

Bonardi, C. (1989). Inhibitory discriminative control is specific to both the response and the reinforcer. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology, 41*(B), 225–242.

Bonardi, C. (1996). Transfer of occasion setting: The role of generalization decrement. *Animal Learning and Behavior, 24,* 277–289.

Bonardi, C. (1998). Conditional learning: An associative analysis. In P. C. Holland & N. A. Schmajuk (Eds.), *Associative learning and cognition in animals: Occasion setting* (pp. 37–67). Washington, DC: American Psychological Association.

Bonardi, C. (2007). Occasion setting is specific to the CS-US association. *Learning and Motivation, 38,* 208–228.

Bonardi, C., & Ward-Robinson, J. (2001). Occasion setters: Specificity to the US and the CS-US association. *Learning and Motivation, 32,* 349–366.

Bouton, M. E. (1990). Context and retrieval in extinction and in other examples of interference in simple associative learning. In L. W. Dachowski & C. F. Flaherty (Eds.), *Current topics in animal learning: Brain, emotion and cognition* (pp. 25–53). Hillsdale, NJ: Erlbaum.

Buhusi, C. V., & Schmajuk, N. A. (1999). Timing in simple conditioning and occasion setting: A neural network approach. *Behavioural Processes, 45,* 33–57.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. *Psychological Review, 107,* 289–344.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84,* 279–325.

Gibbon, J., & Allan, L. (Eds.). (1984). *Time and time perception* (Vol. 423). New York: New York Academy of Sciences.

Gibbon, J., & Balsam, P. (1981). Spreading association in time. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 219–253). New York: Academic Press.

Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar timing in memory. *Annals of the New York Academy of Sciences, 423,* 52–77.

Holland, P. C. (1983). Occasion-setting in Pavlovian feature positive discriminations. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner

(Eds.), *Quantitative analyses of behavior: Discrimination processes* (Vol. 4, pp. 183–206). New York: Ballinger.

Holland, P. C. (1985). The nature of conditioned inhibition in serial and simultaneous feature negative discriminations. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 267–297). Hillsdale, NJ: Erlbaum.

Holland, P. C. (1998). Temporal control in Pavlovian occasion setting. *Behavioral Processes, 44,* 225–236.

Holland, P. C., Hamlin, P. A., & Parsons, J. P. (1997). Temporal specificity in serial feature-positive discrimination learning. *Journal of Experimental Psychology: Animal Behavior Processes, 23,* 95–109.

Jenkins, H. M. (1985). Conditioned inhibition of keypecking in the pigeon. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 327–353). Hillsdale, NJ: Erlbaum.

Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review, 94,* 455–468.

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

Miller, R. R., & Barnet, R. C. (1993). The role of time in elementary associations. *Current Directions in Psychological Science, 2,* 106–111.

Rescorla, R. A. (1972). "Configural" conditioning in discrete-trial bar pressing. *Journal of Comparative and Physiological Psychology, 79,* 307–371.

Rescorla, R. A. (1985). Conditioned inhibition and facilitation. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 299–326). Hillsdale, NJ: Erlbaum.

Rescorla, R. A. (1991a). Combinations of modulators trained with the same and different target stimuli. *Animal Learning and Behavior, 19,* 355–360.

Rescorla, R. A. (1991b). Transfer of inhibition and facilitation mediated by the original target stimulus. *Animal Learning and Behavior, 19,* 65–70.

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.

Savastano, H. I., & Miller, R. R. (1998). Time as content in Pavlovian conditioning. *Behavioral Processes, 44,* 147–162.

Skinner, B. F. (1938). *The behavior of organisms.* New York: Appleton-Century-Crofts.

Staddon, J. E. R., & Higa, J. J. (1993). Temporal learning. In D. Medin (Ed.), *The psychology of learning and motivation* (Vol. 27, pp. 265–294). New York: Academic Press.

Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review, 88,* 135–170.

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.

Swartzentruber, D. E. (1995). Modulatory mechanisms in Pavlovian conditioning. *Animal Learning and Behavior, 23,* 123–143.

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.

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, NJ: Earlbaum.

Received October 6, 2006
Revision received March 2, 2007
Accepted March 5, 2007 ■