

Intertrial unconditioned stimuli differentially impact trace conditioning

Douglas A. Williams¹ · Travis P. Todd¹ · Chrissy M. Chubala¹ · Elliot A. Ludvig²

Published online: 5 August 2016
© Psychonomic Society, Inc. 2016

Abstract Three experiments assessed how appetitive conditioning in rats changes over the duration of a trace conditioned stimulus (CS) when unsignaled unconditioned stimuli (USs) are introduced into the intertrial interval. In Experiment 1, a target US occurred at a fixed time either shortly before (embedded), shortly after (trace), or at the same time (delay) as the offset of a 120-s CS. During the CS, responding was most suppressed by intertrial USs in the trace group, less so in the delay group, and least in the embedded group. Unreinforced probe trials revealed a bell-shaped curve centered on the normal US arrival time during the trace interval, suggesting that temporally specific learning occurred both with and without intertrial USs. Experiments 2a and 2b confirmed that the bulk of the trace CS became inhibitory when intertrial USs were scheduled, as measured by summation and retardation tests, even though CS offset evoked a temporally precise conditioned response. Thus, an inhibitory CS may give rise to new stimuli specifically linked to its termination, which are excitatory. A modification to the microstimulus temporal difference model is offered to account for the data.

Keywords Temporal conditioning · Trace conditioning · Conditioned inhibition · Timing · Temporal-difference models

Electronic supplementary material The online version of this article (doi:10.3758/s13420-016-0240-3) contains supplementary material, which is available to authorized users.

✉ Douglas A. Williams
d.williams@uwinnipeg.ca

¹ Department of Psychology, University of Winnipeg, Winnipeg, MB R3B 2E9, Canada

² University of Warwick, Coventry, UK

It has long been known that *whether* an association is acquired depends on how closely two distinguishable events occur in time. Research in this area owes a debt to Pavlov (1927, pp. 39–41), who reported that the interval between the conditioned stimulus (CS) and unconditioned stimulus (US) determines whether dogs salivate in anticipation of an acid US following the offset of a tactile CS. Short gaps of a few seconds between CS offset and US delivery resulted in a salivary conditioned response (CR), starting during the CS itself and extending into the CS–US gap, which he labeled the *short-trace reflex*. As the gap was lengthened, a CR occurred during the CS–US gap but not during CS itself. This suggested to Pavlov that central nervous system activity tied to the recent removal of the CS, and continuing for a time, could support a CR if the gap were not too long. He called this phenomenon the *long-trace reflex*.

Subsequent research revealed that *when* a CR is initiated or reaches its maximum level is also affected by the CS–US interval (Bitterman, 1964). Although response latencies often migrate closer and closer over trials to the onset of the effective cue (Schneiderman, 1966), the peak of responding typically builds to asymptotic levels at the experimentally designed CS–US interval (Smith, 1968), including during trace conditioning (Kehoe, Ludvig, & Sutton 2009, 2014). Operant trace conditioning experiments have extended the work of Pavlov (1927) by revealing Gaussian-like response functions centered on the time of reinforcement on unreinforced test trials with minimal responding during the stimulus (Buhusi & Meck, 2000). Echoing the conclusions of an earlier review by Mackintosh (1974, pp. 61–62), Balsam, Drew, and Gallistel (2010) cited additional evidence that CRs early in training are frequently first initiated near the arrival time of the US. For example, they noted that when the CS–US interval is shifted before a CR is acquired, the original and shifted CS–US intervals may both support peaks in responding (Ohya & Mauk, 2001).

These familiar temporal conditioning phenomena are often used as a platform for testing real-time computational models. One example of such a theory is a relatively new application (Ludvig, Sutton, & Kehoe, 2008, 2012) of the temporal difference (TD) algorithm (Sutton & Barto, 1987, 1990). This family of algorithms derives its name from the idea that organisms predict future states from information in the current moment and bring their predictions as close to reality as possible using an error-correction approach (e.g., Rescorla & Wagner, 1972). Learning occurs when the organism is surprised by an event it did not expect, such as US occurrence. TD models attracted widespread scientific interest in the 1990s with the discovery that dopamine neurons in the monkey brain respond to surprising reward-predicting events (Schultz, Dayan, & Montague 1997).

To encode whether the US occurred, the TD approach assumes the temporally discounted value of the US serves as corrective feedback during each moment in time, where proportionately greater associative strength is supported closer to the time of anticipated US arrival. This discounting mechanism permits responding to propagate backward over trials toward the onset of the CS from the US delivery time. Excitation is passed by temporal contiguity whereby each preceding time step of the CS gradually acquires a level of associative strength that approaches but never fully equals that obtained by the subsequent time step. To encode when the US occurred, the CS is conceived of as a set of temporally defined stimuli with the potential to enter into association with the US. Early implementations of TD approach assumed the temporally defined stimuli had equal potential across the duration of the CS (Moore, Choi, & Brunzell, 1998; Sutton & Barto, 1987, 1990). This so-called complete serial compound (CSC) representation is limited in that it assumes perfect timing and perfect temporal differentiation (see Ludvig et al., 2012). By contrast, in the microstimulus TD model, the CS is thought to trigger smeared representations of temporally defined stimuli generated at distinct times within the duration of the CS (Ludvig et al., 2008, 2012). Each successive microstimulus is increasingly less intense and less temporally specific as the CS is traversed from beginning to end. This permits the time-locked peak CR to diminish in magnitude with increasing interstimulus intervals and allows for increasing generalization across nearby time points.

According to the TD models, the momentary ability of the CS to support a CR depends on the summed associative strength of the microstimuli active at the prescribed time. Similar to the Rescorla–Wagner model, when multiple memory traces from different CSs are simultaneously present, they compete for associative strength, resulting in cue-competition effects, such as blocking (Barnet, Grahame, & Miller, 1993) and conditioned inhibition (Williams, Johns, & Brindas 2008).

It might not be immediately obvious how a model so dependent on contiguity could ever explain trace conditioning.

Following Pavlov's lead, TD models as a class (Ludvig et al., 2008, 2012; Sutton & Barto, 1987, 1990) assume that the offset of a trace CS activates new stimuli that can be differentiated from those of the intertrial interval (ITI). These new stimuli may then persist into the gap and function like a second CS, becoming a direct associate of the US. As the gap is lengthened, the asymptotic level of conditioning at CS offset should diminish, and there should be less associative strength available to spread from CS offset toward CS onset. Furthermore, as a one-time event, CS offset is not expected to continuously generate stimuli during the trace interval and therefore should not "bridge the gap" as effectively as a nominal CS does (Gibbs, Kehoe, & Gormezano, 1991; Kaplan & Hearst, 1982). In summary, the model treats trace conditioning as if it were a variant of the serial conditioning procedure (Kehoe, 1979). The exteroceptive trace stimulus plays the role of the initial CS and the transient stimulus produced at its termination is the terminal CS, which is then reinforced with the US.

The combination of moment-to-moment learning supposed by the TD approach and cue competition leads to some interesting predictions for trace conditioning. For example, according to the CSC and microstimulus TD models (see Ludvig et al., 2009, for a full derivation of this prediction), the initial part of a trace CS should be either excitatory or neutral, and generally not inhibitory, unless the ambient stimuli of the experimental context are reasonably excitatory. Increasing levels of contextual excitation are expected to cause the early part of the trace CS to become increasingly inhibitory. This leads to an interesting and potentially falsifiable prediction. Without modification, TD models do not permit a trace CS be inhibitory across its whole duration and contemporaneously evoke a temporally specific CR in the gap. Pavlov's (1927) observation of a long trace reflex with minimal responding during the CS suggests the model might be wrong. However, he did not assess the trace CS for inhibition, and the CS might have been neutral right up until it terminated.

The aim of these experiments was to evaluate the foregoing possibility by testing a trace CS (trained with or without intertrial USs) for conditioned inhibition and the ability to trigger a CR in the gap. We chose intertrial USs presented at random times because they are perhaps the most straightforward and effective means to drive up contextual excitation relative to a no-ITI USs baseline (Williams, Lawson, Cook, Mather, & Johns, 2008). All three experiments used an appetitive conditioning procedure with food-restricted rats. Head-entry times are a sensitive measure of temporally defined responding and are normally maximal at the time of US delivery (Williams, Chubala, Mather, & Johns, 2009). Experiment 1 characterized temporally graded responding in trace conditioning in the presence or absence of intertrial USs. Using summation and retardation tests, respectively, Experiments 2a and 2b assessed when conditioned inhibition

developed within a trace conditioning trial in the presence of intertrial USs. As previously derived, the microstimulus TD model predicts that inhibition early in the CS should give way to excitation late in the CS, followed by a further increment at CS offset. On the other hand, it is conceivable the trace CS might become inhibitory as a whole: A temporally specific CR might be observed at the arrival time of the US in the gap, while the bulk of CS is inhibitory.

Experiment 1

Experiment 1 examined the effects of intertrial USs on second-by-second levels of responding in trace conditioning. Previous experiments from our laboratory have used a standard delay conditioning procedure (the US coincident with CS termination; Williams, MacKenzie, & Johns, 2010), and a procedure in which the CS is extended past the occurrence of the US termination (the US is embedded within the CS; Williams et al., 2008). These delay and embedded procedures were included in Experiment 1 as important comparison points. Temporally specific responding develops during both embedded and delay CSs in the presence of intertrial USs. Williams et al. (2009) reported head-entry times in acquisition peaked above the contextual baseline at specific CS–US intervals when pellet delivery occurred either 10, 30, 60, or 90 s after the onset of the 120-s CS. Extinction tests also revealed spike-like changes (during a single second) in head-entry times when CS onset and offset were associated with pellet delivery at 0 and 120 s, respectively, which they attributed to momentary changes in associative strength. Under an embedded CS–US relationship, the number of sessions to acquisition also increases as the rate of intertrial USs is increased (Williams & Lussier, 2011).

Experiment 1 used a 2×3 factorial design in which half of the rats in each group received no intertrial USs, whereas the other half received a long-run average of one probabilistically scheduled intertrial US every 30 s. Groups of rats also differed in the arrival time of the target US after the onset of the 120-s white noise CS: either 110 (embedded), 120 (delay), or 130 s (trace). Figure 1 summarizes the design of the experiment, which also included probe trials without the fixed-time US after acquisition had occurred. Three main hypotheses of the microstimulus TD model were tested in Experiment 1. First, in the presence of intertrial USs, anticipatory responding should decrease during the initial portion of the CS in the trace group as a potential sign of conditioned inhibition. Second, there should be increased responding toward the latter part of the trace CS in anticipation of the upcoming US despite the presence of intertrial USs. Third, we expected that post-CS responding in the trace groups should be bell shaped, centered on the normal arrival time of the US, both with and without intertrial USs. These expectations are clearly related to the

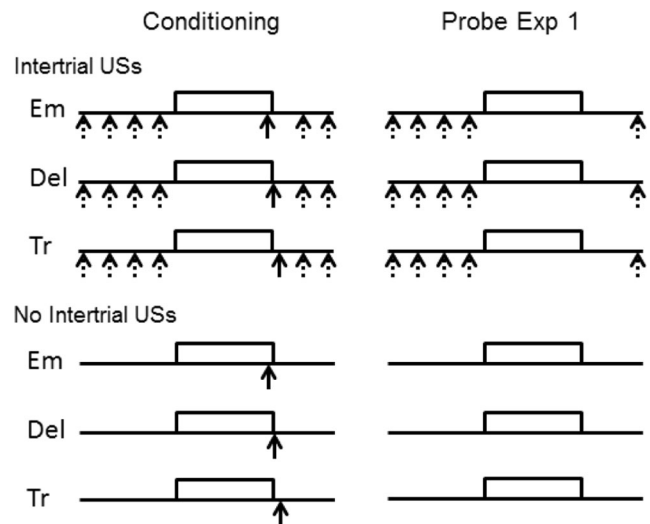


Fig. 1 A schematic depiction of the conditioning and probe trials used in Experiment 1. The *dashed arrows* represent intertrial USs presented randomly in time (1 every 30 s, on average), and the *solid arrows* represent the fixed-time US (withheld on probe trials). Note the fixed-time US appears at slightly different times in the three types of conditioning trials (embedded, delay, and trace)

general prediction that adding intertrial USs should encourage the early part of the trace CS to become inhibitory, but should not undermine anticipatory responding at the target US arrival time.

Method

Subjects

Forty-eight experimentally naïve male Sprague-Dawley rats (*Rattus norvegicus*) served as subjects. They were obtained from Charles River Canada, St. Constant, QC, Canada, and were approximately 90 days old and 250 g upon arrival. The rats were housed in pairs in solid-bottom plastic cages in a colony room with artificial lighting from 0700–1900 hr. Continuous access to water and food occurred during a 2-week acclimatization period. The rats were then food restricted to 80 % of their free-feeding weights. Conditioning sessions occurred during the light portion of the light–dark cycle.

Apparatus

Subjects were trained in one of eight identical chambers (MED Associates, Georgia, VT) measuring 30-cm length \times 25-cm width \times 32-cm height. The chambers were housed in separate, ventilated cabinets (Grason-Stadler, West Concord, MA), which minimized outside light and sound. The front and back panels of the chambers were made of aluminum, and the side panels and ceiling were made of transparent acrylic. The floor consisted of 19 stainless steel rods, 0.5 cm in diameter,

running parallel to the front panel. The CS was white noise (86 dB, Scale A) delivered from a speaker anchored 3 cm above the center point of the ceiling. The rats could gain access to the food pellet US (Formula 21, Bio-Serv, Frenchtown, NJ) through a 5.0×5.0 -cm opening in the middle of the front panel, located 2.0 cm above the grid floor. Pellet deliveries were triggered by a sound-attenuated dispenser (ENV-203, MED Associates, Georgia, VT), which produced a mechanical sound of 54 dB when operated in the absence of the noise of the ventilation fans (64 to 68 dB). A computer equipped with MED-PC IV software (MED Associates, Georgia, VT) controlled the presentation of the CS and the US and monitored head-entry behavior. When the rat's head entered the opening to the recessed food trough, it interrupted an infrared photobeam. The amount of time the photobeam was interrupted on a second-by-second basis served as the dependent variable. On this measure, if the rat's head remained in the food trough for the entire second, it received a maximal score of 1.0 s, whereas a minimum score of 0.0 s was received if the beam was never broken. The number of beam interruptions per second did not matter.

Procedure

Before beginning the experiment, the rats were allowed to consume about 20 food pellet USs in their home cages from a dish. Next, they were randomly assigned to one of six groups ($n = 8$) according to a 2 (intertrial US: no-ITI USs or ITI USs) \times 3 (relationship: embedded, delay, or trace) factorial design. Group labels were no-ITI USs/Em (embedded), no-ITI USs/Del (delay), no-ITI USs/Tr (trace), ITI USs/Em (embedded with ITI pellets), ITI USs/Del (delay with ITI pellets), and ITI USs/Tr (trace with ITI pellets). Groups were counterbalanced across specific chambers and running squads. The target US occurred at 110 (Em), 120 (Del), or 130 s (Tr) after the onset of the 120-s CS. The first 10 s after the CS was free of USs in all groups (serving as the gap in the Tr groups). For the trace groups only, the target US was presented 10-s post-CS. For all groups, the ITI was initiated 10 s after CS termination and averaged 340 s (nonuniform distribution with a 120-s minimum and a 460-s maximum). Thus, the interval between subsequent CS presentations was the same for all groups on average. For the ITI USs condition, a long-run average of one intertrial US occurred every 30 s. Intertrial pellet USs were scheduled by programming US deliveries with a probability of 0.033 during each second of the ITI. Although the number of intertrial pellets received in a given ITI could vary, the long-run average was very close to the desired average interval for all subjects. The no-ITI US groups received a pellet-free ITI. There were eight trials during each of the 36 sessions of conditioning. The conditioning parameters were chosen to be as similar as possible to those used by Williams et al. (2009; Williams et al. 2010). A single

probe trial was introduced at a random point into each of the last eight sessions as a ninth trial. This probe trial was identical to the other trials except for the absence of the target US and the withholding of any intertrial USs normally scheduled in the 30-s post-CS period (triple the value of the 10-s gap used in the Tr groups). Intertrial USs were permitted throughout the ITI until right before CS onset. The normally 4,100-s sessions were lengthened by 590 s (440-s ITI, 120-s CS, 30-s post) to accommodate the extra probe trial.

Results and discussion

Two of the three predictions made at the outset received support. In acquisition, the presence of intertrial USs markedly decreased anticipatory head-entry behavior *during* the CS relative to the ITI, and especially so in the group trained with a trace relationship (Prediction 1). Responding remained very low for the duration of the trace CS and did not increase greatly across the CS, suggesting the entire CS rather than just the initial part might be inhibitory (contra Prediction 2). On unreinforced probe trials, responding spiked at the offset of the CS in the delay group, and there was a bell-shaped distribution of responding in the CS–US gap in the trace groups (Prediction 3).

Acquisition

Figure 2 shows responding during the last four-session block of acquisition with no-ITI USs (top panel) and with ITI USs (bottom panel). To permit within-trial responding to be displayed in its entirety, the data are averaged over 5-s bins. The figure also includes the last 10-s of the ITI. During this period, as might be expected, responding in groups trained with ITI pellets was somewhat greater than those trained without them. A 2 (intertrial US: no-ITI USs or ITI) \times 3 (relationship: Em, Del, Tr) \times 2 (5-s bin) ANOVA performed on the last 10-s of the pre-CS period revealed a marginally significant main effect of intertrial US, $F(1, 42) = 3.63$, $MSE = 0.058$, $p < .07$, but no other main effects or interactions.

The primary finding was that anticipatory head entries during the trace CS, and less so during the delay CS, were reduced in the presence of intertrial USs compared to their absence. To assess temporal conditioning, we checked for the level of head-entry behavior in the 101- to 110-s interval, which spanned the period immediately prior to the delivery of the target US in the embedded groups. A 2 (intertrial US) \times 3 (relationship) \times 2 (bin) ANOVA revealed main effects for intertrial US, $F(1, 42) = 13.55$, $MSE = 0.07$, $p < .0001$, $\eta_p^2 = 0.24$, 95 % CI [.05, .43], and relationship, $F(2, 42) = 7.66$, $MSE = 0.07$, $p < .001$, $\eta_p^2 = 0.27$, 95 % CI [.05, .44], as well as an intertrial US \times relationship interaction, $F(2, 42) = 5.04$, $MSE = 0.07$, $p < .05$, $\eta_p^2 = 0.19$, 95 % CI [.01, .37]. To examine




Fig. 2 Terminal levels of head-entry responding averaged across the last four sessions of acquisition in Experiment 1. Head-entry times are shown relative to the time of CS onset (0 s). Separate panels show data from the pellet free (no-ITI USs) and pellet rich (ITI USs) conditions. The data in each panel are separated by CS–US relationship: embedded (Em), delay (Del), or trace (Tr). *Dashed lines* at 110 s (Em) and 130 s (Tr) show US arrival times, and the *solid line* at 120 s shows both CS termination (all groups) and the arrival of the target US (Del)

the source of the interaction, we tested for the simple effect of relationship within each level of the ITI factor. A simple effect was found only in ITI USs condition, $F(2, 21) = 12.78$, $MSE = 0.07$, $p < .001$, $\eta_p^2 = 0.55$, 95 % CI [.19, .70]. Here, responding in the ITI USs/Em group was greater than in the ITI USs/Del group, $F(1, 14) = 5.47$, $MSE = 0.086$, $p < .05$, $\eta_p^2 = 0.28$, 95 % CI [0, .55], which in turn was greater than in the ITI USs/Tr group, $F(1, 14) = 5.62$, $MSE = 0.07$, $p < .05$, $\eta_p^2 = 0.29$, 95 % CI [0, .56]. Thus, while groups trained without intertrial pellets were not significantly different during this period, responding in the ITI USs/Tr group was depressed relative to the ITI USs/Em and ITI USs/Del groups.

Probe

Figure 3 displays responding in close proximity to the US arrival time during the probe trials included in Sessions 29–36. Responding is shown on a fine-grained second-by-second basis. The figures are centered on the trained arrival time of the US (± 20 s). Note, unlike Fig. 2, the data are aligned to the normal US arrival time. The period covers a limited window of time wherein temporally controlled responding was observed in the Del and Tr groups, thus confirming our third prediction.

We investigated the temporal precision of the CR by searching for the maximum response (y -axis) in individual subjects and then recording the time of the maximum on the x -axis (where 0 = US arrival). A plot of the maximums is

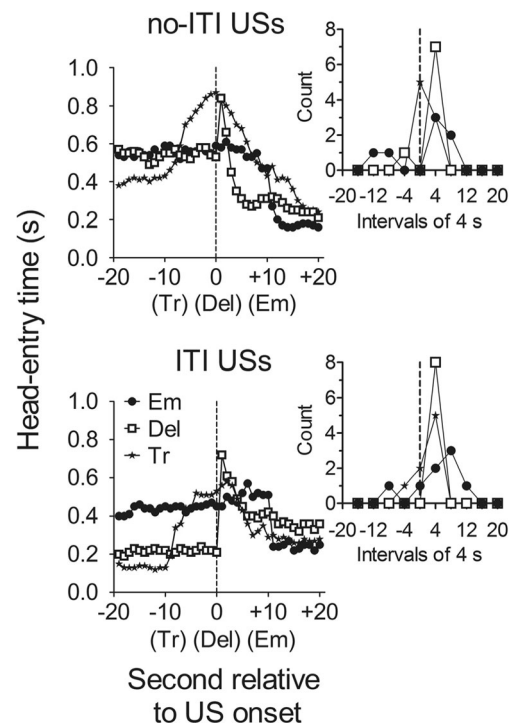


Fig. 3 Responding during the probe trials for the no-ITI USs and ITI USs conditions in Experiment 1 is shown, centered on the scheduled arrival time of the target US (0 s, dashed vertical line). The CS–US relationship, embedded (Em), delay (Del), or trace (Tr) is indicated. The times of CS termination are shown in *brackets* under the x -axis at the 10-s (Em), 0-s (Del), or minus 10-s (Tr) marks. The graph insets are histograms (counts) showing when responding peaked for individual subjects as measured by time of the CR maximum on the x -axis (binned into 4-s periods)

shown in the insets in Fig. 3. Some alternative methods considered included those designed for the operant peak procedure (Church, Meck, & Gibbon, 1994). However, these alternative procedures typically assume a break followed by a reasonably long run of responding, which was only true for the Tr groups. Identical maximums were obtained whether the search window was ± 20 s as used in Fig. 3 or the entire trial window from CS onset to 30 s after CS termination. In all cases, the average x -axis maximums were close to the scheduled arrival time of the US (no-ITI USs/Em = -0.12 s, $SD = 7.57$; no-ITI USs/Del = 0.75 s, $SD = 1.17$; no-ITI USs/Tr = 0.88 s, $SD = 1.81$; ITI USs/Em = 2.75 s, $SD = 5.85$; ITI USs/Del = 1.88 s, $SD = 1.13$; ITI USs/Tr = 0.63 s, $SD = 3.07$). There were no mean differences in the x -axis maximums as a function of the intertrial US and relationship variables. This suggests that responding was temporally defined and peaked near the US arrival time.

To confirm that the spikes in responding at CS termination were restricted to the Del groups, we further compared head-entry times during the last second of the CS and the first second immediately following CS termination. A 2 (intertrial US) \times 3 (relationship) \times 2 (second) ANOVA revealed main effects of intertrial US, $F(1, 42) = 34.93$, $MSE = 0.033$, $p <$

.0001, $\eta_p^2 = 0.45$, 95 % CI [.22, .61], relationship, $F(2, 42) = 8.22$, $MSE = 0.033$, $p < .0001$, $\eta_p^2 = 0.28$, 95 % CI [.06, .45], second, $F(1, 42) = 50.38$, $MSE = 0.008$, $p < .0001$, $\eta_p^2 = 0.55$, 95 % CI [.32, .67], and a relationship \times second interaction, $F(2, 42) = 54.48$, $MSE = 0.008$, $p < .0001$, $\eta_p^2 = 0.72$, 95 % CI [.55, .80]. The interaction was caused by an effect of second in the Del groups, $F(1, 23) = 25.08$, $MSE = .035$, $p < .0001$, $\eta_p^2 = 0.52$, 95 % CI [.24, .69], but not in the Em or Tr groups. It is especially important to note that *both* delay groups showed a spike. Intertrial USs selectively disrupted responding from CS onset in the ITI USs/Del group, but did not disrupt temporal control at CS offset.

In summary, Experiment 1 found that the introduction of intertrial USs attenuated anticipatory responding during the presentation of both the delay and trace CSs. Responding was particularly low across the duration of the trace CS, which is suggestive of conditioned inhibition. The lack of responding to the trace CS in the presence of intertrial USs was strikingly similar to Pavlov's (1927) description of the long trace reflex mentioned in the introduction. In both cases, the stimuli arising at CS termination were strongly associated with the US, although the CS itself evoked little responding. The results of Experiment 1, however, extend those of Pavlov (1927) in an important way. The temporal conditioning occurring after the termination of the trace CS survived the introduction of intertrial USs.

Experiments 2a and 2b

Experiments 2a and 2b used summation and retardation tests (Rescorla, 1969) for conditioned inhibition to assess what portions of the trace CS might be inhibitory in the presence and absence of intertrial USs. Both experiments shared the same 2×2 factorial design as shown in Fig. 4. One factor was whether intertrial USs were present or not during trace conditioning, and the other factor was whether a trace or a novel control CS was tested. Both summation and retardation testing occurred in a distinctive context that had never been associated with random intertrial USs. This aspect of the procedure was expected to minimize any differences in baseline responding at test. No common transfer excitator was trained; thus, the probe trials during retardation of acquisition testing are treated as coming from a separate experiment (Cole, Barnet, & Miller, 1997).

As shown in Fig. 4, the summation test involved separate training of a 150-s transfer excitator in preparation for summation testing, followed by unreinforced probe trials with the transfer excitator presented alone and in combination with the 120-s trace CS (simultaneous onsets). Lower head-entry times to the compound than to the transfer excitator alone would suggest conditioned inhibition, whereas the opposite pattern would suggest conditioned excitation. A switch from conditioned inhibition to

excitation was expected in the ITI USs/Tr condition either during the trace CS or during the gap. The compound in the gap consisted of internal stimuli created by the offset of the trace CS with an actually present excitator. To corroborate the findings of the summation test, Experiment 2b examined whether acquisition would be retarded if the target US were relocated to 10 s after the onset of the trace CS. It would be hard to claim the trace CS simply directed attention away from the transfer excitator, causing suppression during the summation test, if acquisition was also slow on the basis of a relocated US.

Method

Subjects and apparatus

Each experiment included 48 experimentally naive food-restricted rats, cared for in the same manner as in Experiment 1. The experiments were carried out in a set of six conditioning chambers similar to the ones used for Experiment 1. The chambers were narrower, 22.0 cm, and taller, 27.5 cm. A 2.8-W jeweled light served as the visual CS, located 3 cm above the food aperture on the front panel. White noise (86 dB, Scale A) and a 2900-Hz tone (82 dB) served as the trace or control CS in a counterbalanced fashion. Differences in the configuration of the same physical chamber allowed us to create two contexts. One context consisted of a dimly illuminated chamber with alternating 3.7-cm black-and-white-striped sidewalls. Illumination in this context was provided by a 2.8-W shielded houselight, which was located 3 cm beneath the ceiling in the center of the back wall. The floor was comprised of 0.25-cm steel rods spaced 1.50 cm apart, misaligned in a 0.75 cm up-and-down fashion. The second context was unlit and did not have striped walls. The floor in this context was constructed of 0.50-cm steel rods running evenly from one sidewall to the other, spaced 1.10 cm apart. To further increase the distinctiveness of the second context, a 500-ml aluminum bottle filled with frozen water was placed alongside the wall opposite to the chamber door. It provided tactile, thermal, and visual cues to further help the rats discriminate the contexts.

Procedure

All rats first received 36 sessions of trace conditioning. As in Experiment 1, half of the rats received intertrial USs at an average rate of one every 30 s, and the other half received none. The trace conditioning sessions were identical to those described in Experiment 1 with the following two exceptions: First, the white noise and tone CSs were counterbalanced as the trace (Tr) and control (Con) CSs, and two configurations of the conditioning chamber were counterbalanced as the

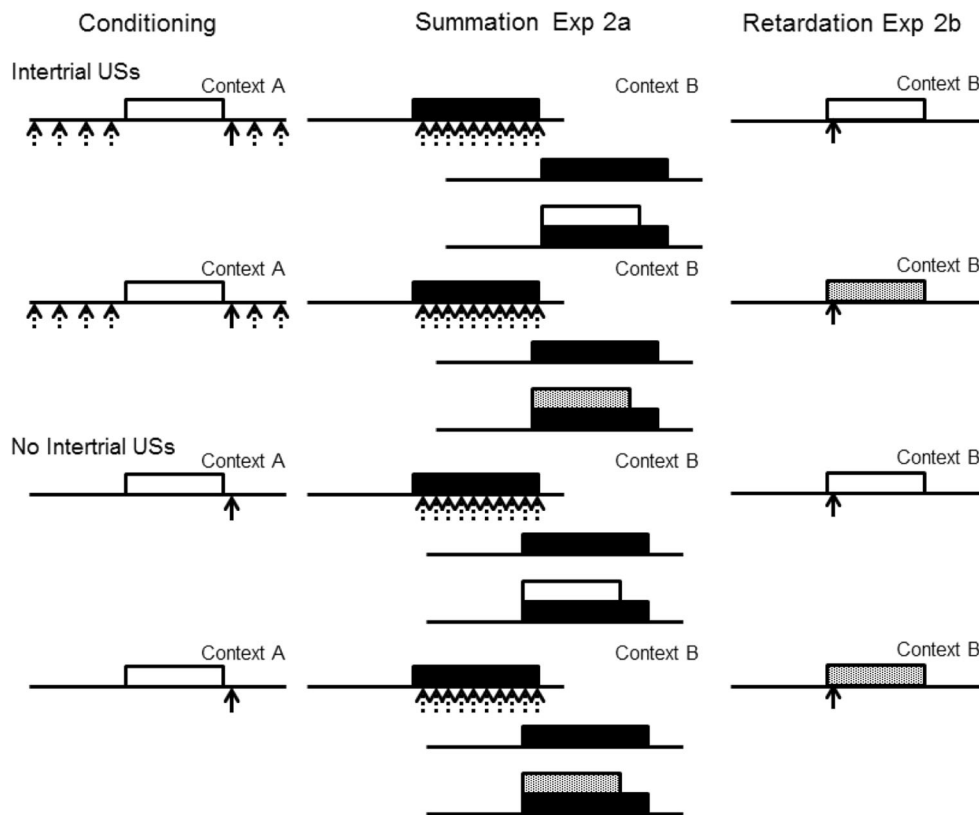


Fig. 4 A schematic of the conditioning and probe trials used in Experiments 2a (summation) and 2b (retardation of acquisition). A trace conditioning stage occurred first in both experiments, followed by either a summation test (excitor is first conditioned followed by unreinforced excitor alone and compound trials) or retardation (US applied 10 s after CS onset rather than after CS offset). The *dashed arrows* represent USs

presented randomly in time (1 every 30 s, on average during conditioning and 1 every 15 s, on average during transfer excitor training; Summation Exp 2a). The *solid arrows* represent the fixed time US. There were four conditions in each experiment (from top to bottom): ITI USs with the Trace stimulus, ITI USs with the control stimulus, no-ITI USs with the Trace stimulus, and no-ITI USs with the control stimulus

conditioning and test contexts. After the completion of 36 sessions of trace conditioning, the rats were assigned to experiment (Experiment 2a or 2b) and tested stimulus (Tr or Con) to equate performance as much as possible while maintaining the counterbalancing of stimuli and contexts. After assignment, they were exposed to the test context for three sessions in preparation for summation and retardation testing. The purpose of the unreinforced context exposure was to reduce the chance that differences in the associative strength of the conditioning context would generalize to the test context. During context exposure, the rats were simply placed in the test context for 1 hr, with no programmed events.

The procedures used in Experiments 2a and 2b differed at this point. In Experiment 2a, the light CS was then conditioned to make it uniformly excitatory. There were eight reinforced presentations of a 150-s light CS in each 1-hr session. During the light CS, the US was delivered at random times at a rate of 0.067 per second (on average one every 15 s). The ITI was US-free. Note, the trace CS was 120 s in duration, making it 30 s shorter than the 150-s light CS. This difference allowed us to then assess the influence of the removal of trace CS on ongoing responding evoked by the light CS during the summation test. The summation test was conducted when the light CS evoked a

moderate and consistent level of responding across its duration, which required seven sessions. The test session began with two reinforced presentations of the light CS during a 15-min “warm-up” period. This was followed by four unreinforced test trials with the light CS (+), and four unreinforced compound trials (-) with the light CS beginning in combination with either the trace CS (for half the subjects) or the novel control CS (for the other half of the subjects). The auditory CS terminated after 120 s, whereas the visual CS terminated after 150 s. This manipulation completed the experimental design, resulting in four groups labeled no-ITI USs/Tr, no-ITI USs/Con, ITI USs/ Tr, and ITI USs/Con. The order of the test trials was randomized.

In Experiment 2b, the retardation test began the day following the unreinforced exposure sessions to the test context. Half of the rats were conditioned with the trace CS, and the other half were conditioned with the control CS. Each session included reinforced and probe trials. On reinforced trials, the target US was delivered 10 s after the onset of the 120-s CS, and no other USs occurred. On probe trials, the target US was omitted and the CS was unreinforced. There were 6 sessions of retardation testing with four reinforced and four probe trials scheduled randomly in each 1-hr session (50 % reinforcement schedule). The same group labels were used as in Experiment 2a.

Results and discussion

Acquisition

The data from unreinforced test trials averaged across Sessions 29–36 are depicted in Fig. 5. The data shown are averaged because the subjects were assigned to experiment and tested stimulus based on their performance.

An initial 2 (intertrial US: present vs. absent) \times 2 (bin: 5-s bins) ANOVA confirmed that pre-CS responding was greater in the presence of intertrial USs than in their absence, $F(1, 94) = 72.83$, $MSE = 0.096$, $p < .0001$, $\eta_p^2 = 0.44$, 95 % CI [.29, .55]. There were no main effects nor interactions involving the bin variable. As shown in Fig. 5, responding dropped shortly after the onset of the trace CS in the ITI USs condition, and there was a lesser increase thereafter in subsequent bins than for the no-ITI USs condition. This observation was supported by a 2 (intertrial US) \times 24 (bin) ANOVA, which produced main effects of intertrial US, $F(1, 94) = 13.91$, $MSE = 1.36$, $p < .001$, $\eta_p^2 = 0.13$, 95 % CI [.03, .26], and bin, $F(23, 2162) = 19.10$, $MSE = 0.29$, $p < .0001$, $\eta_p^2 = 0.17$, 95 % CI [.13, .19], and an interaction of intertrial US \times bin, $F(23, 2162) = 10.49$, $MSE = 0.29$, $p < .0001$, $\eta_p^2 = 0.10$, 95 % CI [.07, .12]. Both the no-ITI USs and ITI USs conditions showed an abrupt increase in responding at CS offset, which subsequently peaked near the trained US arrival time before declining. A 2 (intertrial US) \times 6 (bin) ANOVA applied to the data from the post-CS period found main effects for intertrial US, $F(1, 94) = 6.55$, $MSE = 0.469$, $p < .05$, $\eta_p^2 = 0.07$, 95 % CI [.00, .18], and bin, $F(5, 470) = 76.61$, $MSE = 0.041$, $p < .0001$, $\eta_p^2 = 0.45$, 95 % CI [.38, .50], and an interaction of intertrial US \times bin, $F(5, 470) = 6.92$, $MSE = 0.041$, $p < .0001$, $\eta_p^2 = 0.07$, 95 % CI [.02, .11]. Mean responding in the 126 to 130-s interval (the 130-s mark in Fig. 5, which is just before the trained US arrival time) did not differ in the two conditions, $F(1, 94) = 1.53$, *ns*. Thus, as in Experiment 1, ITI USs reduced responding during the trace CS, which was followed by a robust CR at CS offset.

Summation

Responding to the transfer excitator on its own (+) and in compound (-) with the auditory trace or control CS is shown in Fig. 6. The data are shown in separate panels as a function of whether intertrial USs had occurred in acquisition or not (bottom vs. top panels) and whether the control or trace CS was tested (left vs. right panels). Pre-CS responding was analyzed with a 2 (intertrial US: no-ITI USs vs. ITI USs) \times 2 (stimulus: Con vs. Tr) \times 2 (trial type: "+" vs. "-") \times 2 (5-s bins) ANOVA. This analysis revealed a main effect of intertrial US, $F(1, 44) = 4.12$, $MSE = 0.074$, $p < .05$, $\eta_p^2 = 0.09$, 95 % CI [.00, .26], and an intertrial US \times bin interaction, $F(1, 44) = 4.46$, $MSE = 0.013$, $p < .05$, $\eta_p^2 = 0.09$, 95 % CI [.00, .27]. Pre-CS responding was slightly lower in groups trained previously in another context




Fig. 5 Head-entry responding averaged over the acquisition test trials in Experiments 2a and 2b. Head-entry times are shown relative to the time of trace CS onset (0 s). The label indicates whether the ITI was pellet free (no-ITI USs) or pellet rich (ITI USs). All of the rats received trace conditioning. The dashed line at the 130-s mark shows the learned US arrival time, and the solid line at the 120-s mark shows CS termination

with ITI USs than no-ITI USs, $F(1, 46) = 5.84$, $MSE = .054$, $p < .05$, $\eta_p^2 = 0.11$, 95 % CI [.00, .29]. This unexpected difference is best attributed to the fact that pre-CS responding was near floor levels and was influenced by responding by a few rats. Some type of ‘contrast effect’ between the previously US rich conditioning context and relatively impoverished test context could also have lowered pre-CS responding in the ITI USs groups. The observed pattern is opposite to what might be expected if there had been generalization of excitatory associative strength from the acquisition context to the test context. Nonetheless, the similarity of the results in Experiments 1 and 2a should allay any concern that differences in the excitatory value of the test rather than the acquisition context were responsible for the trace CS becoming inhibitory.

The data of most importance are from the CS period. The main result was that responding in the ITI USs/Tr group

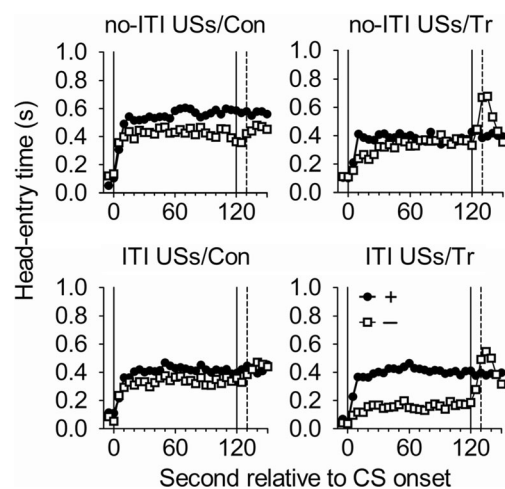


Fig. 6 Head-entry responding averaged over summation test trials (“+” = excitator; “-” = compound) in Experiment 2a. Head-entry times are shown relative to the time of CS onset (0 s). The group labels indicate acquisition treatments (no-ITI USs vs. ITI USs) and whether a control CS (Con) or a trace (Tr) CS was tested as a potential inhibitor. The dashed line at the 130-s mark shows the learned US arrival time for the trace CS, and the solid line at the 120-s mark shows CS termination (trace or control). The transfer excitator begins at the 0-s mark and ends at the 150-s mark

(bottom right panel of Fig. 6), but not the ITI USs/Con group (bottom left panel of Fig. 6), was strongly attenuated on compound trials. A stimulus \times trial type \times bin ANOVA on the data shown in the bottom panels (ITI USs/ Con vs. ITI USs/ Tr) revealed an interaction of stimulus \times trial type, $F(1, 22) = 5.45$, $MSE = .011$, $p < .05$, $\eta_p^2 = 0.20$, 95 % CI [.00, .45]. Responding to the excitator in the ITI USs/Tr group was attenuated throughout by the trace CS compared to the excitator on its own, $F(1, 11) = 31.06$, $MSE = .007$, $p < .001$, $\eta_p^2 = 0.74$, 95 % CI [.32, .85]. No difference between these trials was found in the ITI USs/Con group, $F < 1.0$. There was also less responding on compound trials, but not excitator alone trials, to the trace CS than the control CS, $F(1, 22) = 17.25$, $MSE = 0.008$, $p < .001$, $\eta_p^2 = .44$, 95 % CI [.12, .63]. A stimulus \times trial type \times bin ANOVA on the data shown in the top panels (no-ITI USs/Con vs. no-ITI USs/Tr) revealed an interaction of stimulus \times trial type \times bin, $F(23, 506) = 2.10$, $MSE = 0.031$, $p < .01$, $\eta_p^2 = 0.09$, 95 % CI [.03, .09]. In the no-ITI USs condition, the two-way stimulus \times trial type interaction was not reliable, $F(1, 22) = 1.07$, *ns*, suggesting that negative summation was minimal. Stimulus \times trial type ANOVAs over individual bins only found hints of negative summation at the 10-s, 15-s, and 20-s marks of the trace CS. Here, responding to the compound was less than the excitator alone, smallest $F(1, 11) = 7.28$, $MSE = 0.073$, $p < .05$, $\eta_p^2 = 0.40$, 95 % CI [.01, .65]. Perhaps due to external inhibition (Pavlov, 1927), there was also less responding to the compound than the excitator in no-ITI USs/Con group during a number of bins after the 60-s mark, smallest $F(1, 11) = 4.90$, $MSE = 0.052$, $p < .05$, $\eta_p^2 = 0.31$, 95 % CI [.00, .59].

Statistical analyses confirmed that responding in the post-CS period (last six bins after the 120-s mark) increased rapidly on compound trials in groups tested with the trace CS relative to controls (top panels: no-ITI USs/ Tr vs. no-ITI USs/ Con; bottom panels: ITI USs/Tr vs. ITI USs/Con). Again, stimulus \times trial type \times bin ANOVAs were conducted separately for the top and bottom panels. Both analyses revealed three-way interactions of stimulus \times trial type \times bin, smallest $F(5, 110) = 6.02$, $MSE = 0.010$, $p < .0001$, $\eta_p^2 = 0.22$, 95 % CI [.07, .31]. To investigate the source of the three-way stimulus \times trial type \times bin interactions, we conducted trial type \times bin ANOVAs within each level of stimulus. A trial type \times bin analysis revealed no changes after the removal of the control stimulus (see top left and bottom left panels). On the other hand, after the removal of the trace CS (top right and bottom right panels), there were main effects of bin for both ITI conditions, minimum $F(5, 55) = 4.98$, $MSE = 0.020$, $p < .001$, $\eta_p^2 = 0.31$, 95 % CI [.07, .43], and trial type \times bin interactions for both ITI conditions, minimum $F(5, 55) = 9.10$, $MSE = 0.013$, $p < .0001$, $\eta_p^2 = 0.45$, 95 % CI [.21, .56]. The interactions were caused by more post-trace CS responding to the compound than to the excitator. This can be seen 10 and 15 s after trace CS termination (130- and 135-s marks in Fig. 6) in the no-ITI USs group, smallest $F(1,$

11) = 6.06, $MSE = 0.081$, $p < .05$, $\eta_p^2 = 0.36$, 95 % CI [.00, .62], and at 15 s after CS offset in the ITI USs group, $F(1, 11) = 8.54$, $MSE = 0.020$, $p < .05$, $\eta_p^2 = 0.44$, 95 % CI [.02, .67]. Thus, the sign of the summation effect changed from negative to positive at CS termination in the ITI USs/Tr group, and from neutral to positive in the no-ITI USs/Tr group.

Retardation

The results from the interleaved probe trials of retardation test are shown in Fig. 7 as a function of successive two-session blocks (labeled 1, 2, and 3). The data are displayed only for the times of greatest interest during the early part of the CS. Again, responding is displayed separately for rats previously trained with (bottom panels) and without (top panels) intertrial USs. Left-hand panels show the data from the control CS, whereas right-hand panels show the data from the trace CS. The dashed line is the relocated US delivery time used in the retardation test. As evident from the bottom right panel of Fig. 7, the CS in ITI/Tr group was slower to become associated with the relocated US, corroborating the findings from Experiment 2a that ITI USs resulted in conditioned inhibition during the trace CS.

An examination of pre-CS responding with an intertrial US \times stimulus \times block \times bin ANOVA found main effects for intertrial US, $F(1, 44) = 18.33$, $MSE = 0.033$, $p < .0001$, $\eta_p^2 = 0.29$, 95 % CI [.09, .47], and stimulus, $F(1, 44) = 5.04$, $MSE = 0.033$, $p < .05$, $\eta_p^2 = 0.10$, 95 % CI [.00, .28], and an interaction of intertrial US \times stimulus \times bin, $F(1, 44) = 6.07$, $MSE = 0.001$, $p < .05$, $\eta_p^2 = .12$, 95 % CI [.00, .30]. Although numerical differences in head-entry times were again unimpressive due to the low baseline levels, the pattern was similar




Fig. 7 Retardation of acquisition after a shift in the arrival time of the US is shown for the first (1), second (2), and third (3) block of two sessions. Head-entry times are shown relative to the time of CS onset (0 s), and the dashed line (10 s) shows the US arrival time in test. The labels for each panel indicate whether the ITI in acquisition had been pellet free or pellet rich (no-ITI USs vs. ITI USs), and whether a control or trace (Con vs. Tr) CS was tested as a potential inhibitor

to that found in the summation test with more responding in the no-intertrial USs condition.

Rates of acquisition can be discerned by examining the speed of development of a bell-shaped response distribution centered near the 10-s mark (US delivery time) of the CS over two-session blocks. Separate stimulus (2) \times block (3) \times bin (8) ANOVAs on the data shown in Fig. 7 identified three-way interactions in both the top panels, $F(14, 308) = 3.17$, $MSE = 0.005$, $p < .0001$, $\eta_p^2 = 0.13$, 95 % CI [.03, .16], and in the bottom panels, $F(14, 308) = 3.22$, $MSE = 0.003$, $p < .0001$, $\eta_p^2 = 0.13$, 95 % CI [.03, .16]. To better understand the pattern of differences contributing to the interactions, we examined responding in a single bin just before the arrival time of the US (i.e., the 6- to 10-s time bin, dashed line). These comparisons revealed that responding in the no-ITI USs/Tr group was lower than the no-ITI USs/Con group in Blocks 1 and 2, smallest $F(1, 22) = 8.40$, $MSE = 0.032$, $p < .01$, $\eta_p^2 = 0.28$, 95 % CI [.02, .51]. Thus, the early part of the CS in the no-ITI USs/Tr group appeared somewhat inhibitory, which is consistent with the results of summation testing. By contrast, responding immediately prior to US arrival was consistently lower in the ITI USs/Tr group than the ITI USs/Con group regardless of block, smallest $F(1, 22) = 9.33$, $MSE = 0.023$, $p < .01$, $\eta_p^2 = 0.30$, 95 % CI [.03, .53]. Other comparisons confirmed less responding in the ITI USs/Tr than no-ITI USs/Tr group during all three blocks, smallest $F(1, 22) = 9.33$, $MSE = 0.023$, $p < .01$, $\eta_p^2 = 0.30$, 95 % CI [.03, .53]. This difference suggests that the early portion of the trace CS in the ITI USs/Tr group was strongly inhibitory and much more so than in the no-ITI USs/Tr group.

In summary, both tests led to the same conclusion: The trace CS became strongly inhibitory when trained in the presence of ITI pellets. The results from the summation test also demonstrated trace CS termination evoked a temporally defined CR that positively summated with the transfer excitator. In this case, the termination of the CS caused a shift from negative to positive summation. The added stimulus in the gap on compound trials was presumably the internal trace left by a recently presented inhibitor.

General discussion

These experiments provide an interesting new set of findings about how intertrial USs influence time-based patterns of responding under various CS–US relationships. Experiment 1 found that intertrial USs caused an increasing level of attenuation during the CS itself as the target US was moved in 10-s steps from preceding (embedded relationship) to following CS (trace relationship) the termination of the 120 s. Although intertrial USs caused the most attenuation of responding during the trace CS, a robust CR was still observed in the gap between CS offset and US delivery. This post-CS responding

was not simply the resumption of an expectation of randomly distributed intertrial USs, because the CR peaked exactly 10 s after the trace CS terminated. Our trace conditioning data are reminiscent in some ways of those of Pavlov (1927) with minimal (his) or lesser (ours, ITI USs condition) responding during the CS than the pre-CS period, followed by a CR in the gap. A summation test in Experiment 2a subsequently revealed a trace CS trained in the presence of intertrial USs was strongly inhibitory across its whole duration. Experiment 2b found retardation of acquisition to a US relocated within the early part of the trace CS, providing further evidence of conditioned inhibition. Conditioned inhibition (negative summation) gave way to conditioned excitation (positive summation) only during the gap. Taken together, these data suggest excitatory temporal conditioning triggered by the offset of the trace CS survived the introduction of unsigned USs, although the CS itself became inhibitory.

How well do the predictions of the microstimulus TD model (Ludvig et al., 2008, 2012) described in the introduction fit these data? According to this model, trace conditioning is viewed as a special instance of a serial conditioning procedure (Kehoe, 1979). In serial conditioning, two distinctive CSs follow one another in a fixed sequence, and the terminal CS is reinforced. The suggested parallel is the offset of the trace CS should activate new microstimuli, which function like the terminal CS in a serial conditioning procedure. Associative strength should propagate backward in time following a discount function from the time of US arrival to CS offset (the terminal CS), and then via second-order conditioning from CS offset to CS onset (the initial CS). This state of affairs should change with the introduction of intertrial USs. The CS should become inhibitory, as observed in Experiments 2a and 2b, because it predicts a period of lower reinforcement in the presence of excitatory contextual cues. Post-CS responding is expected to be Gaussian-like because of the consistent occurrence of the target US, irrespective of the presence or absence of random intertrial USs, just as was found in the Tr groups on the probe trials in Experiment 1. By comparison, spike-like responding is predicted at CS termination for the Del groups, as observed in Experiment 1, if it is assumed the subject's receipt of the US occurred slightly after CS termination. However, the more detailed predictions of the model for trace conditioning have mixed consistency with our data. Figure 8 displays formal simulations of the predictions of the microstimulus TD model for Experiment 1 using either a relatively low (right panels = 0.90) or high (left panels = 0.97) discount factor gamma. Gamma is the proportion of associative strength that spreads backward from the current moment to the preceding moment (the discount factor). The effect of changing the value of gamma is readily seen by comparing the left and right upper panels (no-ITI USs controls). With a relatively high gamma, there is less discounting of future reward and more anticipatory responding, which characterizes the




Fig. 8 Simulation of the asymptotic predictions of the microstimulus TD model for Experiment 1. The ITI is modeled as either pellet free (no-ITI USs, top panels) or pellet rich (ITI USs, bottom panels) and the value of future reward either relatively low with more discounting ($\gamma = .90$, left panels) or relatively higher with less discounting ($\gamma = .97$, right panels), as in the original microstimulus TD model. The CS–US relationship is embedded (Em), delay (Del), or trace (Tr). CS onset occurs at 0 s and CS offset occurs at 120 s (solid lines). Stimuli used are CS onset, CS offset, and US. Parameters used are microstimuli per CS = 60; microstimulus width = .08; step size = .005; memory trace decay rate = eligibility trace decay rate = 0.985; stimulus and context presence = 0.8; number of trials = 2,000

data obtained in the no-ITI USs conditions. The MATLAB code for these simulations is available in the [supplemental materials](#).

Unfortunately, the corresponding predictions for the ITI USs condition assuming the same high gamma are less impressive (lower right panel of Fig. 8). Unlike the data, the formal simulations predict a clear temporal pattern, with an increase in associative strength from the initial postonset depression over the course of the CS, irrespective of the CS–US relationship. Decreases in associative strength due to conditioned inhibition during the initial part of the CS are predicted to be followed by an increase to above the contextual baseline for the Em and Del groups. Increasing responding should be seen first in the Em group, then in the Del group, and last in the Tr group. The increase in the Tr group follows because the microstimuli arising later in the CS are paired with the beginning of the trace interval (CS offset), which is excitatory. The expected temporal pattern is consistent with observed responding in the ITI USs/Em group in Experiment 1, but it is somewhat less consistent with responding in the ITI USs/Del group. However, there is a clear mismatch in the ITI USs/Tr group. In Experiment 1, responding did not increase from near-floor levels until very near the termination of the CS, and even then it remained well below pre-CS levels. Likewise, Experiment 2a found negative summation on compound trials throughout the trace CS in the ITI USs/Tr group.

Although experiments from our laboratory have not always found as severe a deficit as depicted in Fig. 3 in delay conditioning (Williams et al., 2010), the microstimulus TD model

clearly underpredicts the attenuation of conditioning that can occur during the late portion of a trace CS. Substantially better fits are not found when the coarseness of the representation of the CS is reduced (e.g., when the number of microstimuli per CS is reduced from 60 to 6; see Ludvig et al., 2012) or the simulations assume preasymptotic data (e.g., 200 vs. 2,000 trials). One solution we have explored assumes that long-duration trace or delay CSs might lose their eligibility to acquire strength via pairings with the secondary reward value conditioned to an upcoming stimulus, the gap. A loss of eligibility in combination with the usual assumption of CS offset as a new event might explain our trace conditioning data.

Consistent with this line of thinking, Ludvig, Sutton, Verbeek, and Kehoe (2009) have recently suggested a related modification to the TD model. In recognition of the need for qualitative change in representation over time, they argue that long-latency temporal elements of the CS might require activation of a collateral brain structure, the hippocampus (Shors, 2004). Hippocampally dependent learning is thought to occur whenever long-latency temporal elements of the CS enter into association with the discounted reward value of the US. These elements are thought to ramp slowly to a low asymptote over the course of the continued presence of a long duration CS and then diminish after reaching their preferred temporal bias point. On the other hand, the short-latency elements are thought to be evoked without hippocampal involvement, and differ in their more narrowed and peaked activation patterns. Thus, one could suppose that serial conditioning is weaker with long-latency microstimuli than with short-latency microstimuli because the hippocampal system is less efficient at secondary learning. This distinction between primary and secondary learning also echoes the primary value and learned value (PVLV) model, a competitor to TD learning for explaining appetitive conditioning in the brain (see O'Reilly, Frank, Hazy, & Watz 2007).

A core principle of TD learning, however, is that primary and secondary learning are identical, making this suggested revision to the model a significant deviation in need of further empirical validation. In addition, the effect of reduced eligibility with trace conditioning should occur both with and without ITI USs. Yet, the animals displayed what would seem to be intact eligibility without ITI USs. Finally, this solution does not address fundamental assumptions about TD learning in this situation but rather the particular way stimuli are encoded in the microstimulus TD model.

Figure 8 might provide a clue to a more promising solution. It is possible the extra USs affected the discount factor gamma. In particular, the introduction of intertrial USs might lead to faster discounting of future reinforcement because pellets are more common in the conditioning session as a whole. Thus, the data of the ITI/USs condition might be more appropriately modeled with a lower gamma ($\gamma = 0.90$; lower left panel of Fig. 8), with a higher gamma used for the no ITI/

USs condition ($\gamma = 0.97$; upper right panel of Fig. 8). Responding still drops post onset with a lower discount rate (as observed), but remain depressed more broadly across the total duration of the CS (as observed). Background levels of responding also diminish somewhat in the ITI USs condition from an otherwise higher level if γ is reduced. It is interesting that the baseline differences caused by the introduction of intertrial USs were not all that large, especially in Experiment 1. This could be taken as further evidence for greater discounting of future reinforcement with frequent intertrial USs. Any model attempting to learn the discounted value of future reward could be modified in this way, including but not restricted to the microstimulus TD model.

We have chosen to highlight the predictions of microstimulus TD model because it makes highly constrained predictions in addition to addressing the larger goal of reconciling error prediction learning with the temporal properties of the CR (Kirkpatrick, 2014). However, our data should also be of interest to those studying temporal learning from other real-time perspectives. Most real-time models, such as the componential standard operating procedure model (e.g., Vogel, Brandon, & Wagner, 2003; Wagner, 1981), also correctly predict that trace procedures are more likely to produce conditioned inhibition when the context is excitatory. These models, however, are also faced with the problem of specifying how an inhibitory CS might trigger a CR in the post-CS period. One possibility is to suppose a reduced level of second-order conditioning in the presence of intertrial USs, following the line of thinking mentioned previously for the PVLV model. This suggestion is tempered by an acknowledgement that the mechanism underlying the effects of intertrial USs remains to be determined. The no-ITI USs versus ITI USs manipulation was employed to increase context excitation, which it did for the most part, but it could also have affected discounting, levels of second-order conditioning, or even the current motivational value of the reinforcer (satiation).

Farther afield are timing models. These models assume that intervals of time are the main content of what is learned (Church et al., 1994; Guilhardi, Yi, & Church, 2007). From this perspective, the ability of the rats to time the arrival of a US after CS termination is at issue. These models make the clear prediction that any event, stimulus onset or offset, may serve to mark the beginning of a fixed interval before US arrival (Buhusi & Meck, 2000). Although these theories are able to accommodate the trace conditioning data found in Experiment 1, it is much less obvious how such an account could handle negative summation (Experiment 2a) or retardation of conditioning (Experiment 2b). Timing models simply do not include the concept of conditioned inhibition (Williams, Johns, et al., 2008) and do not specify when timing will be inhibited. That said, it makes sense that the most proximal stimulus (Fairhurst, Gallistel, & Gibbon, 2003), namely CS offset (Buhusi & Meck, 2000), might have been used by the

rats to time to the arrival of the target US. Such timing might be argued to be a separate and distinct process from the “associative properties” of the CS, perhaps involving different brain mechanisms (Meck, 2006). Given this, Experiments 2a and 2b could be interpreted as providing a striking new dissociation: A CS with inhibitory associative properties concurrently acted as a time marker. Although such a dual-systems approach is certainly a possibility worth acknowledging, the data of these experiments can be explained more parsimoniously through the continuing evolution of a model integrating temporal representations into error-prediction mechanisms.

Author Note This research was supported by a grant from the Natural Sciences and Engineering Research Council of Canada to D. A. Williams.

References

- Balsam, P., Drew, M. R., & Gallistel, C. R. (2010). Time and associative learning. *Comparative Cognition and Behavior Reviews*, *5*, 1–22. doi:10.3819/ccbr.2010.50001
- Barnet, R. C., Grahame, N. J., & Miller, R. R. (1993). Temporal encoding as a determinant of blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, *19*, 327–341. doi:10.1037/0097-7403.19.4.327
- Bitterman, M. E. (1964). Classical conditioning in the goldfish as a function of the CS-US interval. *Journal of Comparative and Physiological Psychology*, *58*, 359–366. doi:10.1037/h0046793
- Buhusi, C. V., & Meck, W. H. (2000). Timing for the absence of a stimulus: The gap paradigm reversed. *Journal of Experimental Psychology: Animal Behavior Processes*, *26*, 305–322. doi:10.1037/0097-7403.26.3.305
- Church, R. M., Meck, W. H., & Gibbon, J. (1994). Application of scalar timing to individual trials. *Journal of Experimental Psychology: Animal Behavior Processes*, *20*, 135–155. doi:10.1037/0097-7403.20.2.135
- Cole, R. P., Barnet, R. C., & Miller, R. R. (1997). An evaluation of conditioned inhibition as defined by Rescorla’s two-test strategy. *Learning and Motivation*, *28*, 323–241.
- Fairhurst, S., Gallistel, C. R., & Gibbon, J. (2003). Temporal landmarks: Proximity prevails. *Animal Cognition*, *6*, 113–120.
- Gibbs, C. M., Kehoe, J. K., & Gormezano, I. (1991). Conditioning of the rabbit’s nictitating membrane response to a CSA-CSB-US serial compound: Manipulations of CSB’s associative character. *Journal of Experimental Psychology: Animal Behavior Processes*, *17*, 423–432. doi:10.1037/0097-7403.17.4.423
- Guilhardi, P., Yi, L., & Church, R. M. (2007). A modular theory of learning and performance. *Psychonomic Bulletin and Review*, *14*, 543–559. doi:10.3758/BF03196805
- Kaplan, P. S., & Hearst, E. (1982). Bridging temporal gaps between CS and US in autoshaping: Insertion of other stimuli before, during, and after CS. *Journal of Experimental Psychology: Animal Behavior Processes*, *8*, 187–203. doi: 1037/0097-7403.8.2.187
- Kehoe, E. J. (1979). The role of CS-US contiguity in classical conditioning of the rabbit’s nictitating membrane response to serial stimuli. *Learning and Motivation*, *10*, 23–38. doi:10.1016/0023-9690(79)90048-1
- Kehoe, E. J., Ludvig, E. A., & Sutton, R. S. (2009). Magnitude and timing of conditioned responses in delay and trace classical conditioning of the nictitating membrane response of the rabbit (*Oryctolagus cuniculus*). *Behavioral Neuroscience*, *123*, 1095–1101. doi:10.1037/a0017112

- Kehoe, E. J., Ludvig, E. A., & Sutton, R. S. (2014). Time course of the rabbit's conditioned nictitating membrane movements during acquisition, extinction, and reacquisition. *Learning & Memory*, *21*, 585–590. doi:10.1101/lm.034504.114
- Kirkpatrick, K. (2014). Interactions of timing and prediction error learning. *Behavioural Processes*, *101*, 135–145. doi:10.1016/j.beproc.2013.08.005
- Ludvig, E. A., Sutton, R. S., & Kehoe, E. J. (2008). Stimulus representation and the timing of reward-prediction errors in models of the dopamine system. *Neural Computation*, *20*, 3034–3054. doi:10.1162/neco.2008.11-07-654
- Ludvig, E. A., Sutton, R. S., & Kehoe, E. J. (2012). Evaluating the TD model of classical conditioning. *Learning & Behavior*, *40*, 305–319. doi:10.3758/s13420-012-0082-6
- Ludvig, E. A., Sutton, R. S., Verbeek, E. L., & Kehoe, E. J. (2009). A computational model of hippocampal function in trace conditioning. *Advances in Neural Information Processing Systems (NIPS-08)*, *21*, 993–1000.
- Mackintosh, N. J. (1974). *The psychology of animal learning*. Oxford: Academic Press.
- Meck, W. H. (2006). Neuroanatomical localization of an internal clock: A functional link between mesolimbic, nigrostriatal, and mesocortical dopaminergic systems. *Brain Research*, *1109*, 93–107. doi:10.1016/j.brainres.2006.06.031
- Moore, J., Choi, J., & Brunzell, D. (1998). Predictive timing under temporal uncertainty: The TD model of the conditioned response. In D. Rosenbaum & A. Collyer (Eds.), *Timing of behavior: Neural, computational, and psychological Perspectives* (pp. 3–34). Cambridge: MIT Press.
- O'Reilly, R. C., Frank, M. J., Hazy, T. E., & Watz, B. (2007). PVLV: The primary value and learned value Pavlovian learning algorithm. *Behavioral Neuroscience*, *121*, 31–49. doi:10.1037/0735-7044.121.1.31
- Ohyama, T., & Mauk, M. D. (2001). Latent acquisition of timed responses in cerebellar cortex. *The Journal of Neuroscience*, *21*, 682–690.
- Pavlov, I. P. (1927). *Conditioned reflexes*. Oxford: Oxford University Press.
- Rescorla, R. A. (1969). Pavlovian conditioned inhibition. *Psychological Bulletin*, *72*, 77–94. doi:10.1037/h0027760
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. R. Prokasy (Eds.), *Classical conditioning II* (pp. 64–99). New York: Appleton-Century-Crofts.
- Schneiderman, N. (1966). Interstimulus interval function of the nictitating membrane response of the rabbit under delay versus trace conditioning. *Journal of Comparative and Physiological Psychology*, *62*, 397–402. doi:10.1037/h0023946
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
- Shors, T. J. (2004). Memory traces of trace memories: Neurogenesis, synaptogenesis, and awareness. *Trends in the Neurosciences*, *27*, 250–256. doi:10.1016/j.tins.2004.03.007
- Smith, M. C. (1968). CS-US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology*, *66*, 679–687. doi:10.1037/h0026550
- Sutton, R. S., & Barto, A. G. (1987). A temporal-difference model of classical conditioning. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society* 355–378.
- 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: Bradford/MIT Press.
- Vogel, E. H., Brandon, S. E., & Wagner, A. R. (2003). Stimulus representation in SOP II: An application to inhibition of delay. *Behavioural Processes*, *62*, 27–48. doi:10.1016/S0376-6357(03)00050-0
- Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale: Erlbaum.
- Williams, D. A., Chubala, C. M., Mather, A. A., & Johns, K. W. (2009). Interstimulus interval and delivery cues influence timing of fixed target pellets by rats. *Learning and Motivation*, *40*, 394–407. doi:10.1016/j.lmot.2009.07.001
- Williams, D. A., Johns, K. W., & Brindas, J. L. (2008). Timing during inhibitory conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*, 237–246. doi:10.1037/0097-7403.34.2.237
- Williams, D. A., Lawson, C. L., Cook, R., Mather, A. A., & Johns, K. W. (2008). Timed excitatory conditioning under zero and negative contingencies. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*, 94–105. doi:10.1037/0097-7403.34.1.94
- Williams, D. A., & Lussier, A. L. (2011). Intertrial pellets influence the acquisition and expression of timed appetitive responding in rats. *Learning and Motivation*, *42*, 300–312. doi:10.1016/j.lmot.2011.06.003
- Williams, D. A., MacKenzie, H. K., & Johns, K. W. (2010). Intertrial unconditioned stimuli preferentially interfere with delay conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *36*, 232–242. doi:10.1037/a0016922