Honey, Rob , Dwyer, Dominic and Navarro, Victor 2025. The spatio-temporal dynamics of conditioned behavior: First-order and higher-order conditioning. Journal of Experimental Psychology: Animal Learning and Cognition 10.1037/xan0000392

The spatio-temporal dynamics of conditioned behavior:

First-order and higher-order conditioning

Short title: Conditioned behavior

Paper accepted for publication in February 2025:

Journal of Experimental Psychology: Animal Learning and Cognition

Address for correspondence: R. C. Honey, School of Psychology, Cardiff University, Tower Building, 70 Park Place, Cardiff, CF10 3AT, UK; Tel: +44 (0)29 20875868; E-mail: Honey@cardiff.ac.uk

Abstract

Pavlovian conditioning procedures generate spatially and temporally distinct behaviors. For example, after rats have received pairings of a lever with food, they approach the food well during the lever (called goal-tracking) and interact with it (called sign-tracking); with these two spatially distinct behaviors being distributed differently across the temporal duration of the lever. Experiment 1 assessed the development of these spatio-temporally defined behaviors during first-order conditioning, as a function of the sequence in which the lever and food occurred (lever→food or food→lever) and the interval between them (1s or 11s). In Experiment 2, the same rats received higher-order conditioning trials in which an auditory stimulus was paired with the lever and the emergence of goal-tracking to the auditory stimulus was assessed. The results of Experiments 1 and 2 revealed dissociations between where and when learning was evident during first- and higher-order conditioning, underscoring the need for models of Pavlovian conditioning to explain both the nature and timing of different conditioned responses.

Keywords: rat, goal-tracking, sign-tracking, timing, HeiDI

Most descriptions of Pavlov's conditioning studies with dogs emphasize the fact that a conditioned stimulus (CS; e.g., a ticking metronome) comes to elicit a similar response to the unconditioned stimulus (US; e.g., the delivery of food into a bowl) as a consequence of their pairing. In these studies, dogs were placed in a harness constraining their movements to enable the accurate collection of drops of saliva from the dogs' mouths during the CS (the conditioned response or CR). However, such descriptions of the consequences of conditioning are a caricature of Pavlov's original observations. For example, the presentation of conditioned stimuli also prompted orienting to the food bowl, and when the dogs were unharnessed they approached the location in which the CS was located as it was moved from one position in the experimental room to another (e.g., Pavlov, 1928; p.168; see also, Zener, 1937). Moreover, the salivary CR, which was assessed across the temporal duration of the CS, was less marked at the start of the CS than at its end (Pavlov, 1928, p. 149); a phenomenon known as inhibition of delay.

The fact that Pavlovian conditioning procedures generate spatio-temporally distinct behaviors, including those related to the nature of the US (e.g., the delivery of food into a bowl) and the CS (e.g., its location in the experimental room), has been confirmed in many preparations (e.g., Holland, 1977, 1984; Iliescu, Dwyer & Honey, 2020; Timberlake & Grant, 1975). Such findings mean that key questions about the conditions, content and mechanisms of Pavlovian conditioning (see Dickinson, 1980; Rescorla, 1988) cannot be answered adequately by measuring a single conditioned behavior: The answers based on one measure might not generalize to other behaviors either within a given preparation (i.e., be internally coherent; see Patitucci et al., 2016) or across preparations (i.e., be externally coherent). Yet this is precisely the approach that has dominated assessment of predictions generated by theoretical models of conditioning processes (e.g., Mackintosh, 1975; Pearce & Hall, 1980; Pearce, 1994; Rescorla & Wagner, 1972; Stout & Miller, 2007; Wagner, 1981), from conditioned suppression in rats (e.g., Dickinson, Hall & Mackintosh, 1976; Hall & Pearce, 1979; Kasprow, Schachtman & Miller, 1987; Rescorla, 1968) to autoshaped key-pecks in pigeons (e.g., Leyland & Mackintosh, 1978; Swan & Pearce, 1987). To be clear, this methodological approach to testing predictions derived from theoretical models has an important limitation: Because the models have little to say about how their underlying constructs (e.g., the strength of a CS→US association, V_{CS}) affect the spatio-temporal properties of conditioned behaviors, it is unclear whether their predictions are internally coherent. Do their predictions hold across all measures of conditioning in each preparation or only a subset, and if only a subset, then why that subset? Recent research reinforces this critique.

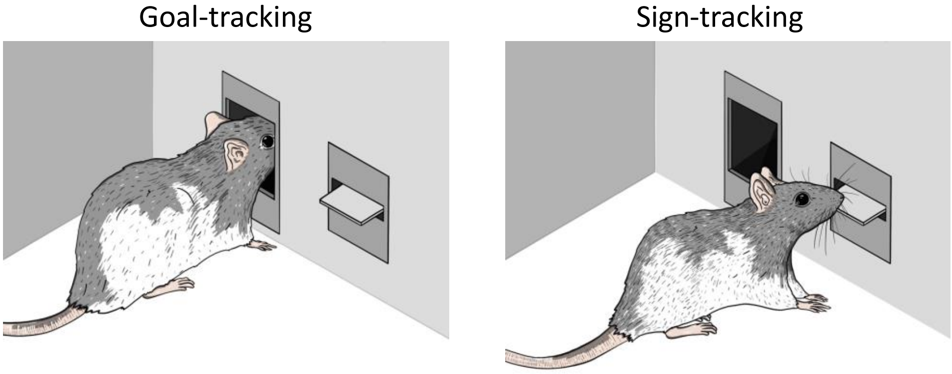

Figure 1. Two responses generated by a Pavlovian conditioning procedure in rats in which the temporary insertion of a lever into an experimental chamber precedes the delivery of food into a recessed food well. The rat in the left panel is entering a food well during the lever (called goal-tracking) and the rat in the right panel is interacting with a lever (called sign-tracking). Adapted from: Navarro, V.M., Dwyer, D. M., & Honey, R.C. (2023). Prediction error in models of adaptive behavior. *Current Biology*, *33*, 4238-4243.

Navarro, Dwyer and Honey (2023) reported the results of a series of experiments using an autoshaping procedure with rats in which the temporary introduction of a lever into the experimental apparatus served as the CS and the delivery of a food pellet into a food well was the US. The procedure is an example of Pavlovian conditioning, with the US being delivered independently of the behavior of the rat. In one experiment, two groups of rats received presentations of two levers, L1 and L2, which were separately introduced into the experimental chamber for 10s and then withdrawn. In group Forward, L1 was followed by the delivery of a food pellet into a food well, whereas L2 was not paired with food. In group Intermixed, L1 was followed by food on half of its presentations and preceded by food on the remainder, and L2 was not paired with food. Discrimination learning was assessed using two behavioral measures: visiting the food well during L1 and L2 (called goal-tracking; e.g., Boakes, 1977; Good & Honey, 1991) and interacting with L1 and L2 (called sign-tracking; e.g., Hearst & Jenkins, 1974; see also, Davey & Cleland, 1982). These two responses are depicted for the presentation of a single lever in Figure 1. L1 came to elicit more goal-tracking than L2 and this difference emerged every bit as readily in groups Forward and Intermixed. L1 also came to elicit more sign-tracking than L2, but this difference was more evident in group Forward than in group Intermixed. This example illustrates the fact that any answer to key questions about the conditions, content and mechanisms of learning will differ depending on the measure of learning that is used: The impact of the two training procedures, forward conditioning versus intermixed forward and backward conditioning, depends on which behavioral index is used (goal-tracking or sign-tracking). This indeterminacy is reinforced by the fact that the temporal profiles of these two conditioned behaviors across the temporal duration of the levers were quite different. In group Forward, goal-tracking increased across the duration of a 10-s lever, whereas sign-tracking declined; and in group Intermixed, the levels of goal-tracking were consistently high and those of sign-tracking were consistently low across L1 (Navarro et al., 2023; see also, Iliescu et al., 2020). Why should learning be more evident in one behavioral measure (e.g., goal-tracking) at one point during the CS and in a different measure (e.g., sign-tracking) at another point?

The questions raised by group-level dissociations in how learning is expressed proliferate once it is recognized that there are also individual differences in how learning is expressed. Pavlov (1927) observed marked quantitative and qualitative individual differences in the effects of conditioning in dogs. More recently, Patitucci et al. (2016; see also, Flagel, Akil & Robinson, 2009) showed that while some rats given forward conditioning trials with L1 and nonreinforced presentations of L2 exhibited learning primarily as differences in goal-tracking during L1 and L2, others exhibited it as differences in sign-tracking. Given the fact that all rats received the same experimental training, why should there be both quantitative and qualitative differences in how learning is expressed? These observations are quite beyond general-process models that make the simplifying assumption that there is an ordinal mapping between theoretical constructs (e.g., associative strength, V) and conditioned behavior (e.g., Rescorla & Wagner, 1972; see also, Mackintosh, 1975; Pearce & Hall, 1980).

The issues raised above have started to be addressed by an associative model of Pavlovian conditioning, HeiDI (Honey & Dwyer, 2022; Honey, Dwyer & Iliescu, 2020; Navarro et al., 2023). Indeed, the results presented by Navarro et al. (2023), involving the efficacy of backward conditioning procedures, provide key support for HeiDI. HeiDI assumes that forward conditioning trials (e.g., lever→food) result in the formation of reciprocal associations between the lever and food (i.e., a lever→food association and a food→lever association, or lever⇄food associations). This assumption means that the strength of both reciprocal associations can contribute to the generation of conditioned behaviors (see also, Asratyan, 1965), and it predicts that backward conditioning trials (e.g., food→lever) will engender conditioned responding when the lever is tested alone, based on food⇄lever associations. Experiment 1 sought to replicate and extend the results reported by Navarro et al. (2023) using a between-subjects design, by examining the impact of (1) forward (lever→food) and backward (food→lever) conditioning trials, and (2) the interval between the lever and food (1s or 11s), on the distribution of goal-tracking and sign-tracking across the duration of the lever. In Experiment 2, the same rats received higher-order conditioning trials in which an auditory stimulus (S1) was paired with the conditioned lever, and we examined the development of goal-tracking to S1 and its temporal distribution across S1. HeiDI provides a basis for the paradoxical prediction that arranging a trace interval between the lever and food, on forward conditioning trials, will increase higher-order conditioning to S1 (cf. Lin & Honey, 2011; Lin et al., 2013). The derivation of this prediction (formally described in Honey & Dwyer, 2022) will be presented in the context of Experiment 2, alongside alternative theoretical analyses.

Experiment 1

The design of Experiment 1 is depicted in Table 1. Rats received trials on which presentations of L1 (e.g., the lever to the left of the food well) were paired with food (L1→Food) and those of L2 (e.g., the lever to the right of the food well) were not (L2→No Food). We manipulated whether rats received forward conditioning (L1→Food) or backward conditioning (Food→L1) trials with L1, and whether the interval between the offset of L1 and the presentation of food (on forward trials), or the presentation of food and the onset of L1 (on backward trials) was 1s or 11s. This factorial design resulted in 4 groups: Forward-1s, Forward-11s, Backward-1s and Backward-11s. All groups received non-reinforced presentations of L1 in each session, providing an assessment of conditioned responding to the lever in the absence of responding generated by the presentation of food (which is a concern in the Backward-1s and Backward-11s groups). Trace conditioning typically impairs the development of a given conditioned behavior, with the impact of the trace interval depending on the conditioning preparation (for a review, see Mackintosh, 1974). In contrast, the development of excitatory Pavlovian conditioning generated by backward conditioning has often been cast as either evanescent (e.g., Heth, 1976) or unreliable (see Mackintosh, 1974, p.60). However, some evidence suggests that the efficacy of backward conditioning trials can be revealed using procedures that allow animals to express the fact that backward associations (like forward associations) encode information about the temporal order of the US and CS (e.g., Arcediano, Escobar & Miller, 2005; see also, Asratyan, 1965, p. 178-179). Moreover, Navarro et al. (2023) showed that even simple backward conditioning trials (in which the presentation of food preceded a lever by 1s) can produce sustained and marked evidence of excitatory conditioning as indexed via goal-tracking but not sign-tracking; while also noting that longer US-CS intervals produce inhibitory conditioning (Delamater, LoLordo, & Sosa, 2003; Taira et al., 2024). Notably, Navarro et al.'s procedure used a relatively short inter-trial interval (60s on average), which might have favoured the expression of goal-tracking (Thomas & Papini, 2020). Here, we used a longer inter-trial interval (130s on average) with the hope of generating more equivalent levels of the two forms of conditioned responding. This change should allow the baseline levels of goal-tracking (conditioned or not) to decline or extinguish (see also, Cinotti, Marchand, Roesch, Girard & Khamassi, 2019).

Table 1. Design of Experiments 1 and 2

Group	Experiment 1		Experiment 2
	Training trials	Probe trials	Probe trials
Forward-1s	L1→1s→Food L2→No Food	L1→No Food L2→No Food	S1→L1 S2→L2
Forward-11s	L1→11s→Food L2→No Food		
Backward-1s	Food→1s→L1 L2→No Food		
Backward-11s	Food→11s→L1 L2→No Food		

Note: L1 and L2 denote two levers and S1 and S2 denote two sounds (a clicker and white noise, counterbalanced); Food indicates the delivery of a food pellet and No food its absence.

Method

Subjects. Sixty-four naïve male Lister Hooded rats (mean ad lib weight = 342g; range: 309-373g; supplied by Envigo, UK) were randomly assigned to each of the 4 groups (*n*=16 per group). No formal power analysis was conducted because it was not possible to estimate the effect sizes for both first-order and higher-order conditioning based on immediately relevant prior research. Instead, the number of subjects was based on prior research that we have conducted using the autoshaping procedure (e.g., Navarro et al., 2023; Patitucci et al., 2016). Rats were housed in pairs in standard home cages. They were maintained between 85% and 95% of their ad lib weights by giving them restricted access to food at the end of each day in these cages, where they had continuous access to water. The room in which these home cages were placed had a 12-hr/12-hr light/dark cycle (lights on at 7 a.m.). The research was conducted in accordance with Home Office regulations under the Animal (Scientific Procedures) Act 1986. Research was conducted in accordance with the Home Office regulations under the Animal (Scientific Procedures) Act 1986 and under the authority of PPL number PP3468526 granted to D. M. Dwyer.

Apparatus. Sixteen identical conditioning boxes (30×24×21 cm: H×W×D; Med Associates, Georgia, VT) were used, with each box being placed in a sound-attenuating shell incorporating a ventilation fan that maintained the background noise at 68 dB(A). The boxes had two aluminum side walls, with front walls, back walls, and ceilings made from clear acrylic. The floor of each box was formed from 19 steel rods (4.8 mm diameter, 16 mm apart) placed above a stainless-steel tray. Food pellets (45 mg; LabDiet, St. Louis, MO, USA) were delivered to a food well (aperture: 5.3×5.3 cm), which was recessed in the center of the left wall at floor level. The food well was equipped with infrared detectors. Both onset and offset of interruption (e.g., by a rat's snout in the food well) of the detector was registered, allowing for duration of food well responses to be recorded. Two retractable levers (4.5×1.8×0.2 cm), located 3 cm to the left and right of the food well, were positioned at a height of 4.6 cm and 1.5 cm from the edge of the walls. As with the food well, we recorded a signal when the levers were depressed by 4mm from their horizontal resting position and again when they were released back to their resting position, which allowed the computation of the duration of each lever press. Each box was equipped with a speaker mounted behind the wall facing the levers/food well and 18 cm above the floor. In Experiment 2, these speakers delivered auditory stimuli (white noise and a 5 Hz clicker) at an intensity of 6 dB above background sound levels. MED-PC software controlled the insertion and retraction of the levers, delivery of food pellets and auditory stimuli, and recorded food well entries and lever presses. Presentations of the levers and auditory stimuli were synchronized to occur simultaneously across all boxes.

Procedure. Rats received one 46-min pre-training session in which 20 food pellets were delivered on a variable-time (VT) 130-s schedule (range: 110-150s). On each of the following 20 days, rats received a single training session that occurred at the same time of day for a given rat (with the earliest session starting at 09:00 AM). The first 10 trials in each session included 5 10-s presentations of L1 and L2, with the assignment of the left and right levers to L1 and L2 counterbalanced. For rats in groups Forward-1s and Forward-11s, a single food pellet was delivered 1s and 11s, respectively, after L1 was retracted from the chamber. For rats in group Backward-1s and Backward-11s, a single food pellet was delivered 1s and 11s, respectively, before L1 was extended into the chamber. For all rats, presentations of L2 were not paired with food. The final 12 trials in each session included 8 trials as described above (4 with L1 and 4 with L2), and 4 non-reinforced probe trials (2 with L1 and 2 with L2) in which no pellets were delivered (i.e., L1 like L2 was neither preceded nor followed by food). The trials were delivered on a variable-time (VT) 130-s schedule (range: 110–150s). The order in which the two levers were presented was random with the constraint that there were no more than two presentations of the same lever in succession.

Data analysis. All analyses were performed in R (R Development Core Team, 2021), using packages brms (Bürkner et al., 2022), and bayestestR (Makowski et al., 2022). To facilitate reproducibility, a singularity image with the R environment used for the analyses is available in the OSF link in the Transparency and openness section. MedPC files containing raw data were first processed in R to calculate the proportions of goal- and sign-tracking. We defined these proportions as the cumulative time in which a response was recorded (the total duration of all nose pokes or lever presses over a time window) divided by the length of time the response was measured (10s for trial-level analyses and 2s bins for stimulus-level analyses; see ahead). Proportions were used rather than response rates because they are less variable, especially for sign-tracking. Background levels of goal-tracking during the inter-trial interval and in the trace intervals between L1 and food were not assessed: Our primary interest was whether L1 and L2 differed (in the context of first-order conditioning in Experiment 1) and whether S1 and S2 different (in the context of higher-order conditioning in Experiment 2). Comparison of the proportions of responding during the stimuli (L1, L2, S1 and S2) and in their absence is complicated by differential stimulus support for the conditioned responses (i.e., goal-tracking and sign-tracking). For example, while the recent presentation of food would encourage goal-tracking in an immediately succeeding interval, the removal of L1 (or L2) would reduce any competition between sign-tracking and goal-tracking and increase goal-tracking. However, these levels of background and trace goal-tracking are available in the datasets available through the OSF link provided below. Response proportions were regressed using Gaussian Processes (GP), whose parameters were estimated under a Bayesian framework. Briefly, GPs model a joint distribution of data points and Gaussian functions, such that each data point is modelled as a linear combination of (indefinite, but data-supported) functions. The estimation of GPs is computationally expensive but has several advantages; crucially for the present work, GPs are a principled way to deal with non-monotonicities in the data, without specifying an explicit model (c.f., Navarro et al., 2023; see Schulz, Speekenbrink, & Krause, 2018, 2018, for an accessible introduction to GPs). In the present approach, we used GPs to jointly model goal- and sign-tracking proportions using a Dirichlet distribution adopting goal-tracking, sign-tracking, and alternative behaviors as mutually exclusive states, but whenever goal-tracking was the only response measured, or when the two types of response were differentially available across the trial (see Experiment 2), we modelled those probabilities against alternative behaviors using a Beta distribution. Each model was estimated via 8 chains of 8000 iterations each (2000 warmup iterations) and then subjected to convergence and posterior predictive checks (see annotated R code in the OSF repository). After the models were estimated, we performed inference using probability of direction (*pd*) tests on expected mean posterior differences (MPD). In this context, the pd describes the maximum proportion of posterior differences smaller or larger than zero (no difference). Notably, the *pd* test is inversely related to p-values, such that larger probabilities map onto smaller p-values (Makowski et al., 2019), with *pd*s = {.950, .975, .995, .9995} roughly mapping to *p*s = {.1, .05, .01, and .001, respectively}, though *pd*s denote the probability of an effect rather than the probability of observing the measured effect under a null distribution. MPDs provide non-standardized estimations of effect size, and we quantify the uncertainty surrounding these effect sizes via 95% Credible Intervals (CI) on each MPD, which are the Bayesian analogue of frequentist confidence intervals.

Transparency and openness. All data, scripts, models (and their posterior samples) used for the data analyses reported in this manuscript are available at: https://osf.io/ntkg3/.

Results

Figure 2 depicts the development of goal-tracking (left-hand panels, A-D) and sign-tracking (right-hand panels, E-H) across training on the non-reinforced probe trials with L1 and L2. Note that for Figure 2 and all remaining figures, individual points correspond to group means, whereas lines and shaded areas correspond to posterior means and credible intervals, respectively. Consider first the left-hand panels. Early in training, the overall levels of goal-tracking were higher in rats given forward training trials (groups Forward-1s and Forward-11s, A and B) than in those given backward training trials (groups Backward-1s and Backward-11s, C and D), MPD = 0.03, 95% CI [0.01, 0.05], pd = .998 in block 1. However, this difference was reversed by the end of training, with overall levels of goal-tracking being higher in backward groups than in forward groups, MPD = -0.04, 95% CI [-0.06, -0.02], pd > .999 in block 5. It is also apparent that while the overall levels of goal-tracking were higher during L1 than L2 in the forward groups (MPD = 0.02, 95% CI [0.01, 0.04], pd = .998), this was not the case in the backward groups, which showed no reliable differences in goal-tracking during L1 and L2 (MPD = 0.001, 95% CI [-0.01, 0.01], pd = .539). Sign-tracking levels were higher to L1 than to L2 in both forward groups (E and F; MPD = 0.06, 95% CI [0.04, 0.08], pd > .999 for group Forward-1s and MPD = 0.02, 95% CI [0.01, 0.04], pd = .997 for group Forward-11s) and the Backward-1s group (G, MPD = 0.02, 95% CI [0.01, 0.04], pd > .999). However, there were no reliable differences in sign-tracking to L1 and L2 in the Backward-11s group (H, MPD = 0.01, 95% CI [-0.004, 0.02], pd = .859). Notably, reliable differences in goal-tracking during L1 and L2 were only evident in group Forward-11s during blocks 1 and 2 (smallest MPD = 0.05, 95% CI [0.01, 0.10], pd = 0.992 in block 2). Conversely, reliable differences in sign-tracking to L1 and L2 were detected more frequently, especially later in training. Group Forward-1s showed reliable differences in blocks 2-5 (smallest MPD = 0.04, 95% CI [0.01, 0.07], pd = 0.992 in block 2), group Forward-11s did so in blocks 4 and 5 (smallest MPD =

0.05, 95% CI [0.02, 0.08], pd > 0.999 in block 4) and group Backward-1s did so in blocks 2-5 (smallest MPD = 0.02, 95% CI [0.003, 0.04], pd = 0.992 in block 2). There were no reliable differences in sign-tracking between L1 and L2 in Group Backward-11s.

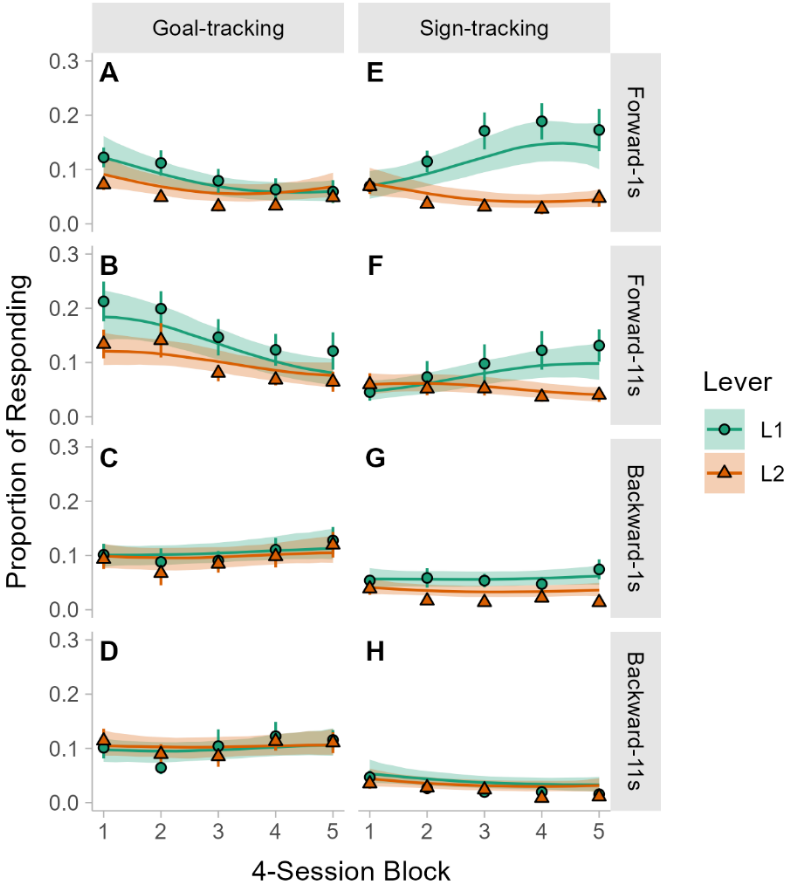

Figure 2. Experiment 1. First-order conditioning in four groups of rats distinguished by whether conditioning trials involved forward or backward pairings of L1 and food, and whether the interval between the L1 and food was 1s or 11s: Forward-1s and Forward-11s, Backward-1s and Backward-11s. In all groups, L2 was not paired with food. Panels A-D show the proportion of goal-tracking across 4-session blocks during non-reinforced probe trials with L1 and L2. Panels E-H show the corresponding proportions of sign-tracking. Results for each group are shown across rows: Forward-1s (A and E), Forward-11s (B and F), Backward-1s (C and H) and Backward-11s (D and H). Filled points denote group means and their whiskers, ± the standard error of the mean (SEM). The lines represent mean posterior predictions from our statistical model, and the shaded areas surrounding them denote their 95% credible interval.

To assess the distribution of responses across the levers, we divided the data from the final 4 sessions of first-order conditioning (block 5 in Figure 2) into successive 2-s bins relative to the onset of the levers. On this block, the levels of goal-tracking and sign-tracking during L1 and L2 appeared to have reached a point of relative stability. Figure 3 shows those distributions during the L1 and L2 trials. In line with performance towards the end of training, neither of the forward groups showed reliably greater goal-tracking to L1 than to L2 at any bin (A, largest MPD = 0.01, 95% CI [-0.02, 0.05], pd = .764 and MPD = 0.01, 95% CI [-0.01, 0.05], pd = .806 for groups Forward-1s and Forward-11s, both at 10s).[1] This pattern of goal-tracking was also true for both backward groups (C-D, largest MPD = 0.02, 95% CI [-0.05, 0.09], pd = .713 for group Backward-1s at 2s and MPD = 0.02, 95% CI [-0.02, 0.06], pd = .835 for group Backward-11s at 10s). Conversely, sign-tracking disclosed reliable differences in nearly all groups. Both forward groups showed reliable differences between L1 and L2 from 2s to 10s of lever presentation (E-F, smallest MPD = 0.04, 95% CI [0.01, 0.08], pd = .994 for group Forward-1s at 2s and MPD = 0.05, 95% CI [0.004, 0.09], pd = .984 for group Forward-11s at 10s). Notably, the Backward-1s group showed increased sign-tracking to L1 starting at 4s (G, smallest MPD = 0.02, 95% CI [0.001, 0.05], pd = .984 at 4s), but the Backward-11s group showed no reliable differences (H, largest MPD = 0.003, 95% CI [-0.02, 0.02], pd = .640 at 10s).

---

[1] The mean levels of goal-tracking in group Forward-11s during the late bins of L1 shown in Figure 3B do not represent accurately the performance of the group: The majority of rats in this group showed little to no goal-tracking during these later bins. Our statistical model correctly reflects this fact.

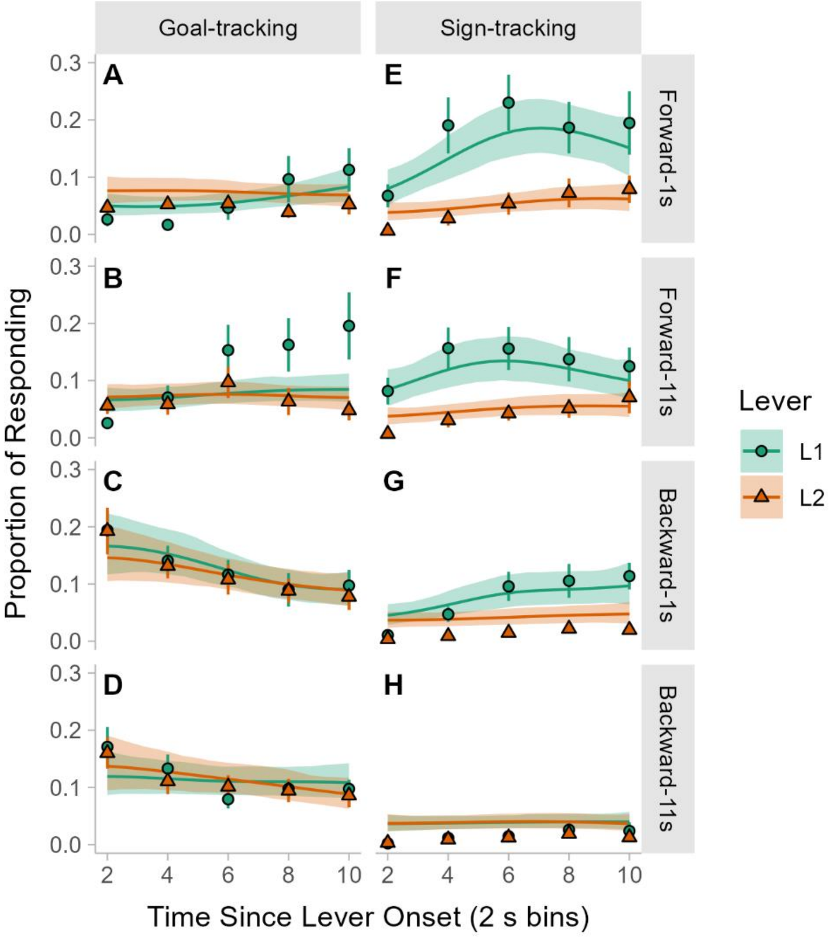

Figure 3. Experiment 1. The distribution of responding across 2-s presentations of L1 and L2 during first-order conditioning trials in four groups of rats distinguished by whether conditioning trials involved forward or backward pairings of L1 and food, and whether the interval between the L1 and food was 1s or 11s: Forward-1s and Forward-11s, Backward-1s and Backward-11s. In all groups, L2 was not paired with food. Each panel shows the proportion of responding across 2s bins, during non-reinforced probe trials given on the last 4 blocks of conditioning. Results for each group are shown across rows: Forward-1s (A and E), Forward-11s (B and F), Backward-1s (C and H) and Backward-11s (D and H). Filled points denote group means and their whiskers, ±SEM. The lines represent mean posterior predictions from our statistical model, and the shaded areas surrounding them denote their 95% credible interval.

Discussion

Experiment 1 examined first-order conditioning during a discrimination task in which one lever (L1) was paired with food and a second lever (L2) was not. Overall, increasing the intertrial interval (relative to that used routinely in our previous research) seemed to increase the levels of sign-tracking relative to those that we observe typically (e.g., Navarro et al., 2023; Patitucci et al., 2016), replicating the findings of Thomas and Papini (2020). For rats given forward conditioning trials (i.e., L1→Food), differential responding to L1 and L2 was rarely evident in goal-tracking but strongly evident in sign-tracking. Moreover, by the end of training the difference in sign-tracking between L1 and L2 was more evident in group Forward-1s than in group Forward-11s. That is, the introduction of an 11s trace interval (in group Forward-11s) impaired the development of differential responding to L1 and L2 (cf. Honey & Hall, 1991; Pavlov, 1927). For rats given backward conditioning trials (Food→L1), differential responding to L1 and L2 was restricted to group Backward-1s and the sign-tracking measure. While this effect was not large, it was sustained across many sessions of training. In Experiment 2, we examined the ability of L1 and L2 to support higher-order conditioning as indexed by goal-tracking to two auditory stimuli, S1 and S2.

Experiment 2

Experiment 2 assessed the extent to which L1 (and L2) supported the development of higher-order conditioning to auditory stimuli S1 (and S2): S1 (e.g., 10-s of white noise) was paired with L1 and S2 (e.g., a 10-s clicker) was paired with L2 (see Table 1). Goal-tracking during S1 and S2 was used to assess higher-order conditioning because it can be measured automatically in the same way as goal-tracking to L1 (and L2); we did not attempt to measure sign-tracking to S1 and S2. While first-order conditioning to a CS is often impaired when there is a long trace interval between that CS and food (see Mackintosh, 1974), a CS established using a trace conditioning procedure can support greater higher-order conditioning than a CS established without such a trace interval (Lin & Honey, 2011; see also, Lin, Dumigan, Dwyer, Good & Honey, 2013). This intriguing observation is beyond standard accounts of higher-order conditioning. It has, however, been addressed in a recent formal model of higher-order conditioning based on HeiDI (Honey & Dwyer, 2022; Navarro et al., 2023).

First, assume that higher-order conditioning reflects the capacity of S1 to evoke the memory of L1, which was itself linked to food during first-order conditioning. Now assume that the L1-food association includes the intensity of L1 at the point when food was delivered. The intensity with which L1 enters association will be lower when there is a trace between L1 and food than when there is no trace interval. As a result of initial pairings of S1 with L1, S1 will retrieve a memory of L1 with a low intensity. This low-intensity L1 representation will be more like the intensity of the L1 conditioned with an interval than to the intensity of the L1 conditioned with no such interval and should more readily activate the memory of food. On this basis, trace conditioning with L1 might be expected to generate more higher-order conditioning than conditioning with no trace interval (or a shorter interval). Honey and Dwyer (2022) used formal simulations, including a requisite similarity function, to demonstrate the range of conditions under which this prediction holds, which is constrained by the strength of the association between S1 and L1. The comparison of groups Forward-1s and Forward-11 allows the generality of the effect demonstrated by Lin and Honey (2011) to be investigated in an autoshaping procedure, which was an important impetus for the development of HeiDI. It is worth noting that the predicted results for Experiment 2 (and those reported by Lin & Honey, 2011) are inconsistent with the temporal coding hypothesis. According to this hypothesis, associations between stimuli encode their temporal order; and different temporally coded associations (e.g., L1→1s→Food and S1→L1) can be integrated through a common referent, in this case the referent would be L1. The integration of the temporally coded associations in

Group Forward-1s, L1→1s→Food and S1→L1, should result in S1 and Food being relatively coincident when S1 is presented (i.e., S1→11s→Food; 11s because L1 is 10s and there is a 1s gap between the offset of L1 and food). In contrast, in Group Forward-11s the integration of S1→L1 and L1→11s→Food should result S1 and Food being less coincident (i.e., S1→21s→Food). The temporal coding hypothesis predicts that trace conditioning should disrupt higher-order conditioning.

The temporal coding hypothesis provides a clear prediction about our backward conditioning procedures. In Group Backward-1s, the Food→1s→L1 association can be integrated with a S1→L1 association through the referent, L1, so that when S1 is presented it aligns with food. Under these conditions, the hypothesis predicts that S1 would generate conditioned behavior. In contrast, for Group Backward-11s integrating the Food→11s→L1 association with a S1→L1 association would result in S1 being aligned with the trace period and would not be predicted to generate conditioned behavior. There is, however, an alternative analysis for pattern of results predicted by the temporal coding hypothesis: S1 could evoke a memory of L1 as the result of the S1→L1 association, and to the extent that backward conditioning results in the formation of reciprocal Food⇄L1 then S1 should generate more conditioned behavior in Group Backward-1s than Backward-11s (Honey et al., 2020; Honey & Dwyer, 2022; Navarro et al., 2023). Finally, to the extent that backward conditioning with a 11-s trace interval in Group Backward-11s results in the development of inhibition to L1 (Delamater et al., 2003; Taira et al., 2024), then one might expect S1 to gain higher-order conditioned inhibition as a result of being paired with L1 (see Rescorla, 1976); with the influences of such inhibition being evident as low levels of conditioned responding during S1.

Method

Subject, apparatus and procedure. The rats and apparatus were those described in Experiment 1. On the 12 days that followed Experiment 1, rats continued to receive conditioning trials with L1 and L2 during the first 10 trials (5 per lever, as described in Experiment 1), but higher-order conditioning trials were intermixed during the last 12 trials. On higher-order conditioning trials, one of the auditory stimuli, S1 (white noise or clicker, counterbalanced) preceded the presentation of L1, and the remaining auditory stimulus, S2, preceded L2; no food pellets was delivered during these trials. The last 12 trials within each session were subdivided into 3, 4-trial blocks, each containing one conditioning trial with L1 (varying across groups in the same way as during first-order conditioning), a non-reinforced presentation of L2, and non-reinforced trials on which the offset of S1 was followed by L1 and the offset of S2 was followed by L2 (i.e., S1→L1 and S2→L2 trials). All other experimental parameters were identical to those described for first-order conditioning in Experiment 1.

Results

Figure 4 shows the proportions of goal-tracking to S1 and S2 across 2-session blocks of higher-order conditioning, when S1 preceded L1 and S2 preceded L2 (i.e., S1→L1 and S2→L2). Higher-order conditioning had markedly different effects in the forward groups. For group Forward-1s, goal-tracking during S1 was not reliably greater than during S2 on any block (A, highest pd = .937 in block 5; but pooling data over more sessions did reveal some significant differences, see ahead). In contrast, for group Forward-11s, goal-tracking to S1 was reliably higher than to S2 from block 2 onwards, being numerically largest and statistically most reliable in block 5 (B, MPD = 0.04, 95% CI [0.01, 0.08], pd = .997). The two backward groups showed similar levels of goal-tracking during S1 and S2 (C-D, highest pds = .830 in block 6 and .716 in block 5 for groups Backward-1s and Backward-11s, respectively).

.

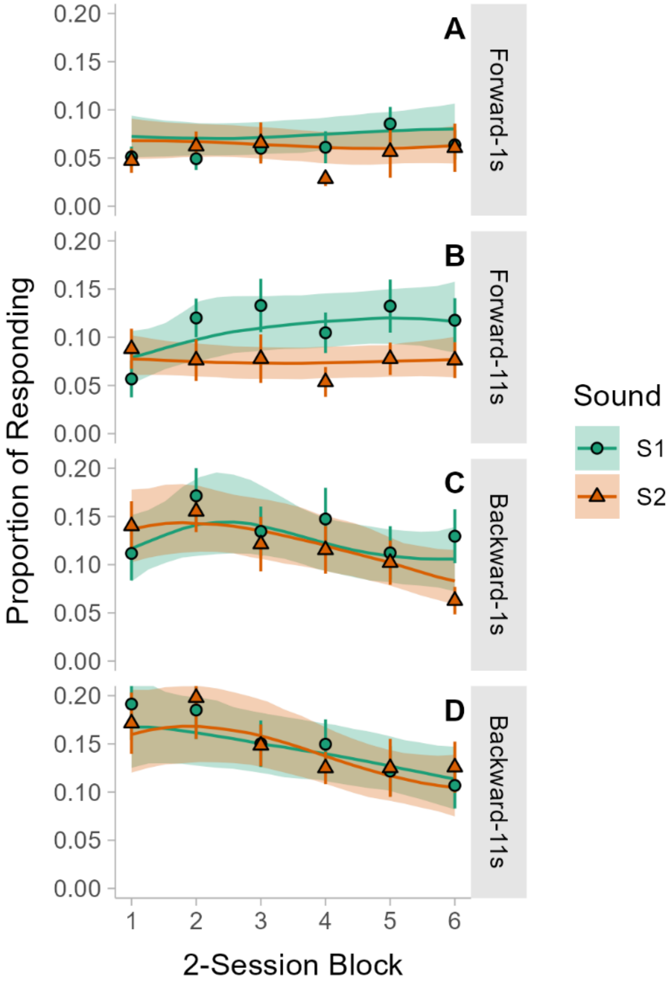

Figure 4. Experiment 2. Each panel shows the proportion of goal-tracking during S1 and S2 across 2-session blocks of non-reinforced higher-order conditioning trials on which S1 was paired with L1 (S1→L1) and S2 was paired with L2 (S2→L2). The four groups are distinguished by whether first-order conditioning trials involved forward or backward pairings of L1 and food, and whether the interval between the two was 1s or 11s: Forward-1s and Forward-11s (panels A and B), Backward-1s and Backward-11s (panels C and D). Filled points denote group means and their whiskers, ±SEM. The lines represent mean posterior predictions from our statistical model, and the shaded areas surrounding them denote their 95% credible interval.

Figure 5 depicts the distributions of responding during the components of the S1→L1 and S2→L2 trials. The data from each trial was divided into successive 2-s bins relative to the onset of the levers and pooled over the final 6 sessions of higher-order conditioning (i.e., blocks 4, 5 and 6 in Figure 4). On these blocks, the levels of goal-tracking and sign-tracking during S1, S2, L1 and L2 appeared to be relatively stable. As the two responses were differentially available across the trial, we assessed these data by estimating two separate models, one for goal-tracking and one for sign-tracking. Goal-tracking in group Forward-1s tended to be low during the initial bins of S1 and S2, but increased midway through S1 significantly surpassing goal-tracking to S2 (A-B, largest MPD = 0.03, 95% CI [0.01, 0.06], pd = .997 at -4s), and later decreased so there was no difference between S1 and S2 (MPD = 0.02, 95% CI [-0.005, 0.04], pd = .931 at 0s). Goal-tracking in Group Forward-11s tended to be similar throughout S1 (and significantly greater than throughout S2; smallest MPD = 0.03, 95% CI [0.01, 0.06], pd = .992 at 0s). After the onset of the lever stimuli, goal-tracking to L1 was suppressed, with group Forward-1s goal-tracking to L1 being significantly lower than to L2 at 6s (A, MPD = -0.02, 95% CI [-0.03, -0.0004], pd = .978, likely due to the high-levels of sign-tracking in this group during that period) and group Forward-11s failing to show any significant differences between L1 and L2 (B, largest MPD = -0.01, 95% CI [-0.03, 0.01], pd = 0.861 at 6s). Sign-tracking in both forward groups was significantly greater to L1 than to L2 throughout the lever duration (E-F, smallest MPD = 0.05, 95% CI [0.02, 0.09], pd > 0.999 for group Forward-1s, and MPD = 0.06, 95% CI [0.03, 0.09], pd > 0.999 for Forward-11s, both at 2s).

The two backward groups showed markedly different patterns of responding during these higher-order conditioning trials. Goal-tracking in group Backward-1s was reliably higher to S1 than to S2 at 0s, during the last 2s of the auditory stimuli (C, MPD = 0.04, 95% CI [0.01, 0.07], pd = 0.998) and reliably higher to L1 than to L2 during 2-4s (smallest MPD = 0.04, 95% CI [0.01, 0.07], pd = .993 at 4s). This is surprising, given the probe trials for L1 and L2 during

Experiment 1 revealed no such differences in goal-tracking. In contrast, group Backward-11s showed small but reliable differences during the last 2s of the auditory stimuli and the first 2s of the levers, with goal-tracking towards the end of S2 being reliably higher than S1 (D, smallest MPD = -0.03, 95% CI [-0.06, -0.001], pd = 0.979 at -2s) and goal-tracking during the beginning of L2 being reliably higher than the equivalent to L1 (MPD = -0.03, 95% CI [-0.05, -0.001], pd = 0.979 at 2s). Most notably, sign-tracking in both groups given backward conditioning increased and then decreased across L1, becoming reliably higher than sign-tracking to L2 at 8s of the lever presentation (G-H, MPD = 0.02, 95% CI [0.001, 0.05], pd = 0.979 for group Backward-1s and MPD = 0.03, 95% CI [0.01, 0.05], pd = 0.997 for group Backward-11s) but showing no reliable differences at10 s (MPD = 0.03, 95% CI [-0.002, 0.06], pd = 0.963 for group Backward-1s and MPD = 0.02, 95% CI [-0.0001, 0.05], pd = 0.974 for group Backward-11s).

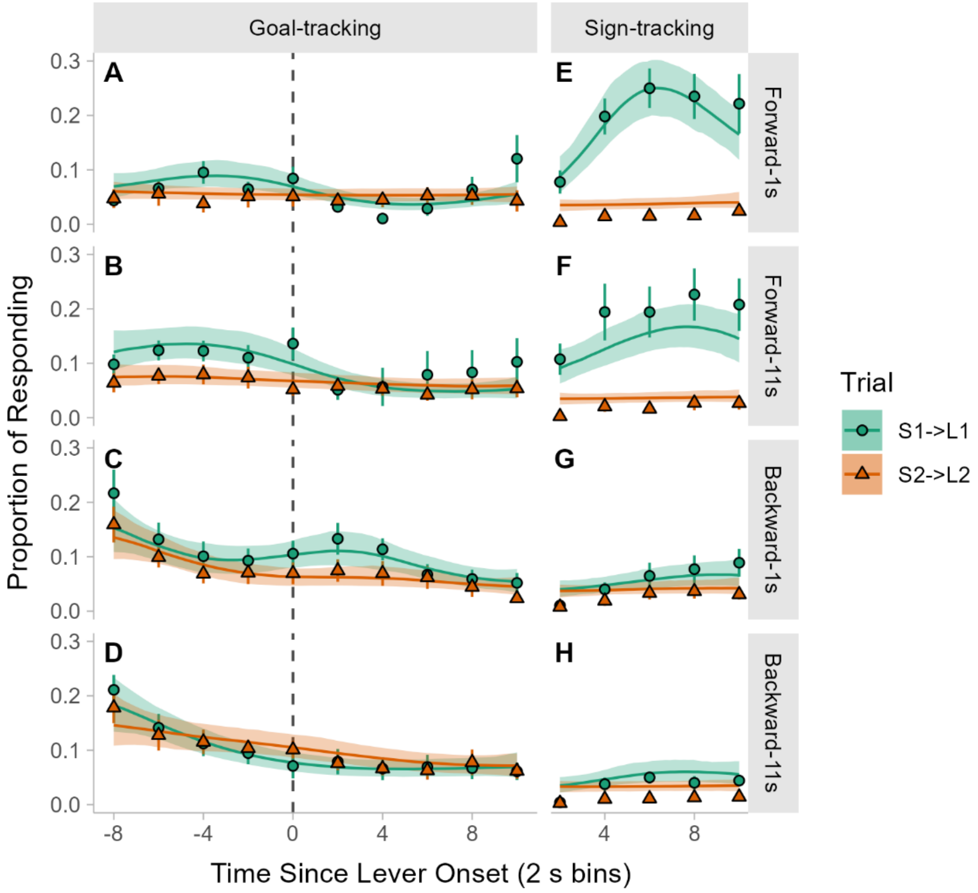

Figure 5. Experiment 2. The distribution of responding across the 10-s components of the non-reinforced higher-order conditioning trials in Experiment 2 (i.e., S1→L1 and S2→L2) in four groups of rats distinguished by whether conditioning trials involved forward or backward pairings of L1 and food, and whether the interval between the L1 and food was 1s or 11s: Forward-1s and Forward-11s, Backward-1s and Backward-11s. In all groups, L2 was not paired with food. Each panel shows the proportion of responding across 2s bins, on the last 6 sessions of higher-order conditioning. Panels A-D show the mean proportions of goal-tracking on S1→L1 and S2→L2 trials. Panels E-H show the corresponding proportions of sign-tracking during L1 and L2. The vertical dashed lines denote the offset of the 10s sounds (S1 and S2) and the onset of the levers (L1 and L2). Filled points denote group means and their whiskers, ±SEM. The lines represent mean posterior predictions from our statistical model, and the shaded areas surrounding them denote their 95% credible interval.

Discussion

The patterns of higher-order conditioning in groups given either forward or backward first-order conditioning raise several interesting issues, which will be considered in turn. Excitatory higher-order conditioning, as measured by more goal-tracking during S1 than S2, was evident in group Forward-11s, but not clearly seen in group Forward-1s (Figure 4, A-B). This observation is surprising given the fact that first-order conditioning with L1 was less evident, as indexed by sign-tracking, in group Forward-11s than in group Forward-1s (Figure 2, E-F). However, as already noted, these results have a precedent. Lin and Honey (2011; see also, Lin et al., 2013) showed that longer trace intervals generated more higher-order conditioning (second-order conditioning and sensory preconditioning) than shorter intervals. As outlined when introducing Experiment 2, HeiDI provides one potential analysis for this finding (see Honey & Dwyer, 2022). The higher-order conditioning in group Forward-11s has an interesting feature: Differences in goal-tracking to S1 and S2 were evident at a point where (differential) first-order conditioning to L1 and L2 was most apparent in sign-tracking. This fact suggests that higher-order conditioning reflected the operation of an associative chain (i.e., S1→L1→Food) or a mediated association between S1 and food (i.e., S1→Food) rather than a stimulus→response association (i.e., S1→interact with lever; cf. Holland, 1977; Rizley & Rescorla, 1972). These observations are consistent with studies of second-order conditioning in pigeons (e.g., Stanhope, 1992) and sensory preconditioning in rats (e.g., Dwyer, Burgess & Honey, 2012), which showed that the nature of higher-order conditioned responding can be quite different from that of first-order conditioning. These results are also consistent with HeiDI, which provides a formal analysis of higher-order conditioning, the conditions under which it occurs and how it is manifest in behavior (Honey et al., 2020; Honey & Dwyer, 2022; Navarro et al., 2023). We will return to HeiDI in the General Discussion.

Both groups given backward conditioning trials, Backward-1s and Backward-11s, showed similar levels of goal-tracking during S1 and S2 (Figure 4, C-D). However, by the later stages of training, group Backward-1s showed more goal-tracking during the last seconds of S1 than S2, and more goal-tracking during the initial bins of L1 than L2 (see Figure 5, C-D). These results provide evidence that backward conditioning, at least in the case of group Backward-1s, can produce sustained changes in both sign-tracking (Experiment 1) and goal-tracking (Experiment 2). These results join those reported by Navarro et al (2023) in showing that backward conditioning can result in sustained changes in excitatory conditioning (cf. Mackintosh, 1974, p. 60; see also, Heth, 1976). A final intriguing observation is that in group Backward-11s, S1 generated less goal-tracking than S2, and L1 generated less goal-tracking than L2. The lower levels of responding to S1 than S2 is consistent, at least, with the development of inhibitory higher-order conditioning to S1 (Rescorla, 1976), although additional controls would be necessary in order to confirm this interpretation. According to this analysis, it must be assumed that L1 had gained inhibitory properties, perhaps because it was non-reinforced in a context which had momentarily gained associative strength as the result of the presentation of food (see also, Delamater et al., 2003; Taira et al., 2024); and that S1 borrowed these properties through being paired with L1. Of course, according to some models, such borrowing of (inhibitory) associative strength would need to outweigh the fact that the direct link between S1 and food should have been strengthened on S1→L1 trials: Because S1 was non-reinforced in the presence of an inhibitor, L1, the pooled error term within the Rescorla and Wagner (1972) model would be positive, which should have resulted in S1 developing an excitatory association with food. Of course, this is the complement of the argument that is required to explain excitatory second-order conditioning within such a model: The excitatory associative chain (i.e., S1→L1→Food) would need to counteract the development of the inhibitory link between S1 and food on non-reinforced S1→L1 trials (for further discussion, see Honey & Dwyer, 2022). The fact that L1 resulted in more sign-tracking than L2 in group Backward-11s seems inconsistent with it having inhibitory properties. However, this difference in sign-tracking, which was not evident throughout training in Experiment 1, could be a secondary consequence of the difference in goal-tracking between S1 and S2 rather than evidence of excitatory conditioning. In contrast, for group Backward-1s there was consistent evidence of excitatory first-order conditioning from higher levels of both sign-tracking and goal-tracking during L1 than L2 in Experiment 2.

General Discussion

Pavlovian conditioning is the most widely recognized phenomenon in psychology. However, while some of its properties are well-established, they are often overlooked both descriptively and theoretically. For example, conditioning procedures generate behaviors that not only reflect the nature of the US but also the CS; and these behaviors are temporally distributed across the duration of the CS. Here, we examined the spatio-temporal characteristics of conditioned responding during first-order Pavlovian conditioning (Experiment 1) and higher-order conditioning (Experiment 2) using autoshaping procedures with rats: We measured the development of goal-tracking and sign-tracking across both sessions and the duration of lever presentations, and goal-tracking during the auditory stimuli.

One general conclusion that the results of Experiments 1 and 2 reinforce is that the impact of both group-level and within-subjects manipulations critically depends on the nature of the conditioned response that is measured (goal-tracking or sign-tracking) and when it is measured across both sessions and the duration of the stimuli (early or later). These dissociations represent a challenge to models of Pavlovian conditioning that assume a monotonic mapping between learning and performance, irrespective of whether they are trial-based (e.g., Mackintosh, 1975; Pearce, 1994; Pearce & Hall, 1980; Rescorla & Wagner, 1972) or time-based (e.g., Gallistel & Gibbon, 2000; Stout & Miller, 2007; Wagner, 1981). Some of these dissociations are highlighted in the next paragraph.

Consider first Experiment 1. Rats received trials on which one lever was paired with food (L1→Food) and another lever was not paired with food (L2→No Food). Across Pavlovian the later conditioning sessions, differences in responding between L1 and L2 were more evident in sign-tracking than goal-tracking (i.e., Forward-1s > Forward-11s > Backward-1s > Backward-11s). The observation that backward like forward conditioning can generate *sustained* excitatory conditioning is interesting in its own right (cf. Heth, 1976, Mackintosh, 1974, p. 60) and the finding that such conditioning is only evident (here) in sign-tracking does not simply reflect a general lack of sensitivity of goal-tracking as a measure. Most notably, Navarro et al. (2023) used procedures that tended to result in more goal-tracking than sign-tracking to levers (cf. Thomas & Papini, 2020) and observed sustained evidence of both forward and backward conditioning in goal-tracking, but not sign-tracking (see also, Iliescu et al., 2018; Patitucci et al. 2016). We turn now to Experiment 2. Rats from Experiment 1 received higher-order conditioning where one auditory stimulus, S1, was paired with L1 (i.e., S1→L1), and S2 was paired with L2 (i.e., S2→L2). Excitatory higher-order conditioning was evident in differential goal-tracking during S1 and S2, with this difference being most clear in group Forward-11s (cf. Lin & Honey, 2013; Lin et al., 2013). Evidence consistent with inhibitory higher-order conditioning was evident exclusively in group Backward-11s. Moreover, while goal-tracking declined across S1 (and S2) in the backward groups, it did not do so in the forward groups; similarly, goal-tracking during L1 declined in backward groups and increased in forward groups (see also, Navarro et al., 2023).

How might models of Pavlovian conditioning accommodate such dissociations? We have developed one approach that seems promising. HeiDI assumes that bidirectional associations form during pairings of one CS (X) and US (i.e., X⇄food), and that pairing a second CS (A) with the first allows A to enter into a direct inhibitory association with the US and to access the first-order bi-directional associations (i.e., X⇄food; see Honey et al., 2020; Honey & Dwyer, 2022). To explain differences in the spatial distribution of responding, HeiDI assumes that the relative perceived intensities of the CS and (retrieved) US determines whether CS-oriented responding (e.g., sign-tracking) or US-oriented responding (e.g., goal-tracking) will dominate; with the balance shifting from goal-tracking to sign-tracking as the perceived intensity of the CS increases relative to the US. To explain differences in the temporal distribution of responding, HeiDI assumes that the perceived intensity of the CS declines across its duration and that the perceived intensity at the point when the US arrives becomes associated with that US. One consequence of these assumptions is that associative strength will accrue to a lower perceived intensity during forward than backward conditioning (and contribute to their distinct profiles of responding); with other perceived intensities of the same CS being capable of generating responding to the extent that they are similar (see Navarro et al., 2023). We have provided formal simulations illustrating the general utility of this approach across a broad range of phenomena involving both first-order conditioning and higher-order conditioning -

phenomena that were difficult for alternative analyses to address[2]. These phenomena, which have counterparts in Experiments 1 and 2, included the facts that (1) backward conditioning procedures can support sustained excitatory conditioning, which has a different temporal profile to forward conditioning (Navarro et al., 2023); and (2) excitatory higher-order conditioning can be more effective after first-order trace conditioning (Lin & Honey, 2013; Lin et al., 2013). For example, in group Backward-1s there was a marked decline in goal-tracking across the temporal extents of both L1 and L2 (Figure 3C, see also, Navarro et al., 2023), which was complemented by a modest but selective increase in sign-tracking across L1 (Figure 3G). Whereas in group Forward-1s, goal-tracking remained low throughout L1 and L2 (Figure 3A) and there was a marked increase in sign-tracking across L1 (Figure 3E). Also, first-order forward conditioning (assessed using sign-tracking) was more evident in group Forward-1s than group Forward-11s (Figures 2E and 2F, respectively), but higher-order conditioning (as assessed using goal-tracking) was more evident in group Forward-11s than group Forward-1s (Figures 4B and 4A, respectively).

We will close this discussion by considering one feature of the HeiDI model that we have given little formal or indeed informal consideration to in the past: The role of the inter-trial interval in learning and performance. In keeping with some others, we have assumed that such intervals provide (further) opportunities for learning about the constellation of features collectively referred to as the experimental context. The perceived intensity of these features might decline (i.e., habituate) across an experimental session, gain and lose associative strength, and generate a range of different behaviors (e.g., exploration, orienting, sniffing).

---

[2] Navarro (2023) has developed an open-access R package (calmr) for generating simulations of HeiDI and other models of Pavlovian conditioning (https://cran.r-project.org/web/packages/calmr/index.html).

However, renewed consideration of the role of the inter-trial interval is prompted by the observation that in contrast to our previous autoshaping research (with short inter-trial intervals, ITIs), here we observe marked levels of sign-tracking relative to goal-tracking (with longer ITIs; see also, Thomas & Papini, 2020). There are several potential explanations for this observation (e.g., within HeiDI). For example, longer inter-trial intervals will allow the extinction of any direct (and indirect) associations involving the context (e.g., involving both CSs and USs) and produce changes in the (sign- and goal-tracking) behaviors generated by the context. However, perhaps the most obvious consequence of lengthening the inter-trial interval is to reduce the impact of short-term adaptation or habituation on the perceived intensity of the lever and thereby behaviour directed toward the lever (i.e., sign-tracking; see Honey et al., 2020). Evaluating this possibility experimentally and presenting formal modelling is beyond the scope of this paper (but see, Iliescu et al., 2020; Navarro et al., 2023). Nevertheless, being able to shift the distributions of goal-tracking and sign-tracking has the desirable consequence of allowing theoretically important findings (e.g., features of backward conditioning and higher-order conditioning) to be observed with both response measures.

Author note

All authors contributed to the ideas presented in this article, the design and analysis of the experiments, and to writing the paper. This research was undertaken using the supercomputing facilities at Cardiff University operated by Advanced Research Computing at Cardiff (ARCCA) on behalf of the Cardiff Supercomputing Facility and the HPC Wales and Supercomputing Wales (SCW) projects. We acknowledge the support of the latter, which is part-funded by the European Regional Development Fund (ERDF) via the Welsh Government. The research was supported by the BBSRC (UK; BB/T004339/1; PI: RCH). VN was partially supported by the UKRI under the UK government's Horizon Europe funding guarantee (EP/Y026489/1). Address for correspondence: Robert C. Honey, School of Psychology, Cardiff University, Tower Building, Park Place, Cardiff, CF10 3AT, UK; Tel: +44 (0)29 20875868; Email: Honey@cardiff.ac.uk.

References

Arcediano, F., Escobar, M., & Miller, R.R. (2005). Bidirectional associations in humans and rats. *Journal of Experimental Psychology: Animal Behavior Processes*, *31*, 301-318. https://doi.org/10.1037/0097-7403.31.3.301

Asratyan, E.A. (1965). *Compensatory adaptations, reflex activity and the brain* (pp. xv, 194). Pergamon Press.

Boakes, R.A. (1977). Performance on learning to associate a stimulus with positive reinforcement. In H. Davis, & H.M.B. Hurwitz (Eds.). *Operant-Pavlovian interactions* (pp. 67-97). NJ: Lawrence Erlbaum Associates.

Bürkner, P.-C., Gabry, J., Weber, S., Johnson, A., Modrak, M., Badr, H. S., Weber, F., Ben-Shachar, M. S., & Rabel, H. (2022). brms: Bayesian regression models using "Stan" (2.17.0) [Computer software]. https://CRAN.R-project.org/package=brms

Cinotti, F., Marchand, A.R., Roesch, M.R., Girard, B., & Khamassi, M. (2019). Impacts of inter-trial interval duration on a computational model of sign-tracking vs goal-tracking behaviour. *Psychopharmacology*, *236*, 2373-2388. https://doi.org/10.1007/s00213-019-05323-y

Davey, G.C.L., & Cleland, G.G. (1982). Topography of signal-centered behavior in the rat: Effects of deprivation state and reinforcer type. *Journal of the Experimental Analysis of Behavior*, *38*, 291-304. https://doi.org/10.1901/jeab.1982.38-291

Delamater, A.R., LoLordo, V.M., & Sosa, W. (2003). Outcome-specific conditioned inhibition in Pavlovian backward conditioning. *Learning & Behavior*, *31*, 393-402. https://doi.org/10.3758/BF03196000

Dickinson, A. (1980). *Contemporary Animal Learning Theory*. Cambridge University Press.

Dickinson, A., Hall, G., & Mackintosh, N.J. (1976). Surprise and the attenuation of blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, *2*, 313-322. https://doi.org/10.1037/0097-7403.2.4.313

Dwyer, D.M., Burgess, K.V., & Honey, R.C. (2012). Avoidance but not aversion following sensory-preconditioning with flavors: A challenge to stimulus substitution. *Journal of Experimental Psychology*: *Animal Behavior Processes*, *38,* 359-368. https://doi.org/10.1037/a0029784

Flagel, S.B., Akil, H., & Robinson, T.E. (2009). Individual differences in the attribution of incentive salience to reward-related cues: Implications for addiction. *Neuropharmacology*, *56*, 139-148. https://doi.org/10.1016/j.neuropharm.2008.06.027

Gallistel, C. R., & Gibbon, J. (2000) Time, rate, and conditioning. *Psychological Review 107,* 289–344. https://doi.org/10.1037/0033-295x.107.2.289.

Good, M., & Honey, R.C. (1991). Conditioning and contextual retrieval in hippocampal rats. *Behavioral Neuroscience*, *105*, 499-509. https://doi.org/10.1037/0735-7044.105.4.499

Hall, G., & Pearce, J.M. (1979). Latent inhibition of a CS during CS–US pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*, 31-42. https://doi.org/10.1037/0097-7403.5.1.31

Hearst, E., & Jenkins, H.M. (1974). Sign tracking: The stimulus-reinforcer relation and directed action. Austin: TX: Psychonomic Society.

Heth, C.D. (1976). Simultaneous and backward fear conditioning as a function of number of CS-UCS pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, *2*, 117-129. https://doi.org/10.1037/0097-7403.2.2.117

Holland, P.C. (1977). Conditioned stimulus as a determinant of the form of the Pavlovian conditioned response. *Journal of Experimental Psychology: Animal Behavior Processes*, *3*, 77-104. https://doi.org/10.1037//0097-7403.3.1.77

Holland, P.C. (1984). Origins of behavior in Pavlovian conditioning. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 18, pp. 129-174). Academic Press. https://doi.org/10.1016/S0079-7421(08)60361-8

Honey, R.C., & Dwyer, D.M. (2022). Higher-order conditioning: A critical review and computational model. *Psychological Review*, *129*, 1338-1357. https://doi.org/10.1037/rev0000368

Honey, R.C., Dwyer, D.M., & Iliescu, A.F. (2020). HeiDI: A model for Pavlovian learning and performance with reciprocal associations. *Psychological Review*, *127*, 829-852. https://doi.org/10.1037/rev0000196

Iliescu, A.F., Dwyer, D.M., & Honey, R.C. (2020). Individual differences in the nature of conditioned behavior across a conditioned stimulus: Adaptation and application of a model. *Journal of Experimental Psychology. Animal Learning and Cognition*, *46*, 460-469. https://doi.org/10.1037/xan0000270

Iliescu, A.F., Hall, J., Wilkinson, L.S., Dwyer, D.M., & Honey, R.C. (2018). The nature of phenotypic variation in Pavlovian conditioning. *Journal of Experimental Psychology. Animal Learning and Cognition*, *44*, 358-369. https://doi.org/10.1037/xan0000177

Kasprow, W.J., Schachtman, T.R., & Miller, R.R. (1987). The comparator hypothesis of conditioned response generation: Manifest conditioned excitation and inhibition as a function of relative excitatory strengths of CS and conditioning context at the time of testing. *Journal of Experimental Psychology: Animal Behavior Processes*, *13*, 395-406. https://doi.org/10.1037/0097-7403.13.4.395

Leyland, C.M., & Mackintosh, N.J. (1978). Blocking of first- and second-order autoshaping in pigeons. *Animal Learning & Behavior*, *6*, 391-394. https://doi.org/10.3758/BF03209633

Lin, T.E., Dumigan, N.M., Dwyer, D.M., Good, M.A., & Honey, R.C. (2013). Assessing the encoding specificity of associations with sensory preconditioning procedures. *Journal of Experimental Psychology: Animal Behavior Processes*, *39*, 67-75. https://doi.org/10.1037/a0030662

Lin, T.E., & Honey, R.C. (2011). Encoding specific associative memory: Evidence from behavioral and neural manipulations. *Journal of Experimental Psychology: Animal Behavior Processes*, *37*, 317-329. https://doi.org/10.1037/a0022497

Mackintosh, N.J. (1974). *The psychology of animal learning*. London: Academic Press.

Mackintosh, N.J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276-298. https://doi.org/10.1037/h0076778

Makowski, D., Lüdecke, D., Ben-Shachar, M. S., Patil, I., Wilson, M. D., Wiernik, B. M., Bürkner, P.-C., Mahr, T., Singmann, H., Gronau, Q.F., & Crawley, S. (2022). bayestestR: Understand and describe Bayesian models and posterior distributions (0.12.1) [Computer software]. https://CRAN.R-project.org/package=bayestestR

Navarro, V.M., Dwyer, D. M., & Honey, R.C. (2023). Prediction error in models of adaptive behavior. *Current Biology*, *33*, 4238-4243. https://doi.org/10.1016/j.cub.2023.08.043

Patitucci, E., Nelson, A.J.D., Dwyer, D.M., & Honey, R.C. (2016). The origins of individual differences in how learning is expressed in rats: A general-process perspective. *Journal of Experimental Psychology: Animal Learning and Cognition*, *42*, 313-324. https://doi.org/10.1037/xan0000116

Pavlov, I.P. (1928). *Lectures on conditioned reflexes: Twenty-five years of objective study of the higher nervous activity (behaviour) of animals*. NY: International Publishers Co., Inc..

Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review, 101*, 587–607.

Pearce, J.M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*, 532-552. https://doi.org/10.1037/0033-295X.87.6.532

R Development Core Team. (2021). R: A language and environment for statistical computing. (3.3.2) [Computer software]. R Foundation for Statistical Computing. http://www.R-project.org

Rescorla, R.A. (1976). Second-order conditioning of Pavlovian conditioned inhibition. *Learning and Motivation*, *7*, 161-172.

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. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory*. Appleton-Century-Crofts.

Rescorla, R.A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66*, 1-5. https://doi.org/10.1037/h0025984

Rescorla, R.A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, *43*, 151-160.

Rizley, R.C., & Rescorla, R.A. (1972). Associations in second-order conditioning and sensory preconditioning. *Journal of Comparative and Physiological Psychology*, *81*, 1-11. https://doi.org/10.1037/h0033333

Schulz, E., Speekenbrink, M., & Krause, A. (2018). A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions. *Journal of Mathematical Psychology*, *85*, 1-16. https://doi.org/10.1016/j.jmp.2018.03.001

Stanhope, K.J. (1992). The representation of the reinforcer and the force of the pigeon's keypeck in first- and second-order conditioning. *Quarterly Journal of Experimental Psychology*, *44B*, 137-158. https://doi.org/10.1080/02724999208250607

Stout, S.C., & Miller, R.R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review*, *114*, 759-783. https://doi.org/10.1037/0033-295X.114.3.759

Swan, J.A., & Pearce, J.M. (1987). The Influence of Predictive Accuracy on Serial Autoshaping: Evidence of Orienting Responses. *Journal of Experimental Psychology: Animal Behavior Processes*, *13*, 407-417. https://doi.org/10.1037/0097-7403.13.4.407

Taira, M., Millard, S.J, Verghese, A., DiFazio, L.E., Hoang, I.B., Jia, R., Sias, A., Wikenheiser, A., & Sharpe, M.J. (2024). Dopamine release in the nucleus accumbens core encodes the general excitatory components of learning. *Journal of Neuroscience*, *44*, e0120242024.

Thomas, B.L., & Papini, M.R. (2020). Shifts in intertrial interval duration in autoshaping with rats: Implications for path dependence. *Learning and Motivation*, *72*, 101687.

Timberlake, W., & Grant, D.L. (1975). Auto-shaping in rats to the presentation of another rat predicting food. *Science*, *190*, 690-692. https://doi.org/10.1126/science.190.4215.690

Wagner, A.R. (1981). SOP: A model of automatic memory processing in animal behavior. In N.E. Spear & R.R. Miller, *Information processing in animals: Memory mechanisms* (pp. 5-48). Hillsdale, NJ: Erlbaum. https://doi.org/10.4324/9781315798820-6

Zener, K. (1937). The significance of behavior accompanying conditioned salivary secretion for theories of the conditioned response. *American Journal of Psychology*, *50*, 384-403. https://doi.org/10.2307/1416644