Original Article

Corresponding Author:

Arthur Prével, Laboratoire SCALab, UMR 9193, Université de Lille Nord de France, Campus de Lille 3, Domaine Universitaire du Pont de Bois, B.P. 60149

59653 Villeneuve D'Ascq Cedex France

Email: arthur.prevel@univ-lille3.fr

Excitatory Second-Order Conditioning Using a Backward First-Order Conditioned Stimulus: A Challenge for Prediction Error Reduction

Arthur Prével[1], Vinca Rivière[1], Jean-Claude Darcheville[1],

Gonzalo P. Urcelay[2], & Ralph R. Miller[3]

[1]Laboratoire SCALab, UMR 9193, Université de Lille Nord de France Campus de Lille 3, Domaine Universitaire du Pont de Bois, B.P. 60149 59653 Villeneuve d'Ascq Cedex France

[2]Department of Neuroscience, Psychology, and Behaviour, University of Leicester, Leicester, LE1 7HA, UK

[3]Department of Psychology, State University of New York at Binghamton, Binghamton, NY 13902-6000 USA

Abstract

Prével and colleagues (2016) reported excitatory learning with a backward CS in a conditioned reinforcement preparation. Their results add to existing evidence of backward CSs sometimes being excitatory, and were viewed as challenging the view that learning is driven by prediction error reduction, which assumes that only predictive (i.e., forward) relationships are learned. The results instead were consistent with the assumptions of both Miller's Temporal Coding Hypothesis and Wagner's Sometimes Opponent Processes (SOP) model. The present experiment extended the conditioned reinforcement preparation developed by Prével et al. (2016) to a backward second-order conditioning preparation, with the aim of discriminating between these two accounts. We tested whether a second-order CS can serve as an effective conditioned reinforcer, even when the first-order CS with which it was paired is a backward conditioned stimulus that elicits no responding. Evidence of conditioned reinforcement was found, despite no CR being elicited by the first-order backward CS. The evidence of second-order conditioning in the absence of excitatory conditioning to the first-order CS is interpreted as a challenge to SOP. In contrast, the present results are consistent with the Temporal Coding Hypothesis and constitute a conceptual replication in humans of previous reports of excitatory second-order conditioning in rodents with a backward CS. The proposal is made that learning is driven by 'discrepancy' with prior experience as opposed to '*prediction* error.'

Keywords

Conditioned reinforcement; prediction error reduction; second-order conditioning; SOP model; Temporal Coding Hypothesis.

Excitatory Second-Order Conditioning

Using a Backward First-Order Conditioned Stimulus:

A Challenge for Prediction Error Reduction

Toward better understanding the necessity of prediction error in learning, we recently reported evidence of excitatory conditioning by humans with a backward conditioned stimulus (CS) in a conditioned reinforcement preparation (Prével, Rivière, Darcheville, & Urcelay, 2016). Backward conditioning is of theoretical importance in that it is seemingly an instance of error-correction learning in which the error is not ‘predictive,’ which is contrary to the widely-held view that ‘prediction’ error is essential for learning to occur. In Prével et al., the conditioning phase was implemented in a computer-based operant task in which the participants had to complete a fixed-ratio 12 (FR12) schedule to obtain presentations of the primary reinforcer and subsequently the CSs. The conditioning was carried out with two CSs, one CS paired with a specific magnitude of the primary reinforcer (unconditioned stimulus [US]) and the second CS with no reinforcement (no-US). Conditioned reinforcement was demonstrated in a choice task (i.e., a concurrent fixed-ratio 12 [FR12] schedule) in which the participants allocated their responses to a response key that delivered the conditioned reinforcer in the absence of the primary reinforcer. The effect was observed in comparison with a number of different control conditions that ruled-out the possibility of forward trace conditioning as suggested by Spetch, Wilkie, and Pinel (1981).

Prével, et al.’s (2016) results replicated with human participants a previous experiment by Urushihara (2004) that demonstrated conditioned reinforcement with a backward CS using rats as subjects. Together, the results of these studies lend support to previous claims of excitatory conditioning following backward pairings (e.g., Ayres, Haddad, & Albert, 1987; Burkhardt, 1980; Heth, 1976; Mahoney & Ayres, 1976; Spetch, Terlecki, Pinel, Wilkie, & Treit, 1982; Tait & Saladin, 1986). Interestingly, conditioned reinforcement in Prével et al.

experiment was of similar strength following forward and backward pairings, a result in contrast with most of evidence of excitatory backward conditioning (e.g., Mahoney & Ayres, 1976). The evidence of excitatory backward conditioning in a conditioned reinforcement preparation represents a new challenge to traditional associative models (e.g., Kamin, 1969; Mackintosh, 1975; Rescorla, 1968, 1972; Rescorla & Wagner, 1972), and more generally to the informational hypothesis principle that prediction error reduction is necessary for learning to occur (e.g., Bush & Mosteller, 1951; Stout & Miller, 2007; often elaborated today as total prediction error reduction, e.g., Egger & Miller, 1962, 1963; Niv & Schoenbaum, 2008; Rescorla & Wagner, 1972; Schultz & Dickinson, 2000; but see Witnauer, Urcelay, & Miller, 2014, for a critical review). According to all versions of this hypothesis, Pavlovian learning is driven by the difference between the outcome predicted during a trial and the outcome actually experienced by the subject. The evidence of conditioned excitation with a backward CS challenges this assumption because the backward CS does not have a predictive relationship with the US, and consequently should not result in excitatory learning (for discussion, see Arcediano & Miller, 2002; Savastano & Miller, 1998).

In contrast, models based on temporal integration such as Miller's Temporal Coding Hypothesis (e.g., Arcediano & Miller, 2002; Savastano & Miller, 1998) assume that conditioned excitation with a backward CS is possible. The Temporal Coding Hypothesis states that: 1) temporal contiguity between two events (rather than a predictive relationship) is necessary and sufficient for the formation of associations, 2) the temporal relationship between the two events is learned as part of an association (i.e., cognitive temporal maps are formed), and 3) the encoded temporal relationship plays a critical role in the topology of the response elicited when the CS is subsequently presented. Finally, 4) a subject has the ability of superimpose temporal maps when they share a common element, thereby creating a temporal relationship between elements even when they were never physically paired. As a result, the Temporal Coding Hypothesis is distinguished from the prediction error reduction hypothesis and traditional associative models by the assumption that associations can be learned not only from forward pairings, but simultaneous and backward pairings as well, as long as good temporal proximity exists between the CS and the US. Moreover, the Temporal Coding Hypothesis makes a distinction between what is learned by a subject and the expression of that learning, and assumes that the expression of learning depends on the predictive value of a CS at the time of testing. Thus, it predicts that a backward CS would support excitatory responding to a second CS provided, through temporal integration, a predictive relationship is arranged between the second CS and the US (e.g., Barnet & Miller, 1996). Applying these assumptions to the Prével et al. (2016) experiment, the participants presumably learned the US → CS relationship during the conditioning phase based on the strong temporal contiguity used in that experiment (i.e., a 100-ms inter-stimulus interval). Then, when the backward CS was delivered alone as a consequence of completing the FR12 schedule during the test phase, its presentation may have activated a representation of the US through the association learned during the conditioning treatment, which presumably made the participants expect the presentation of the US during the presses made toward completing the FR12 schedule, thereby conferring a predictive value to the operant response.

Another model that predicts excitatory backward conditioning is Wagner' Sometimes Opponent Processes (SOP) model (1981; see also Mazur & Wagner, 1982). Thus, it provides an alternative to the Temporal Coding Hypothesis as an explanation of excitatory backward conditioning without assuming temporal integration. SOP is a connectionist model in which CSs and USs are each represented by nodes composed of a large but finite number of elements, and in which elements are connected through associative links. Each element can be discretely in one state or another of three different states: a state of inactivity (I), a state of primary activity (A1), and a state of secondary activity (A2). When a stimulus (e.g., a US) is presented, SOP assumes that for each moment the stimulus is presented there is a constant probability P(I) that any element of the node that is in the inactivity state I will transfer to the state of primary activity A1. Then, any element in the A1 state will “decay” to the A2 state with the probability P(d1), and from the A2 state to the I state with the probability P(d2). According to SOP, an excitatory association is formed when elements of two nodes representing the CS and the US are co-activated in the A1 state in the subject’s memory. Once an excitatory association is formed, presentation of the CS causes probabilistic activation of associated US node elements that are in the I state into the A2 state. Conditioned responding to a CS is a function of the number of the proportion of US node elements activated from the I state to the A2 state. Because the model assumes that at termination of the presentation of a US, elements of the US node may continue to reside in the A1 state for some time until they decay into the A2 state, good temporal contiguity during backward US → CS pairings may result in some overlap of the CS and US nodes in the A1 state and, consequently produce an excitatory CS – US association despite the pairings having been backward. However, according to SOP when a CS node is in the A1 state and a US node has decayed into the A2 state, an inhibitory CS – US association is formed. Consequently the model’s prediction of excitatory backward conditioning depends largely on the decay rate parameter for the US node to go from A1 to A2. Slow decay should favor excitatory conditioning, whilst fast decay should lead to inhibitory conditioning. Applying this assumption to the Prével et al. (2016) experiment, and assuming that the rules for conditioned reinforcement are similar to the rules assumed for CR, the high temporal contiguity between the US and the CS should have supported the co-activation of their representative nodes in the A1 state, resulting in conditioned excitation and consequently the conditioned reinforcement effect that was observed.

Because the Prével et al. (2016) experiment was unable to discriminate between Miller's Temporal Coding Hypothesis and Wagner' SOP accounts of the observed excitatory backward conditioning with conditioned reinforcement, the present study was intended to replicate and extend the results of that experiment, but using a design in which the two explanations would make different predictions. In addition, the present study aimed to investigate again the circumstances under which a conditioned stimulus will support excitatory conditioning, as a means of testing the assumption that associative learning in Pavlovian conditioning scenarios is driven by *prediction* error reduction (Niv & Schoenbaum, 2008; Rescorla & Wagner, 1972; Schultz & Dickinson, 2000). In order to achieve this, we decided to conceptually replicate with humans an instrumental (i.e., conditioned reinforcement) preparation modeled after prior studies by Barnet and Miller (1996) and Cole and Miller (1999) which used rats in a Pavlovian second-order conditioning (SOC) preparation. In the procedure developed by Barnet and Miller (1996) and Cole and Miller (1999), *backward pairings* between the US and the first-order CS (i.e., US → CS1 pairings) were administered during Phase 1. The backward CS1 was seen to pass summation and retardation tests for conditioned inhibition, a result commonly found when many backward US → CS pairings are administered and when is assessed by the ability of the backward CS to decrease a CR to an excitatory CS (see for example Heth, 1976; Moscovitch & Lolordo, 1968; Siegel & Domjan, 1971, 1974). But interestingly, the first-order backward CS (i.e., CS1) also supported second-order excitatory conditioning of CS2 following CS2 → CS1 trials, with excitatory responding observed to CS2. These results are of interest for two reasons. First, the evidence of excitatory conditioning to CS2 implies that an association was formed (or at least learning occurred) between the US and the backward first-order conditioned stimulus (i.e., CS1) during Phase 1, a result unexpected if Pavlovian learning was driven by *prediction* error reduction. Second, such a result is difficult to explain based on Wagner' SOP because, according to that model, SOC depends on a CS2 → CS1 → US excitatory chain, with the CS1 → US association being excitatory (see Wagner & Brandon, 1989). As a consequence, the conditioned inhibitory control by CS1 observed in Barnet and Miller's (1996) and Cole and Miller's (1999) experiments is challenging for this account of SOC when Phase 1 is backward paired.

As previously stated, the present study was designed to contrast the account of SOP with that of the Temporal Coding Hypothesis concerning the results of the Prével et al. (2016) experiment. More generally, the present study aimed to illuminate the mechanisms underlying excitatory higher-order conditioning following backward pairings, as they speak to the assumption that Pavlovian conditioning is driven by *prediction* error reduction. To achieve this, we embedded an extension of the Barnet and Miller (1996) and Cole and Miller (1999) procedure within an instrumental conditioned reinforcement preparation. In addition to the above-mentioned theoretical question, the present experiment aimed to replicate previous findings with humans concerning backward conditioning (e.g., Arcediano, Escobar, & Miller, 2003), but now exploit the sensitivity of conditioned reinforcement in an instrumental situation. The present experiment was conducted with human participants, used a cover story in which 'points' served as the 'primary' reinforcer, and involved five phases (see Figure 1). Phase 1 was intended to demonstrate the appetitive value of the primary reinforcer using a concurrent fixed-ratio schedule similar to the Prével et al. (2016) study. Phase 2 consisted of first-order conditioning, and Phase 3 was concerned with demonstrating the absence of a CR to the backward first-order CS. In Phase 4, the critical CS2 → CS1 pairings occurred, and the final phase consisted of testing whether CS2 could serve as an effective conditioned reinforcer, despite the absence of conditioned excitation to CS1 as measured in Phase 3. This test was conducted using a concurrent fixed-ratio schedule without the primary reinforcer being presented (i.e., in extinction), and like in the experiments in Barnet and Miller (1996) and Cole and Miller (1999), evidence of conditioned reinforcement to CS2 would be interpreted as a challenge to both the necessity in learning of *prediction* error reduction and Wagner's SOP account of SOC with backward conditioning in Phase 1. In other words, such a result would support models based on temporal integration like the Temporal Coding Hypothesis, and suggest that a predictive relationship between a CS and a US is not necessary for Pavlovian conditioning to occur, but may be necessary for responding at the time of testing.

Method

Participants

Twelve undergraduate students at the University of Lille, Nord de France participated in this study. The participants were given credits towards completing a course requirement, and an informed consent form was read and signed by the participants.

Apparatus

The experiment was run on a PC and programmed using Matlab. It was conducted in a small, darkened, quiet room, and participants were seated on a chair such that their eyes were roughly 40 cm from the center of the computer monitor (1600 x 900 pixels). Points served as the appetitive US, and participants were instructed that they had to earn 80 points to complete the task. Participants left the experiment when they completed the task, even if they failed to earn 80 points. The first-order CSs consisted of five 4 cm diameter colored circles (blue, green, red, yellow, and purple), and the second-order CSs were white, capital letters (A, B, C, D, and E; the font was Arial) with a font size value of 80 based on Matlab programming language (i.e., approximately 2.8 cm). The screen consisted of a black background with a points tally in the top right corner, and two 3.5 x 3.5 cm grey squares separated by an 11 cm space on which participants could left click a computer mouse were used as response keys during Phases 1 and 5. The two response keys were reinforced on a FR12 schedule of reinforcement.

Procedure

Unlike the Prével et al. (2016) experiment, in this study we required an anticipatory response during the concurrent FR phases (Phases 1 and 5), as well as during the conditioning phases (Phases 2 and 4) and test for first-order conditioning (Phase 3; see Figure 1 for the complete design): On-screen targets "←", "→", and "SPACE" were used as discriminative stimuli (S^D) and were presented at the center of the monitor. Arrows were 3.5 cm long, and the word SPACE was written with a font size of 60 based on the Matlab programming language. In Phases 1 to 5, during presentations of these targets, participants were required to press a key on the keyboard matching the target presented. Targets lasted 1200 ms, and a correct press (i.e., pressing the key matching the S^D) was immediately followed by a 'positive' tone (i.e., 440 Hz, 200 ms) and points were added to the participant's point tally. When participants made an incorrect press (i.e., pressing an incorrect key or pressing a key before or after presentation of a discriminative stimulus), a 'negative' tone (i.e., 220 Hz, 200 ms) immediately occurred and there was no change in points. US1 and US2 equaled 0.2 and 2.0 points, respectively. Left and right arrows (i.e., "←"; "→") were discriminative stimuli (S^{D1} and S^{D2}) assigned to US1 and US2 (counterbalanced across participants). The target "SPACE" (S^{D3}) was always assigned to US3. US3 resulted in the addition of 0.5 points. In addition to the temporal arrangement between the CSs and the S^Ds, magnitude of reinforcement was also manipulated in the experiment. The reason for the magnitude of reinforcement manipulation resided in the possibility of that we were creating two distinct predictions between the Temporal Coding Hypothesis and SOP with respect to the test for the excitatory second-order conditioning in Phase 5. Strictly speaking, the test consisted of a concurrent FR12 schedule leading to delivery of a backward second-order CS and a forward second-order CS[1]. The backward second-order CS was associated (through the first-order backward CS) with S^{D2}, and the forward second-order CS was associated (through the first-order forward CS) with S^{D1}. As a consequence, the backward second-order CS was associated to a larger magnitude of reinforcement (i.e., 2.0 points) than the forward second-order CS (i.e., 0.2 points). If responding were solely guided by US magnitude, this should be reflected in stronger responding to the backward CS1 during Phase 3. Thus, a unique prediction of the Temporal Coding Hypothesis is that higher responses to the backward second-order CS should be observed during Phase 5 despite lower responses to the backward first-order CS during Phase 3. In contrast, SOP would predict higher responses to the forward second-order CS in Phase 5 if an absence of responding is observed to the first-order backward CS (and responding is seen to the first-order forward CS) because SOP assumes that second-order conditioning results from the first-order CS (i.e., CS1) being excitatory.

Thus, the experiment consisted of the five phases as depicted in Figure 1, and used a fully within-subject design. First-order and second-order CSs presentations lasted 1200 ms. Below are the general instructions (translated from French) that were presented to the participants at the beginning of the experiment:

"In this experiment, your objective is to collect points by responding correctly to targets presented on the computer screen. The experiment will end when you will reach a score of 80 points, so the time that you will spend in this experiment will depend on your performance in responding to the targets. Targets consist of an arrow oriented toward the left, an arrow oriented toward the right, and the word "SPACE". Responding consists of pressing the key on the keyboard corresponding to the target presented. For instance, if the word "SPACE" is presented on the screen, you will have to press the space bar on the keyboard to collect points. Note that targets are presented for 1200 ms and you can collect points only if

[1] The terms 'backward' and 'forward' mean that the first-order CSs were 'backward' and 'forward' CSs, respectively.

you press a correct key while a target is presented. One of the two arrows will earn 2.0 points, whereas the other arrow will earn 0.2 points. The word "SPACE" will always be associated with 0.5 points earned. Your score will be shown on the top-right corner of the screen. Correct responses will be accompanied by a positive sound and incorrect responses by a negative sound. Finally, the experiment is divided into several phases, and you can choose to stop the experiment at any moment if you want."

[insert Figure 1]

Phase 1 – Concurrent schedule with appetitive USs. Phase 1 consisted of training on two concurrent FR12 schedules of reinforcement with US1 (i.e., +0.2 points) for making Response 1 (R1) and US2 (i.e., +2.0 points) for making Response 2 (R2; see Figure 2). Phase 1 served two purposes. First, it allowed participants to gain experience on the FR schedules; second, it provided independent evidence that participants were sensitive to the different values of the two USs. The left and right squares on the screen were used as response keys, and responses consisted of clicking on one of the squares after pointing to it with the computer mouse. The following instructions were presented to the participants before Phase 1:

"In the first phase of the experiment, you will have to make choices between two buttons presented on the left and right sides of the screen. A choice consists in clicking twelve times with the computer mouse on one of the two buttons, and it will be automatically followed by one of the two arrows. For instance, after twelve clicks on the left side button, the arrow oriented to the right might be presented for 1200 ms. The same target is indicative of a specific button throughout the phase. Remember that one of the two targets earns 2.0 points, whereas the other arrow earns 0.2 points will automatically be added to your score. Finally, note that the program only records your first key press on each trial, so pay attention and only press a key when a target is presented on the screen."

The sequence during Phase 1 was: (1) after an intertrial interval of 4 s, (2) left and right response keys (i.e., grey squares) appeared simultaneously on the monitor and participants responded by clicking on the squares with the computer mouse; after ratio completion (i.e., twelve responses) on one of the two squares; (3) the S^Ds appeared (i.e., either the left arrow ['←'] appeared immediately following ratio completion on the left response key and the right arrow [i.e. '→'] appeared immediately following ratio completion of the right response key). Target presentations (i.e., left or right arrows) did not automatically lead to earned points; matching responses on the keyboard were necessary for participants to earn points. (4) pressing the left arrow key on the keyboard once during this discriminative stimulus S^{D1} presentation added 0.2 points (i.e., US1) to the counter, whereas pressing the right arrow key of the keyboard once during that discriminative stimulus S^{D2} added 2.0 points (i.e., US2) to the counter. The discriminative stimuli (i.e., left or right arrows) leading to US1 and US2 were counterbalanced across participants. Phase 1 lasted for 10 trials with one of the two S^D presented during each trial. Note that, to avoid participants' consistently pressing the keys, an intertrial interval and the following trial were considered as a unit during which the program only recorded the first response emitted by the participant. Subsequent responses during a unit (i.e., pressing a key of the keyboard after the first press of a trial) had no consequences. As a result, a participant's pressing one key (e.g., the left arrow key) during the ITI or during FR12 completion precluded the participant's earning points during the presentation of an S^D. This constraint was applied throughout the experiment, and participants were informed of this constraint during the instructions.

[insert Figure 2]

Phase 2 – First-order conditioning. First-order conditioning was conducted during Phase 2. Forward pairings were conducted between CS1_F and S^{D1} (CS1_F → S^{D1}), backward pairings between CS1_B and S^{D2} (S^{D2} → CS1_B), simultaneous pairings between CS1_S and S^{D3}

(CS1_S – S^{D3}), while CS1_N1 and CS1_N2 were presented alone. Simultaneous CS1_S – S^{D3} pairings were of no theoretical interest; they were intended to improve discrimination between the stimuli, and to train participants to press a key on the keyboard during a CS and not only during an S^{D}. CS1_N1 and CS1_N2 trials served as control conditions needed to assess first-order (and in Phase 5, second-order) conditioning. The following instructions were presented to the participants before Phase 2:

"In the second phase, targets will now be automatically presented during series of trials that will be divided into three sub-phases. During this phase, it is unnecessary for you to click on any buttons on the screen with the computer mouse. A trial will start with a brief fixation cross presented on the screen. The fixation cross will then be followed by one of the three targets after a short interval, but note that on some trials no target will be presented. Again, you will have to respond to the targets by pressing the key on the keyboard corresponding to the target presented. Note that the targets will be accompanied by circles of different colors and you will have to pay attention to these stimuli because they are relevant to the following phases of the experiment. Again, the program will record only the first key press that you will make during a trial."

We divided Phase 2 in three conditioning sub-phases. Each CS and US were presented four times per sub-phase, with a total of twelve presentations of each CS and US pair by the end of Phase 2. Figure 3 summarizes a forward CS1_F → S^{D1} conditioning trial. A conditioning trial consisted of: (segment 1) a fixation cross with a 250 ms duration signaling a new trial, (segment 2) an interval of 4000 ms of a blackout screen, (segment 3) CS1_F presented during 1200 ms, followed by (segment 4) S^{D1} for 1200 ms, and finally (segment 5) a blackout screen for 1200 ms. Trials were separated by a 2000-ms intertrial interval. During backward S^{D2} → CS1_B pairings, a blackout screen replaced the CS presentation during segment 3, and CS1_B was presented during segment 5. During the simultaneous pairing trials, CS1_S and S^{D3} were presented in segment 4 and no CS was presented in segments 3 and 5. CS1_N1 and CS1_N2 trials were equivalent to forward and backward conditioning trials except that no US was presented. The order of CSs presentations was randomized across participants within sub-phases except that a CS was never followed by the same stimulus on successive trials. The targets (i.e., S^D) were presented at the center of the screen, and the CSs were presented above the S^D with no overlap between the stimuli. Like in Phase 1, during S^D presentations participants had to press the correct key of the keyboard in order to collect points (i.e., USs; e.g., pressing the left arrow key of the keyboard during the left arrow discriminative stimulus presentation). This is illustrated in the bottom-right box of the Figure 3, with a US delivered during S^D presentation provided the participant pressed the correct key during that stimulus.

[insert Figure 3]

Phase 3 – Test of first-order conditioning. The procedure used to test first-order conditioning was similar to Phase 2 except that no S^D was presented, nor were the sounds for correct and incorrect responses presented. The points earned by the participants were also hidden. To keep subject's motivation constant, they were informed that the S^Ds were hidden, but that they had to continue to earn points by guessing when the S^Ds were presented, based on the CS presentations. The following instructions were presented to the participants before Phase 3:

"Phase 3 is similar to Phase 2 except that you will no longer see when the targets are presented nor the score you have earned. Moreover, the sounds signaling a correct or an incorrect response will no longer be delivered. Consequently, in order to collect points, you will be required to guess when the targets are functional based on the circles presented during the trials. Note that you will not be able to perfectly anticipate a target during a trial. We ask you to respond only when you think that a target should be present."

The program recorded the keys pressed by the participants during segments 3, 4, and 5 as well as during the 1200 ms following each trial. The test phase ended with four presentations of each CS, distributed in four blocks of each CS being presented once per block and order of the CSs within a block being random. We expected that the participants would press the key associated with S^{D1} after CS1_F, and not press the key associated with S^{D2} during or after CS1_B. In addition, we expected participants to press the key associated with S^{D3} during CS1_S, and not to press either key during or after CS1_N1 or CS2_N2.

Phase 4 – Second-order conditioning. The procedure used for second-order conditioning was similar to Phase 3 but with second-order CSs added. The second-order CSs were presented in Phase 4 paired with the first-order CSs (see Figure 4), with CS1_F → CS2_F pairings, CS2_B → CS1_B pairings, CS1_S – CS2_S simultaneous pairings, CS1_N1 → CS2_N1 pairings, and CS2_N2 → CS1_N2 pairings. The following instructions were presented to the participants before Phase 4:

“Phase 4 is divided in two sub-phases. This phase is similar to Phase 3 except that a letter will accompany the presentation of the circles during trials. Like in Phase 3, you have to guess when the targets should be presented in order to collect points, based on the stimuli appearing on the screen. Again, the score you have earned will no longer be shown and the sounds signaling a correct or an incorrect response will not be delivered.”

The second-order conditioning trials were arranged to produce a hypothetical temporal overlap between the forward, backward, and simultaneous CS2s and their respective CS1-evoked S^Ds representations, that is, critically, temporal overlap of CS2_F and the representation of S^{D1} evoked by CS1_F, and CS2_B and the representation of S^{D2} evoked by CS1_B. In each of the five conditions, the CS1s and CS2s were paired eight times by the end of Phase 4, distributed in eight blocks with each CS being presented once per block and the order of CSs within a block being random.

[insert Figure 4]

Phase 5 – Concurrent schedules with second-order CS as a test of second-order conditioning. A conditioned reinforcement test to assess the success of Phase 4 in producing SOC was conducted in Phase 5 and consisted of two sub-phases with second-order forward and backward CSs tested in one sub-phase and the second-order neutral CSs tested in the other sub-phase. The order of the two sub-phases was randomized across participants, and each of the two sub-phases consisted of concurrent FR12 schedules of reinforcement with the second-order CSs. In contrast to Phase 1, discriminative stimuli were not presented during Phase 5, and consequently conditioned reinforcers (i.e., second-order stimuli) were presented alone. In one of the two sub-phases, CS2_F was delivered as a consequence of FR12 completion on the R2 key and CS2_B was delivered as a consequence of FR12 completion on the R1 key. The following instructions were presented to the participants before Phase 5:

“Phase 5 is divided into two sub-phases with choices similar to Phase 1, except that the letters presented during Phase 4 will now be presented. In the first sub-phase, you will have to complete a series of choices with the help of the computer mouse in which two letters are presented. Like in Phase 1, the same stimulus will always be paired with a button during the entire sub-phase, and like in Phases 3 and 4, you are asked to guess when targets should be presented and which targets they are. The second sub-phase is exactly the same as the first one, except that two others letters will be presented.”

The contingency was similar to Phase 1 and can be summarized as follows: After an inter-trial interval of 4 s, left and right response keys appeared in the monitor. CS2_F immediately followed ratio completion on one response key, and CS2_B immediately followed ratio completion on the other key. Keep in mind that forward CS1_F → S^{D1} pairings had been carried-out in Phase 2 and followed by CS1_F → CS2_F pairings in Phase 4, with a hypothetical temporal overlap between representations of S^{D1} and CS2_F. Moreover, a backward S^{D2} → CS1_B has been conducted in Phase 2 as well as a CS2_B → CS1_B in Phase 4, with a hypothetical temporal overlap between representations of S^{D2} and CS2_B. Because S^{D1} was associated with a smaller magnitude of reinforcement than S^{D2} (US1 was 0.2 points and US2 was 2.0 points), evidence for second-order conditioning to CS2_B could be assessed as the number of responses allocated to the response key delivering that stimulus. As before, the test sub-phase consisted of 10 trials and the number of responses on the two keys was recorded. In the second test sub-phase, CS2_N1 and CS2_N2 were delivered as a consequence of ratio completion with a similar contingency (FR12). This test served as a control condition, and no significant differences between the two keys were expected because the stimuli had both been paired with neutral first-order CSs. However, some responses on the two keys were expected because the sub-phase consisted of 10 forced choices, requiring the participants to complete one of the two FR12 schedules during ten trials in order to complete the task. Consequently, the absence of conditioned reinforcement effects with respect to CS2_N1 and CS2_N2 was evidenced by the absence of differences in responding on the two keys.

Results

During Phase 1, we measured how the participants allocated their responses across concurrent FR12 schedule, in which responses on one response key delivered S^{D1} (associated with US1, i.e., +0.2 points) and responses on the other delivered S^{D2} (associated with US2, i.e., +2.0 points). Figure 5 summarizes the results showing the mean number of responses measured on the two keys. In Figure 5 and all subsequent figures, the SEMs depicted have been adjusted for the within-subject design of the experiment (Cousineau, 2005). The participants allocated more responses to the key delivering S^{D2} (i.e., US2), and an analysis of variance (ANOVA) revealed an effect of the US delivered, $F(1, 11) = 2209.$, $p < .05$, $\omega^2 = 0.99$. This indicates that participants were sensitive to the different values of the USs. As a reminder, participants were informed through the instructions of the two FR12 schedules running on the response keys. Consequently, a strong preference for the response key delivering S^{D2} was expected if participants were sensitive to the different values of the USs.

[insert Figure 5]

During Phase 3, we measured the responses that the participants emitted during the first-order CSs presentations. The results summarized in Figure 6 represent the mean percentage of CRs in the presence of each of the first-order CSs. The criterion for a CR was a response (i.e., key pressing) emitted during a CS or during the 1200 ms following that CS. The results show that participants emitted responses to CS1_F and CS1_S, reflecting their respective US values (i.e., US1 and US3), but the participants did not respond appreciably to CS1_B, despite it having being paired with the largest number of points (i.e., +2.0 points). Note also that for CS1_B the participants never emitted CRs during the 1200 ms preceding that stimulus. Finally, participants did not respond to CS1_N1 or CS1_N2, the two neutral stimuli. An ANOVA conducted on the CRs observed to CS1_F (US1), CS1_S (US3), and CS1_B (US2) revealed a difference across CSs, $F(2, 22) = 187.63$, $p < .05$, $\omega^2 = 0.89$. This pattern suggests that participants learned the *predictive* value of each CS1, and consequently responded robustly to the forward (CS1_F) and simultaneous (CS1_S), but not to the backward CS1_B presumably because it signaled the impending absence of a US. Critically, this difference was observed despite CS1_B having been paired with 2.0 points, whereas CS1_F was paired with only 0.2 points (and CS1_S was paired with 0.5 points). Temporal order had a greater effect on Phase 3 responding than did magnitude of reinforcement.

[insert Figure 6]

Figure 7 (Top panel) depicts how the participants allocated their responses across the concurrent FR12 schedule during Phase 5 when CS2_F was delivered for responding on one response key and CS2_B was delivered for responding on the second response key.

Specifically, Figure 7 summarizes the total number of responses measured on each of the two keys. The participants allocated distinctly more responses on the key delivering CS2_B than CS2_F. This was supported by an ANOVA which revealed an effect of CS, $F(1, 11) = 61.84$, $p < .05$, $\omega^2 = 0.84$.

[insert Figure 7]

Finally, Figure 7 (Bottom panel) also shows how the participants allocated their responses across the concurrent FR12 schedule during Phase 5 when one response key delivered CS2_N1 and the other delivered CS2_N2. No significant difference was detected between responding on the two keys, $F(1, 11) = 0.28$, $p < .05$.

Discussion

To summarize our findings, three different US magnitudes (US1 < US3 < US2) were paired with forward, simultaneous, and backward CS1s, respectively, in Phase 2. Despite US2 being larger than US1 and participants demonstrating sensitivity to this difference in Phase 1, in Phase 3 CS1_B elicited less responding than CS1_F presumably because of CS1_B being a backward as opposed to forward CS. During Phase 4, participants were exposed to second-order conditioning with CS1_F → CS2_F pairings, CS2_B → CS1_B pairings, CS1_S – CS2_S simultaneous pairings, as well as CS1_N1 → CS2_N1 and CS2_N2 → CS1_N2 pairings. A test for second-order conditioning was carried-out in Phase 5 with CS2_F as a consequence of one response and CS2_B as a consequence of the other response in one sub-phase, and CS2_N1 and CS2_N2 as consequences of the two responses in the other sub-phase. Participants allocated more responses to the response key delivering CS2_B (relative to CS2_F), but no significant difference was observed between CS2_N1 and CS2_N2, suggesting that CS2_B served as a conditioned reinforcer.

The critical finding here is more responding during Phase 5 for CS2_B than for CS2_F (i.e., more responding on R1 than R2, see Figure 7). Importantly, in Phase 1, the larger magnitude of reinforcement for R2 (US2 = 2.0 points) relative to R1 (US1 = 0.2 points) resulted in *more* responding on R2 (see Figure 5). Thus, the preference for CS2_B in Phase 5 cannot simply be attributed to the difference in reward magnitude during Phase 1.

The present experiment conceptually replicated the backward SOC experiments with rats reported by Barnet and Miller (1996) and Cole and Mille (1999) in which excitatory responding to a second-order CS was observed while no responding to the paired first-order CS was observed. Here we extended those findings to an instrumental conditioned reinforcement preparation with human participants. Moreover, the present experiment replicated the recent demonstrations of conditioned reinforcement with a backward CS by Prével et al. (2016). The results found by Prével et al. were interpreted as a challenge to the view that Pavlovian conditioning is driven by prediction error reduction (e.g., Niv & Schoenbaum, 2008; Schultz & Dickinson, 2000) and as consistent with both Miller's Temporal Coding Hypothesis (e.g., Arcediano & Miller, 2002; Savastano & Miller, 1998) and Wagner' (1981) SOP model. The present experiment extended the results of Prével et al. (2016) and contrasted the two accounts of those data. The evidence of conditioned reinforcement with the second-order backward CS, even when no CR was observed to the first-order backward CS, like the results of Barnet and Miller (1996) and Cole and Miller (1999) with rats, is inexplicable in terms of the SOP model as well as models of learning that assume learning is driven by a reduction in predictive error.

Both Wagner's (1981) SOP model and prediction error models of learning erroneously anticipate no responding to CS2_B in light of the absence of excitatory responding to CS1_B. As previously stated, prediction error models (e.g., Niv & Schoenbaum, 2008; Schultz & Dickinson, 2000) assume that learning in Pavlovian conditioning is driven by the difference between the outcome predicted by cues presented during a trial and the outcome actually experienced by the subject. Whereas these models correctly anticipate an absence of CR to the backward CS (i.e., CS1_B, Phase 3) because the backward CS does not have a predictive relationship with the US, the evidence of conditioned reinforcement value of the second-order CS (i.e., CS2_B) during Phase 5 is challenging for these accounts. Indeed, it is hard to see how a conditioned stimulus can support excitatory second-order conditioning when, at the same time, it was not involved in (positive) prediction error during first-order conditioning due to the backward arrangement, unless one assumes that an association was learned between SD^2/US2 and CS1_B during Phase 2. But such an assumption is inconsistent with a prediction error being necessary for learning to occur. In contrast with prediction error models, Wagner’s (1981) SOP assumes that such a backward association can be learned under some circumstances. However, the results of the present experiment, like those of Barnet and Miller (1996) and Cole and Miller (1999), also challenge this account. As previously stated, SOP assumes that SOC is based on an excitatory value of CS1 (or at least the early portions of CS1; see Williams & Hurlburt, 2000), and responding to CS2 is assumed to result from an excitatory CS2 → CS1 → US associative chain (Wagner & Brandon, 1989). Thus, a direct CS2 → US association is not expected to have been established during Phase 4 according to the model because of the US node being in the A2 state due to presentation of CS1. The contiguous pairings during first-order conditioning are assumed to have co-activated the CS1 and US nodes in the A1 state and consequently produced excitatory conditioning of CS1, thereby providing the possibility for CS1 presentation to provoke a transition of the US node to the A2 state and produce a CR. With the assumption that CSs in either A1 or A2 states can contribute to the spread of activation (see Mazur & Wagner, 1982), the CS2 → CS1 pairing during second-order conditioning permits the establishment of a CS2 → CS1 → US excitatory associative chain and, consequently, conveys to CS2 the potential of CS1 to elicit CR. However, it is unclear how SOP could handle at the same time the absence of CR to CS1 and excitatory responding to CS2. Even if one relaxes SOP’s rules and allow A2 – A1 excitatory learning (see Holland, 1983), the presumed absence of the US in the A2 state in response to CS1 presentation (because of the absence of CRs recorded to CS1_B) seems to preclude the possibility for a CS2 → CS1 → US associative chain.

In contrast to Wagner's (1981) SOP model, Miller's Temporal Coding Hypothesis can account for the present observations. The present evidence of second-order conditioning supported by a backward CS is consistent with predictions of the Temporal Coding Hypothesis (Arcediano & Miller, 2002; Savastano & Miller, 1998), and with the more general notion of temporal integration. According to this account, the participants formed bidirectional associations between the first-order CSs and USs including their temporal relationship during Phase 2, and between the first-order CSs and second-order CSs including their temporal relationship during Phase 4. The hypothetical temporal overlap between representations of second-order CSs and USs representations, mediated by the first-order CSs, presumably conveyed to at least the early components of the second-order CS a predictive relationship regarding the later components of the USs as well as a conditioned reinforcing value in Phase 5. The larger amount of points associated to US2 would convey to CS2_B its larger appetitive value (measured by the preference observed during Phase 1). Alternatively stated, based on the previously described integration of temporal maps which here result in CS2_B and CS2_F having the same simultaneous temporal relationship with their respective S^Ds, the model anticipates equal responding on R1 and R2 if the procedural differences between these two conditions were only their temporal relationships in Phases 2 and 4. But the greater magnitude of reinforcement associated with R2 (i.e., US2) than R1 (i.e., US1) leads to an expectation that R2 will be preferred over R1, which is what was observed. That is, US1 being smaller than US2 is presumably the basis for more responding for CS2_B than CS_F in Phase 5. These results are also consistent with recent considerations of Thrailkill and

Shahan (2014) on the application of temporal integration to conditioned reinforcement preparations.

Note that in the Prével et al. (2016) paper we also proposed an account of backward conditioned reinforcement based on the Bio-behavioral model of Donahoe and Palmer (2004, p. 100). This model assumes that the ability of a second-order CS to function as an eliciting stimulus or a conditioned reinforcer depends on the excitatory value of CS1. The evidence found in this experiment, that CS2 serves as a conditioned reinforcer even when conditioned excitation is not evident in CS1, challenges the assumptions of the Bio-Behavioral model as well as that of Wagner's (1981) SOP. Moreover, the present results seem to represent a challenge for some theories of timing (see Arcediano et al., 2003) like the Scalar Expectancy Theory (SET, e.g., Gibbon et al., 1984). In SET, responding is determined by a response rule in which the system compares the absolute discrepancy between the current time in working memory from the onset of a [forward] stimulus with the time to reinforcement recorded in long-term memory from prior pairings. The evidence of backward SOC is difficult to explain in terms of SET because there is no reason to expect in terms of that model that the subjects learn to time a US with respect to a backward CS. A solution to explain the evidence of backward SOC in terms of SET could reside in the assumption that the subjects learned the time interval between the first-order backward CS at trial n and the US at a trial n + 1. However, our use of intermixed pairings in which the stimuli presented at a trial were never followed by the same stimuli at a next trial precluded this possibility.

Finally, it could be argued that the main results of this study, that is higher responding for CS_B over responding for CS_F, results from CS_F becoming inhibitory during Phase 4. In other words, it could be argued that CS2_F developed inhibitory properties during Phase 4 due to presentations of CS1_F and CS2_F, together with the omission of S^{D1}/US1 (similar to the A+/AX- design). Larger responding observed to CS2_B during Phase 5 would be the result of participants avoiding the presentation of that inhibitory CS2_F instead of the acquisition of conditioned reinforcing properties by CS2_B. However, this interpretation in terms of inhibitory properties acquired by CS2_F is unlikely for different reasons. First, Phase 4 arranged only eight second-order conditioning trials per CS, and conditioned inhibition classically appears when many trials are arranged (Yin, Barnet, & Miller, 1994). Moreover, second-order conditioning trials were not interspersed with CS1 - S^{D}/US reinforced trials in our experiment, and CS1_F - CS2_F were serially presented, which delays the development of conditioned inhibition (Stout, Escobar, & Miller, 2004). Finally, just as CS2_F second-order conditioning trials omitted the presentation of S^{D1}/US1, CS2_B second-order conditioning trials also omitted the presentation of S^{D2}/US2. Consequently, if CS2_F acquired inhibitory properties during Phase 4 CS2_B should also have acquired those properties. Thus, it is difficult to understand why CS2_B resulted in much more responding than CS2_F if the two were inhibitory. For all these reasons, it seems that an interpretation in terms of inhibitory acquisition by CS2_F is unlikely to account for the observed results. However, future experiments will have to include the recording of CR to the second-order CS as a better control.

In conclusion, the present demonstration of backward conditioning along with other prior examples of backward conditioning is problematic for the widely held view that learning occurs only when there is an error in *prediction*, a position taken both by the informational hypothesis and specific instances of the informational hypothesis such as the Bush and Mosteller (1951) model (also see Mackintosh's [1975] and Stout & Miller's [2007] comparator hypothesis for application of Bush and Mosteller's version of the informational hypothesis) and the total error reduction assumption of the Rescorla-Wagner (1972) model (for a discussion of additional problems with the error reduction hypothesis, see Witnauer et al. [2014]). Nevertheless, there seems to be considerable value in explaining various learning phenomena (e.g., the seemingly diminishing benefit of added training trials) in the notion of error reduction. However, what is challenged by backward excitatory conditioning is not the notion of error reduction, but that the error has to have a predictive value. Perhaps it is time to reframe the informational hypothesis from requiring '*prediction* error' into requiring 'discrepancy with prior experience.' That is, learning appears to depend on a discrepancy from prior experience (or an absence of any relevant experience), but the discrepancy need not be predictive.

Funding statement

The research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of interest

The authors declare that there is no conflict of interest.

References

Arcediano, F., & Miller, R. R. (2002). Some constraints for models of timing: A temporal coding hypothesis perspective. *Learning and Motivation*, *33*, 105-123.

Arcediano, F., Escobar, M., & Miller, R. R. (2003). Temporal integration and temporal backward associations in human and nonhuman subjects. *Learning & Behavior*, *31*, 242-256.

Ayres, J. J. B., Haddad, C., & Albert, M. (1987). One-trial excitatory backward conditioning as assessed by conditioned suppression of licking in rats: Concurrent observations of lick suppression and defensive behaviors. *Animal Learning & Behavior*, *15*, 212-217.

Barnet, R. C., & Miller, R. R. (1996). Second-order excitation mediated by a backward conditioned inhibitor. *Journal of Experimental Psychology: Animal Behavior Processes*, *22*, 279-296.

Burkhardt, P. E. (1980). One-trial backward fear conditioning in rats as function of US intensity. *Bulletin of the Psychonomic Society*, *15*, 9-11.

Bush, R. R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review*, *58*, 313-323.

Cole, R. P., & Miller, R. R. (1999). Conditioned excitation and conditioned inhibition acquired through backward conditioning. *Learning and Motivation, 30*, 129-156.

Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson’s method. *Tutorials in Quantitative Methods for Psychology*, *1*, 42-45.

Donahoe, J. W., & Palmer, D. C. (2004). *Learning and complex behaviors*. Vivian-Packard Dorsel.

Egger, M. D., & Miller, N. E. (1962). Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology*, *64*, 97-104.

Egger, M. D., & Miller, N. E. (1963). When is a reward is reinforcing?: An experimental study of the information hypothesis. *Journal of Comparative and Physiological Psychology*, *56*, 132-137.

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

Heth, C. D. (1976). Simultaneous and backward fear conditioning as a function of number of CS-US pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, *2*, 117-129.

Holland, P. C. (1983). Representation-mediated overshadowing and potentiation of conditioned aversions. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*, 1-13.

Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B.A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279-296). New York, NY: Appleton-Century-Crofts.

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

Mahoney, W. J., & Ayres, J. J. B. (1976). One-trial simultaneous and backward fear conditioning as reflected in conditioned suppression of licking in rats. *Animal Learning & Behavior*, *4*, 351-362.

Mazur, J. E., & Wagner, A. R. (1982). An episodic model of associative learning: Acquisition. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative Analyses of Behavior, 3,* pp. 3-39. Cambridge, M. A.: Ballinger Publishing.

Moscovitch, A., & Lolordo, V. M. (1968). Role of safety in the Pavlovian backward fear conditioning procedure. *Journal of Comparative and Physiological Psychology, 66*, 673-678.

Niv, Y., & Schoenbaum, G. (2008). Dialogues on prediction error. *Trends in Cognitive Sciences, 12*, 265-272.

Prével, A., Rivière, V., Darcheville, J.C., & Urcelay, G. P. (2016). Conditioned reinforcement and backward association. *Learning and Motivation, 56*¸ 38-47.

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66,* 1-15.

Rescorla, R. A. (1972). Information variables in Pavlovian conditioning. *Psychology of Learning and Motivation, 6,* 1-46.

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* (pp. 64-69). New York: Appleton-Century-Crofts.

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

Schultz, W., & Dickinson, A. (2000). Neural coding of prediction errors. *Annual Review of Neuroscience*, *23*, 473-500.

Siegel, S., & Domjan, M. (1971). Backward conditioning as an inhibitory procedure. *Learning and Motivation, 2,* 1-11.

Siegel, S., & Domjan, M. (1974). The inhibitory effect of backward conditioning as a function of the number of backward pairings. *Bulletin of the Psychonomic Society*, *4*, 122-124.

Spetch, M. L., Terlecki, L. J., Pinel, J. P. J., Wilkie, D. M., & Treit, D. (1982). Excitatory backward conditioning of defensive burying in rats. *Bulletin of the Psychonomic Society, 19*, 111-114.

Spetch, M. L., Wilkie, D. M., & Pinel, J. P. J. (1981). Backward conditioning: A reevaluation of the empirical evidence. *Psychological Bulletin*, *89*, 163-175.

Stout, S., Escobar, M., & Miller, R. R. (2004). Trial number and compound stimuli temporal relationship as joint determinants of second-order conditioning and conditioned inhibition. *Animal Learning & Behavior*, *32*, 230-239.

Stout, S. C., & Miller, R. R. (2007). Sometimes competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review, 114*, 759-783.

Tait, R. W., & Saladin, M. E. (1986). Concurrent development of excitatory and inhibitory associations during backward conditioning. *Animal Learning & Behavior*, *14*, 133-137.

Thrailkill, E. A., & Shahan, T. A. (2014). Temporal integration and instrumental conditioned reinforcement. *Learning & Behavior*, *42*, 201-208.

Urushihara, K. (2004). Excitatory backward conditioning in an appetitive conditioned reinforcement preparation with rats. *Behavioural Processes*, *67*, 477-489.

Wagner, A. R. (1981). SOP: A model of automatic processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5-47). Hillsdale, NJ: Erlbaum.

Wagner, A. R. & Brandon, S. E. (1989). Evolution of a structured connectionist model of Pavlovian conditioning (AESOP). In S. B. Klein & R. R. Mower (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 149-189). Hillsdale, NJ: Erlbaum.

Williams, D. A., & Hurlburt, J. L. (2000). Mechanisms of second-order conditioning with a backward conditioned stimulus. *Journal of Experimental Psychology: Animal Behavior Processes*, *26*, 340-351.

Witnauer, J. E., Urcelay, G. P., & Miller, R. R. (2014). The error in total error reduction. *Neurobiology of Learning and Memory*, *108*, 119-135.

Yin, H., Barnet, R. C., & Miller, R. R. (1994). Second-order conditioning and Pavlovian conditioned inhibition: Operational similarities and differences. *Journal of Experimental Psychology: Animal Behavior Processes*, *20*, 419-428.

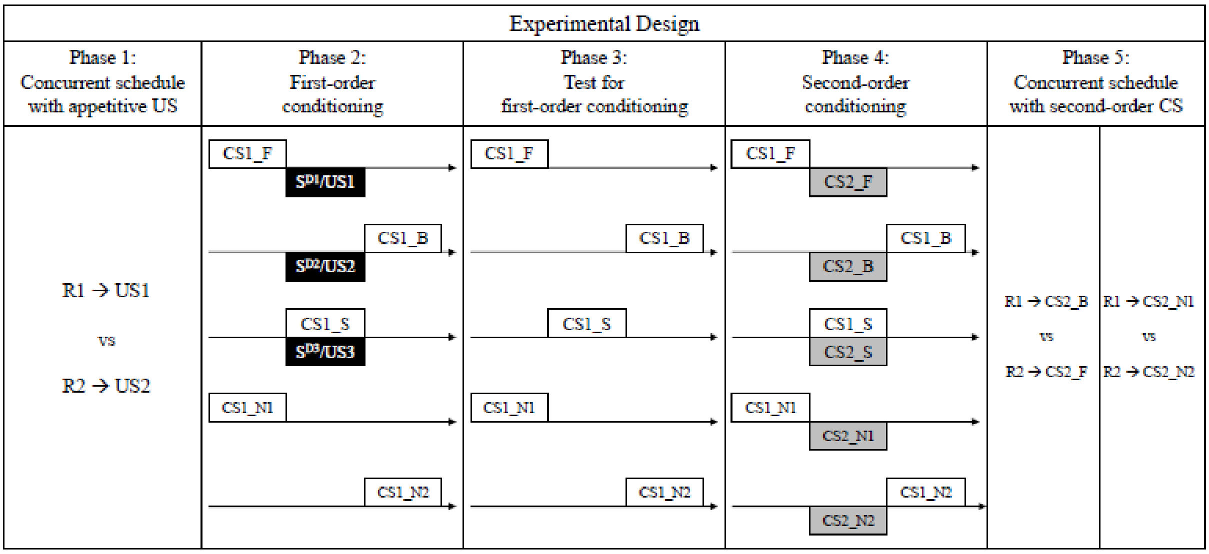

Figure 1. Summary of the experimental design. In Phase 1, participants were trained on concurrent FR12 schedules with R1 and R2 delivering US1 and US2, respectively. Phase 2 consisted of first-order conditioning of CS1 with various conditioning arrangements, Forward (F), Backward (B), Simultaneous (S), CS-alone at the beginning of a trial (N1), and CS-alone at the end of a trial (N2). S^D/US indicates the discriminative stimuli (S^D) were presented and the appropriate US was earned for responding to that S^D. Phase 3 provided a test for first-order conditioning of the CS1s. Phase 4 consisted of second-order conditioning of the CS2s. Phase 5 provided the test for second-order conditioning on the two concurrent FR12 schedules with R1 and R2 delivering CS2_B and CS2_F, respectively, during one sub-phase, and CS2_N1 and CS2_N2 during the other sub-phase.

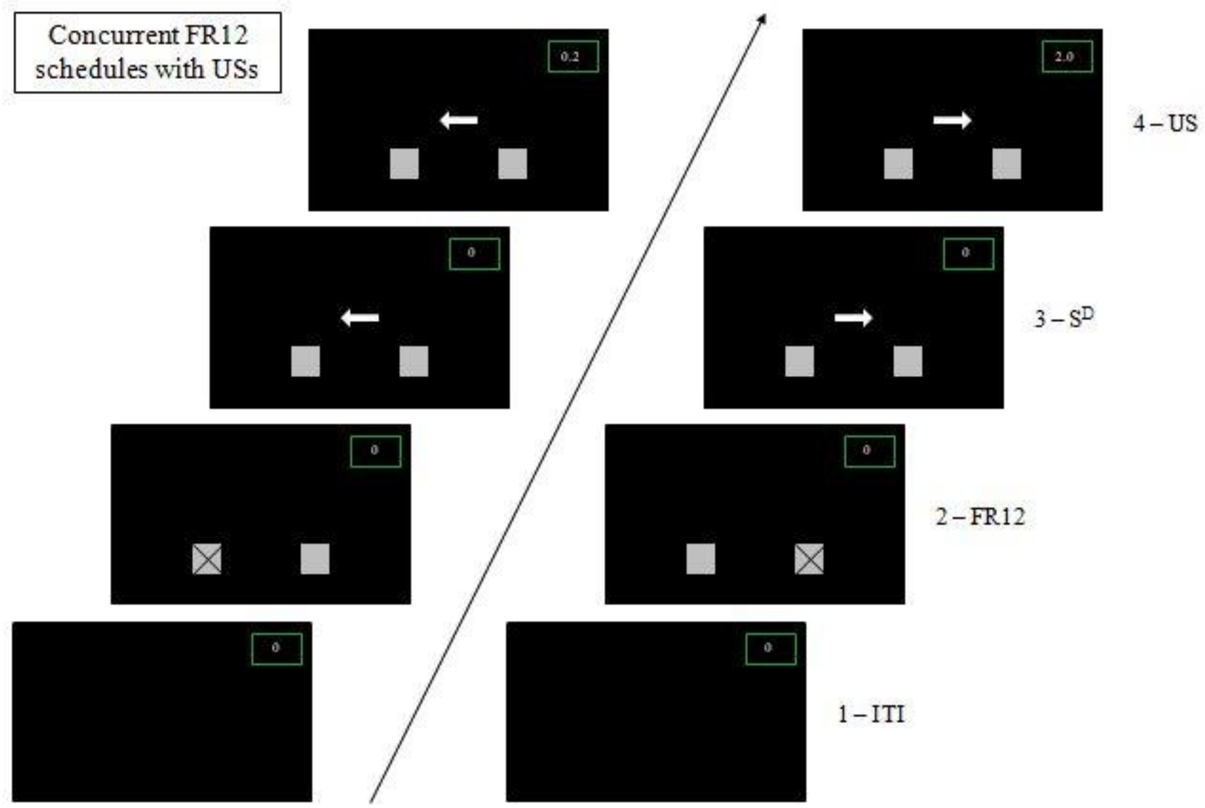

Figure 2. Concurrent FR12 schedules with US1 (i.e., +0.2 points) or US2 (i.e., +2.0 points) as US in Phase 1. The crosses signal the response key on which the FR12 is completed.

Prével, Rivière, Darcheville, Urcelay, & Miller – Figure 2

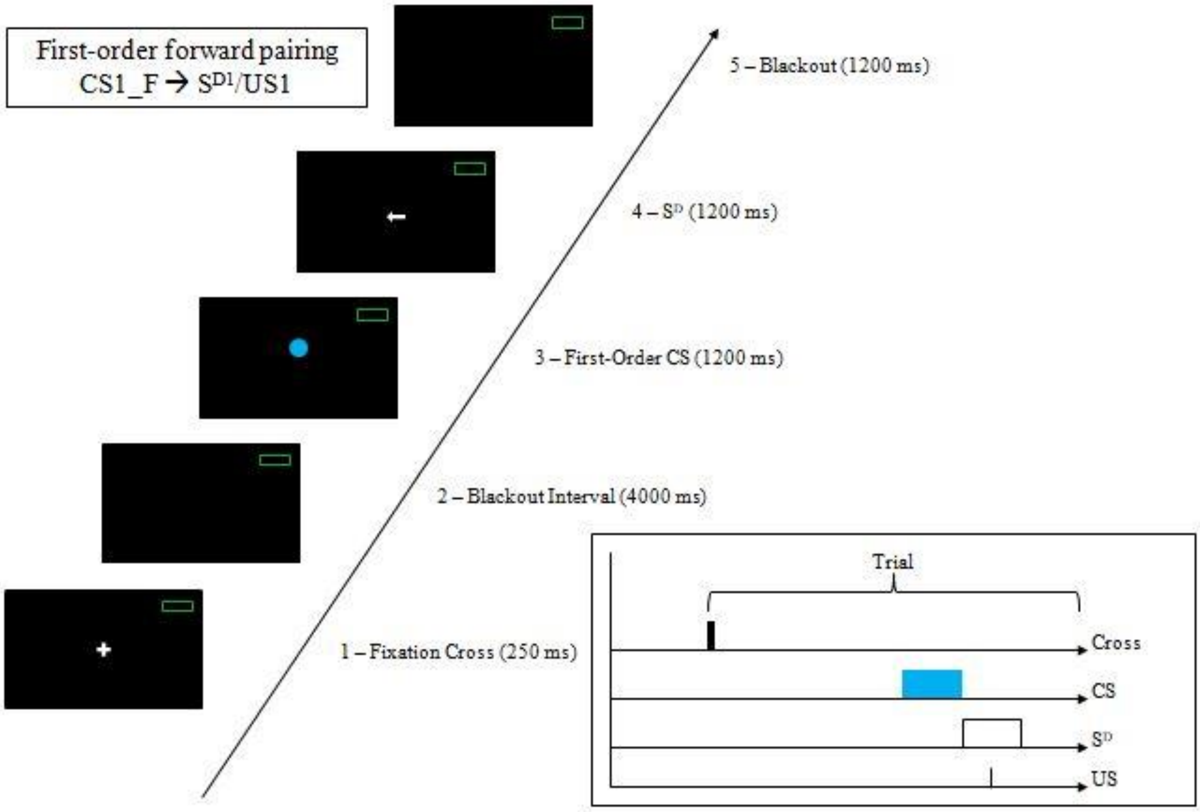

Figure 3. Illustration of a first-order conditioning trial conducted during Phase 2. There was no opportunity for clicking on the screen during Phase 2.

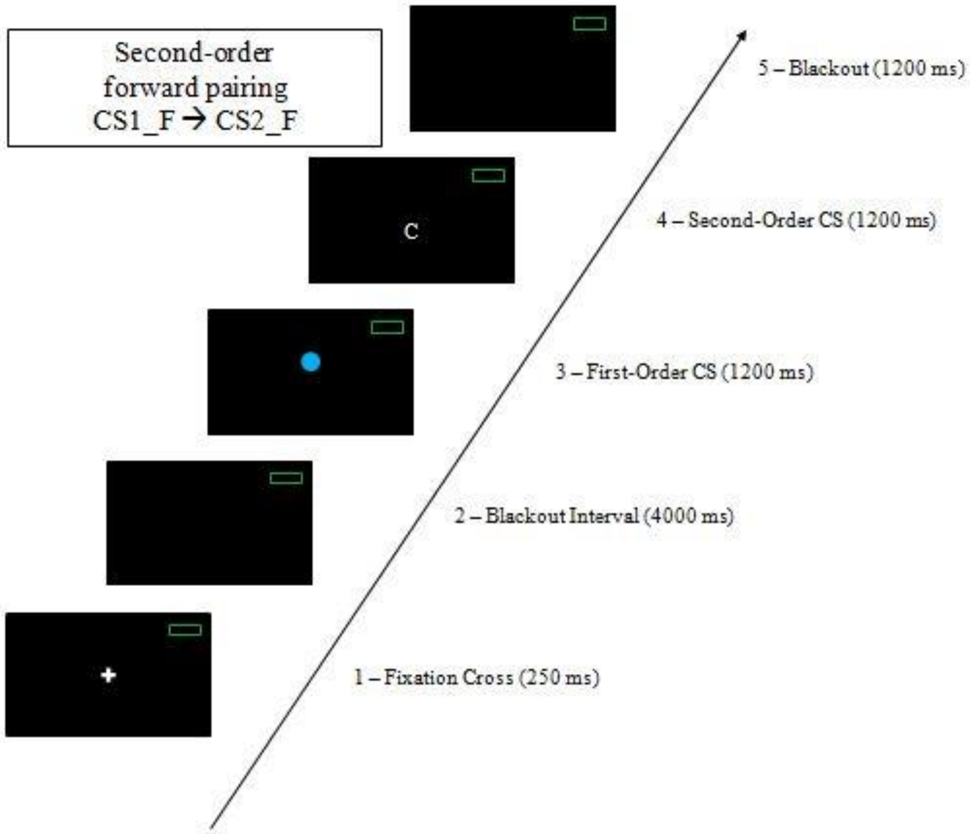

Figure 4. Illustration of a second-order conditioning trial conducted during Phase 4. There was no opportunity for clicking on the screen during Phase 4.

Prével, Rivière, Darcheville, Urcelay, & Miller – Figure 4

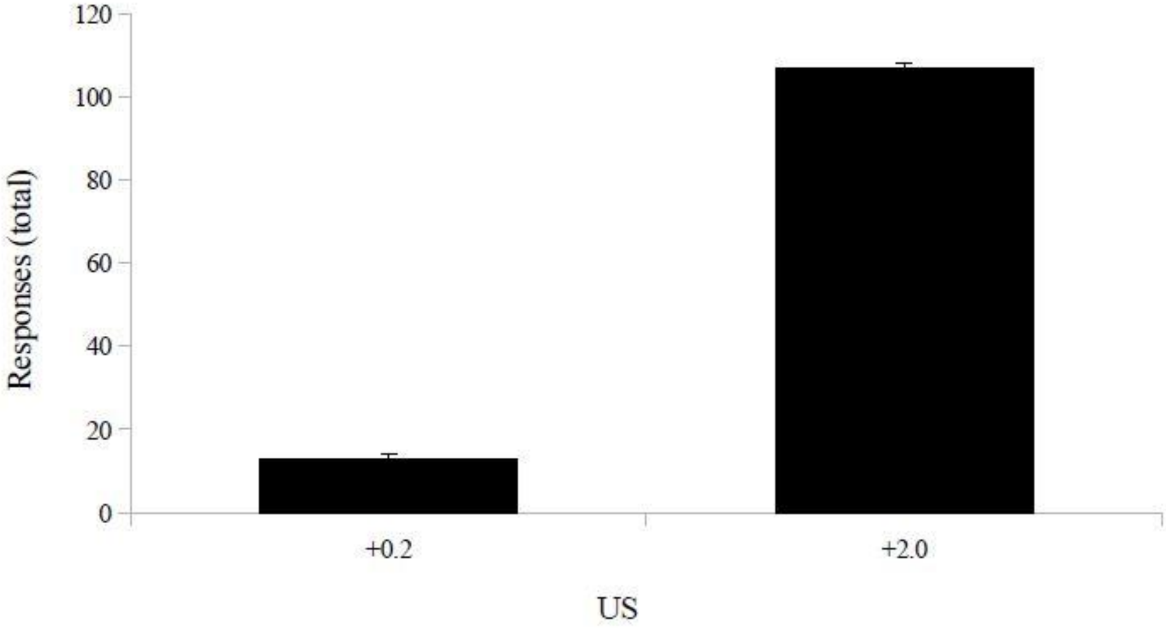

Figure 5. Mean number of responses on the keys delivering US1 (i.e., +0.2 points) and US2 (i.e., +2.0 points) during Phase 1. Error bars represent the SEM.

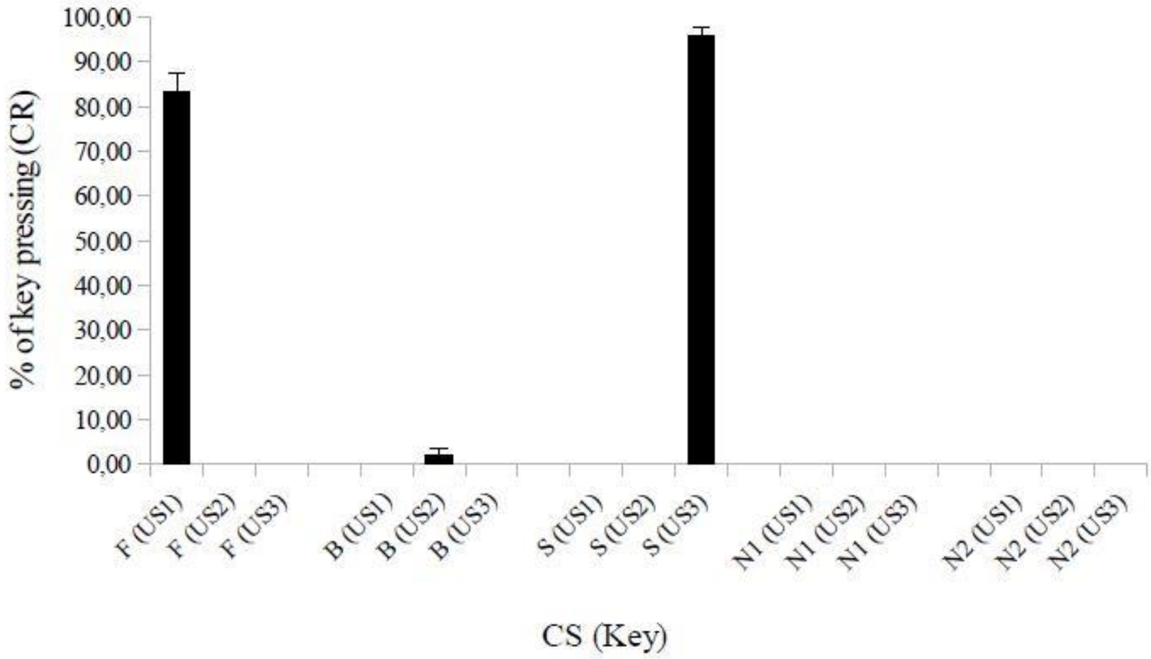

Figure 6. Mean percentage of CRs across keys associated with US1, US2, and US3, in the presence of the first-order CSs (i.e., F, B, S, N1, N2) during Phase 3. Error bars represent the SEM.

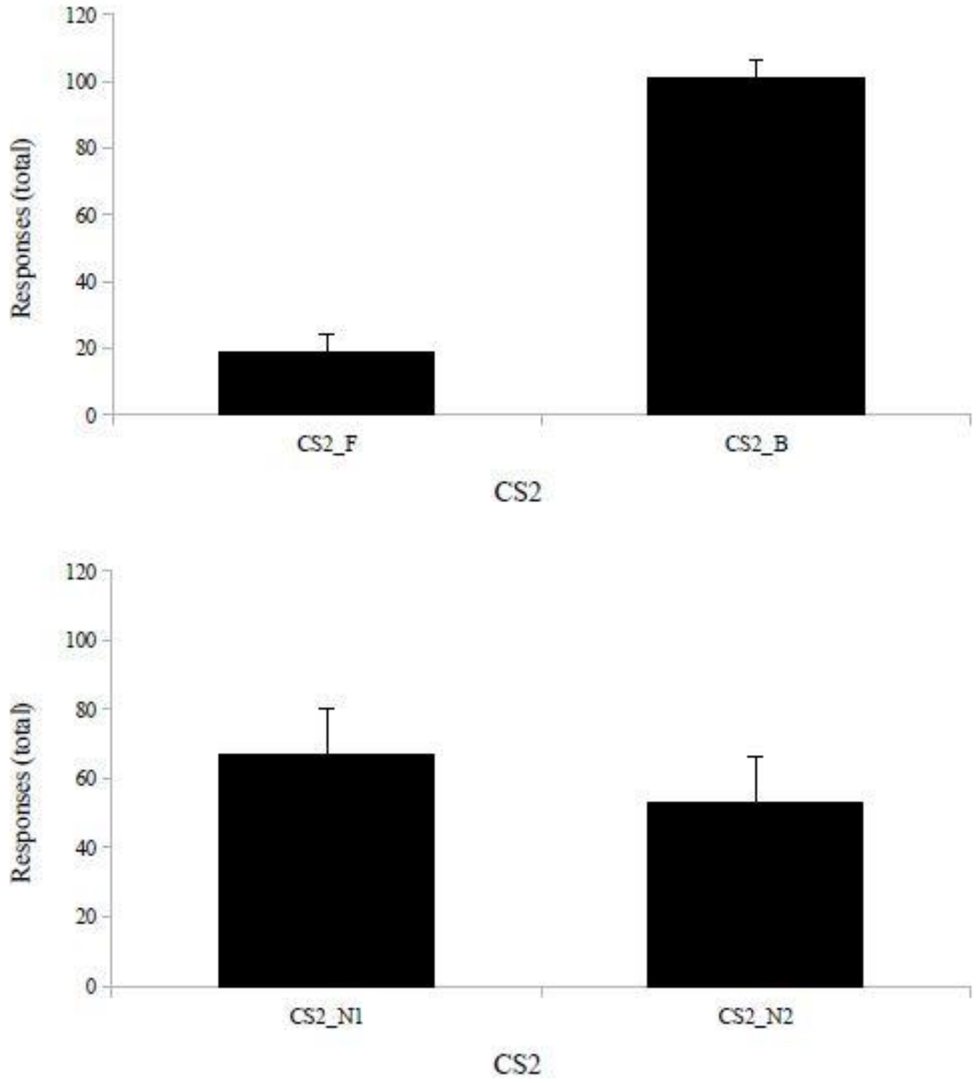

Figure 7. (Top) Mean number of responses on the key delivering CS2_F and CS2_B during Phase 5. (Bottom) Mean number of responses on the key delivering CS2_N1 and CS2_N2 during Phase 5. Error bars represent the SEM.

Prével, Rivière, Darcheville, Urcelay, & Miller – Figure 7