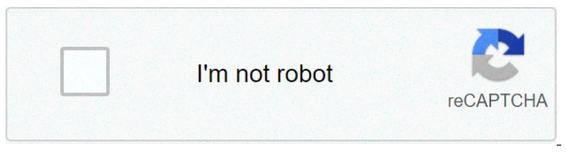

Continue

Second order conditioning

In classical conditioning, second-order conditioning or higher-order conditioning is a form of learning in which a stimulus is first made meaningful or consequential for an organism through an initial step of learning, and then that stimulus is used as a basis for learning about some new stimulus. For example, an animal might first learn to associate a bell with food (first-order conditioning), but then learn to associate a light with the bell (second-order conditioning). Honeybees show second-order conditioning during proboscis extension reflex conditioning.[1] Three phases in second-order conditioning In the SOC procedure, there are three phases. In the first training phase, a conditioned stimulus, (CS1) is followed by an unconditioned stimulus (US). In the second phase, a second-order conditioned stimulus (CS2) is presented along with CS1. Finally, in the test phase, CS2 is presented alone to the subjects while their responses are recorded.[2] Models of second-order conditioning Theoretical models for how second-order conditioning (SOC) works have a basis in associative learning theories. There are four broad models based on the associations formed during SOC. The first model suggests that the second-order stimulus (CS2) and the conditioned response (CR) form a direct link which is strengthened by the presence of the first-order stimulus (CS1). The second model suggests that in successful SOC an associative representation of each stimulus is created. The presentation of the CS2 would evoke a representation of the CS1, which would evoke a representation of the unconditioned stimulus (US), thus leading to the CR. The third model suggests a direct link between the CS2 and a representation of the US which leads to the CR. The fourth model suggests that the CS2 elicits a CR through a CS1 representation because a connection exists between the CS2 and the CS1 representation.[3] Second- Order conditioning helps explain why some people desire money to the point that they hoard it and value it even more than the objects it purchases. Money is initially used to purchase objects that produce gratifying outcomes, such as an expensive car. Although money is not directly associated with the thrill of a drive in a new sports car, though second- order conditioning, money can become linked with this type of desirable quality.[4] In fear conditioning It has been demonstrated in an associative fear conditioning chain, such as CS2 --> CS1 --> US, that extinction of freezing responses to the first-order stimulus (CS1) leads to responding impairments in CS2, but extinction of the second-order stimulus (CS2), does not have any effect on CS1 (Debiec et al.). In the same study, the effect of activation (memory retrieval) on such an associative chain has been examined. Results demonstrated that protein synthesis inhibition after exposure to a single CS1 impairs responses to both CS1 and CS2, but protein synthesis inhibition after exposure to a single CS2, only disrupts CS2 and leaves CS1 freezing intact. Therefore, it is believed that when the first-order association is directly activated, it is placed into a labile state (as we would expect from reconsolidation research) which may affect dependent associations. However, when the first-order association is only indirectly activated (through the associative chain), it appears that there is not sufficient stimulation to kick off cellular processes which would place it in a labile state, so it remains fixed.[5] References ^ Bitterman et al. 1983. Classical Conditioning of Proboscis Extension in Honeybees (Apis mellifera). J. Comp. Psych. 97: 107-119. ^ Jara, E., Vila, J., & Maldonado, A. (2006, August). Second-order conditioning of human causal learning. Learning and Motivation, 37(3), 230-246 . Retrieved from UTSC Library database. ^ Jara, Elvia; Vila, Javier; Maldonado, Antonio (2006). "Second-Order Conditioning of Human Causal Learning". Learning and Motivation. 37 (3): 230–246. doi:10.1016/j.lmot.2005.12.001. ^ Wegner, Daniel; Schacter, Gilbert (2011). Psychology (2nd ed.). p. 268. doi:10.1016/j.lmot.2012.12.002. ^ Debiec, J., Doyere, V., Nader, K., LeDoux, J.E. (February 28, 2006). Directly reactivated, but not indirectly reactivated, memories undergo reconsolidation in the amygdala. PNAS, Volume 103, Number 9, 3428-3433. Retrieved from " Open Access Peer-reviewed In human causal learning, excitatory and inhibitory learning effects can sometimes be found in the same paradigm by altering the learning conditions. This study aims to explore whether learning in the feature negative paradigm can be dissociated by emphasising speed over accuracy. In two causal learning experiments, participants were given a feature negative discrimination in which the outcome caused by one cue was prevented by the addition of another. Participants completed training trials either in a self-paced fashion with instructions emphasising accuracy, or under strict time constraints with instructions emphasising speed. Using summation tests in which the preventative cue was paired with another causal cue, participants in the accuracy groups correctly rated the preventative cue as if it reduced the probability of the outcome. However, participants in the speed groups rated the preventative cue as if it increased the probability of the outcome. In Experiment 1, both speed and accuracy groups later judged the same cue to be preventative in a reasoned inference task. Experiment 2 failed to find evidence of similar dissociations in retrospective revaluation (release from overshadowing vs. mediated extinction) or learning about a redundant cue (blocking vs. augmentation). However in the same experiment, the tendency for the accuracy group to show conditioned inhibition and the speed group to show second-order conditioning was consistent even across sub-sets of the speed and accuracy groups with equivalent accuracy in training, suggesting that second-order conditioning is not merely a consequence of poorer acquisition. This dissociation mirrors the trade-off between second-order conditioning and conditioned inhibition observed in animal conditioning when training is extended. Citation: Lee JC, Livesey EJ (2012) Second-Order Conditioning and Conditioned Inhibition: Influences of Speed versus Accuracy on Human Causal Learning. PLoS ONE 7(11): e49899. Reginald Frederick Westbrook, University of New South Wales, Australia Received: September 9, 2012; Accepted: October 17, 2012; Published: November 28, 2012Copyright: © 2012 Lee, Livesey. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.Funding: This research was supported by a University of Sydney Bridging Support Grant awarded to EJL. No external sources of funding contributed to this research. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.Competing interests: The authors have declared that no competing interests exist. In a typical human causal learning experiment, cues are presented that may increase or decrease the likelihood of a particular outcome and the participant's task is to assess to what degree each cue either causes or prevents that outcome. Various authors have suggested that this process involves elementary associative learning mechanisms because results from many of these experiments bear strong resemblance to animal conditioning phenomena (e.g. [1], [2]). Conditioned inhibition – or learning about a cue that has a negative contingency with an outcome – is one such example. Conditioned inhibition results from experience with a feature negative (FN) discrimination, where one cue leads to an outcome (A+), but when it is paired with a second cue, no outcome occurs (AX−). After sufficient training with these contingencies, the test stimulus (X) typically acquires inhibitory properties, such that its presence reduces responding in animal conditioning [3] or lowers ratings of causation or contingency in human learning [4]. In other words,×becomes a conditioned inhibitor as a consequence of its negative contingency with the outcome. When paired with another cue that has previously signaled the outcome (e.g. B+) the conditioned inhibitor reduces behavioral anticipation of the outcome that would normally be elicited by B (i.e. a summation test; [3]). In human causal learning, some doubt has been cast over several experiments that purport to show conditioned inhibition because of the choice of appropriate controls (see [5]). However, several experiments have found evidence of conditioned inhibition using a conservative test in which the ratings for the critical summation test compound BX are compared to ratings for a compound of B and a neutral or novel stimulus [5], [6]. In these studies, ratings for BX were substantially diminished, indicating that learning about×reduces causal ratings above and beyond what would be expected from a simple external inhibition effect; the reduction in ratings produced by pairing B with any other stimulus that has not been paired with the outcome [5]. Thus, like several other phenomena, conditioned inhibition appears to be common to a range of very different learning paradigms from Pavlovian conditioning to human causal judgment. The general conclusion that human judgments of causation have an associative basis has been challenged on several grounds, including parsimony [7]. Humans display cognitive abilities such as deductive reasoning (e.g. [8]) and rule abstraction [9], [10] that could succinctly explain many of the causal learning results without recourse to primitive learning mechanisms. The task of separating the contributions of associative learning from other forms of cognition is made difficult by the fact that most experimental results in causal reasoning and contingency judgement are consistent with multiple explanations. Conditioned inhibition, for instance, could be explained as the formation of an inhibitory link between the conditioned inhibitor×and the outcome, which negates excitatory associations between other cues and the outcome. Such explanations follow naturally from the mechanisms described in many associative learning models (e.g. [11]). However, alternatively one could interpret this as the participant forming an inference that cue×prevents the occurrence of the outcome [12]. These explanations are by no means mutually exclusive but both effectively account for the learned properties of the conditioned inhibitor. Given this general problem of dissociating psychological processes from one another, the FN paradigm is particularly interesting because under some circumstances, the cue (X) that possesses a negative contingency with the outcome actually appears to acquire excitatory rather than inhibitory properties [13], [14], [15]. This effect is often referred to as second-order conditioning because×acquires excitatory properties via its association with an excitatory cue (A) that is paired directly with the outcome. Several animal learning studies suggest that a transition from second-order conditioning to conditioned inhibition occurs through the course of training, with inhibition developing slowly. For example, Yin, Barnet, and Miller [15] have shown that conditioned inhibition manifests only with extended training with the FN contingencies intermixed, while second-order conditioning is evident with fewer training trials, presented either interspersed or in a blocked (A+ then AX−) design. Second-order conditioning is noteworthy because normative and inferential models predict that×should not be treated as a cause of the outcome, given its negative contingency (X never appears with the outcome). For this reason, the mere fact that second-order conditioning occurs is viewed as being potentially diagnostic of the psychological mechanisms involved in learning [16]. Evidence for this effect in human causal learning can be found in a study reported by Karazinov and Boakes [17], who found second-order conditioning by limiting participants' time to think on each trial. Each participant completed a causal scenario in which they played the role of a doctor attempting to discover which foods consumed by a fictitious patient were causing migraine headaches. Participants in one group completed the training phase of the experiment in a self-paced fashion (as is usually the case in causal learning tasks), whereas another group were limited to three seconds to respond on each training trial. Embedded amongst several other contingencies, the participants were given a FN discrimination (P+/PX−), where the addition of×to P prevented a migraine from occurring. However, instead of judging the test stimulus (X) to be preventative of the outcome, as did the self-paced group, in both experiments the paced group gave the test cue a higher causal rating than they did a non-causal control cue (M) trained in compound (LM−). Results from the typical summation tests – comparing×to M in compound with a trained excitor (T+) – suggested a similar pattern. Experiment 1 revealed a group interaction whereby TX was rated higher than TM in the paced group, but neither conditioned inhibition nor second-order conditioning was evident in the unpaced group. In Experiment 2, the unpaced group rated TX lower than TM (consistent with conditioned inhibition) but no group interaction was evident and the paced group did not rate TX higher than TM. Shanks ([16]; see also Mitchell et al., [7]) has recently cited this result as a compelling example of causal learning taking a form that defies any obvious explanation in terms of rational inference, suggesting instead the operation of associative processes in human causal learning. The result is particularly noteworthy because effects indicating excitatory and inhibitory learning were revealed with training on the same contingencies, albeit not within the same experiment. Other cue competition effects are known to be sensitive to the conditions of learning in a seemingly similar fashion. These include retrospective revaluation effects (e.g. mediated extinction versus release from overshadowing; [18]) and the evaluation of a redundant cue (e.g. blocking versus augmentation; [19]), which will be briefly discussed in relation to Experiment 2. However, by and large, studies rarely observe cue contingency effects of this nature occurring in both excitatory and inhibitory directions on the basis of a single manipulation. Karazinov and Boakes' [17] results constitute the best evidence for a non-rational second-order conditioning effect in human causal learning. However, even in their study, excitatory and inhibitory simple effects were not found in the same experiment. The potential significance of the effect and the somewhat equivocal nature of Karazinov and Boakes' result make it all the more important to replicate this dissociation and to examine its properties. The primary aim of this study was to garner further evidence for Karazinov and Boakes' [17] dissociation in the FN paradigm by varying additional training parameters in addition to their pacing manipulation, providing a stronger impetus to respond either as quickly or as accurately as possible. However, unlike Karazinov and Boakes, we wished to obtain the dissociation using an identical set of test stimuli to find effects consistent with conditioned inhibition and second-order conditioning. Both experiments used a between-subjects design to manipulate trial time (unpaced versus paced trials), accompanied by instructions and feedback that emphasized the importance of either accuracy or speed during learning. Participants given self-paced trials and instructions to be as accurate as possible were expected to show learning consistent with conditioned inhibition, as has been observed in similar causal learning tasks previously (e.g. [4]). Participants given trial time limits and instructions emphasizing speed were expected to show second-order conditioning, consistent with Karazinov and Boakes' [17] findings. In each experiment, participants assumed the role of a pharmaceutical researcher learning about the effects of different drugs that could cause potential side-effects. The cues were novel drug names (e.g. Slevoral, Melixil), and the possible outcomes were the occurrence of migraine (Experiments 1 and 2), nausea (Experiment 1 only), or no outcome. Experiment 1 focused on the feature negative contingencies in a complex causal learning task involving multiple outcomes. Experiment 2 examined the effect of trial time restriction on other cue contingency effects in addition to the FN discrimination. To test the claim that normative and inferential models do not predict second-order conditioning [17], an inference test in Experiment 1 aimed to show that conditioned inhibition was the rational judgement that should have resulted in the speed group. Experiment 1 primarily aimed to dissociate excitatory and inhibitory learning resulting from acquisition of the FN discrimination, using instructions, feedback and trial time limits to emphasise either speed or accuracy during training. In addition to the stimuli directly involved in the FN paradigm, other stimuli were included to assess transfer of learning and to function as filler cues (Table 1). The experiment used a scenario in which two possible side effects could occur as outcomes. Thus, each trial type was associated with "migraine", "nausea", or "no outcome". Each participant completed two sets of FN discrimination and related control trials, one set involving migraine as the potential outcome, the other involving nausea (see Table 1). After training, both groups were given a self-paced ratings test, in which they were shown drug cues (or combinations of cues) and had to indicate the degree to which they expected each of the two side-effects to occur. The ratings test yielded two kinds of scores: outcome-specific ratings (specifically using the rating for the associated outcome during training) and the ratings difference scores (the difference between the ratings for the associated outcome and the alternative outcome). For example, the outcome-specific score for A1 was the rating for outcome 1 only, and the difference score was obtained by subtracting the rating for outcome 2 from the rating for outcome 1. The difference scores were included as a means of gauging outcome specificity in learning, allowing for learning that "X causes/prevents O1" to be distinguished from the generalised learning of "X causes/prevents a side-effect", which would manifest as a change in ratings for both scales (e.g. see [20]). To assess learning, non-causal cues C1 and C2 were combined with trained excitors (B1 and B2) to form a novel control compound, which would then be compared with a novel compound consisting of the test cues (X1 and X2) and the same trained excitors (B1 and B2). Thus, the presence of conditioned inhibition or second-order conditioning was assessed via a summation test by comparing these critical test stimuli B1X1 and B2X2, to controls B1C1 and B2C2. If participants had genuinely learned that the test stimuli (X1 and X2) were inhibitors, they should rate the probability of their respective illnesses occurring as being low when they are paired with different excitors, compared to when the excitors are paired with the non-causal (but also non-preventative) control cues (C1 and C2). This was thought to be a conservative but necessary measure of conditioned inhibition, since it is known that combining a trained excitor with another stimulus results in lower predictive ratings due to reasons other than conditioned inhibition (see [5], [6], [21]). Since the aim was to obtain the group interaction on the same test cues, the choice of control cue was driven by the need to compare excitatory and inhibitory learning with an unambiguously non-causal cue. Conversely, a higher rating for BX than for BC indicates second-order conditioning has occurred as it suggests that the presence of×has an excitatory rather than an inhibitory relationship with the outcome. This is an atypical measure for second-order conditioning, which has conventionally involved testing individual stimuli. However, it is appropriate in this case for two reasons. First, both BX and BC are novel compounds and any effect on ratings generated by uncertainty about new combinations of drugs will affect both. Second, it provides a direct comparison with the evidence for conditioned inhibition. By any conventional analysis based on associative learning principles, the excitatory strength of B should not inflate ratings of BX any more than BC and thus if BX receives a higher rating than BC, it should be based on the participant's evaluation of×vs. C. Following from both the animal literature and Karazinov and Boakes' [17] results, it was expected that conditioned inhibition would be evident in the accuracy group. The question of most interest was whether this effect would interact with the group manipulation and, more specifically, whether second-order conditioning would occur in the speed group, where the opportunity to reflect on each trial is restricted. A self-paced inference test at the end of the experiment sought to clarify whether conditioned inhibition was considered a rational judgement, and specifically, whether the speed group would still show second-order conditioning when given the opportunity to reason about the contingencies. Fifty-two first-year psychology students from the University of Sydney participated in exchange for partial course credit. Five participants who scored below 35% (slightly above chance) accuracy for the feature negative stimuli (mean of A1, A2, A1X1 and A2X2) in the last quarter of the training phase were excluded, leaving 23 participants in the speed condition, and 24 in the accuracy condition (37 female, mean age=19.8 years). All participants gave written informed consent and the procedure was approved by the University of Sydney Human Research Ethics Committee. The experiment was programmed using Psych Toolbox for Matlab [22], [23] and run on Apple Mac Mini desktop computers connected to 17 inch CRT monitors, refreshed at a rate of 85 Hz. Participants made their responses using a standard Apple keyboard and mouse. Testing was conducted in individual cubicles in groups of up to five, with sound feedback delivered via personal headphones. In the training phase, participants were asked to assume the role of a pharmaceutical researcher whose job was to determine the effects of different drugs using trial and error. On each trial a drug or combination of drugs was presented and participants were asked to predict which of three possible outcomes they thought might occur (migraine, nausea or no outcome) by clicking on one of the buttons below the drug names. When an answer was selected, the box surrounding the outcome turned yellow, the three buttons disappeared and were replaced by the correct answer while the drug names remained on the screen. The drug names appeared in one of 3 colours (blue, green or red) and either a picture of a sad face or medicine was displayed on the feedback screen if the correct outcome was one of the illnesses. The choice of cue colour and picture was not systematically related to particular cues or outcomes. Participants in the accuracy group were told to do the task as accurately as they could and to take their time, receiving a buzzer tone and the word 'INCORRECT' on the top of the screen if they made an error, as well as the word "correct" in smaller font if they chose correctly. Participants in the speed group were told to complete the task as fast as they could and were given only 1.5 seconds to respond, after which a buzzer tone was heard and the word 'FASTER' appeared at the top of the screen and no response recorded. The speed group were not given any feedback as to whether they were correct or incorrect and were only shown the correct answer. All contingencies were consistent throughout and therefore each stimulus presentation fully predicted a particular outcome. There were 8 blocks of 24 trials presented continuously without break for the entire training phase (192 trials in total). Within each block there were 2 repetitions of the 12 trial types (see Table 1), with their order of appearance randomised within each block. The spatial presentation of stimuli within each compound was counterbalanced so equal numbers of each were seen (e.g. AX and XA). In the ratings test, participants were asked to rate the likelihood of each of the two outcomes occurring given the presence of one or two of the drug cues. On each trial, the drug name(s) appeared at the top of the screen, followed by two linear analogue scales appearing next to each of the outcome names (i.e. one scale for migraine, one for nausea). The end points of each scale were labelled "definitely will not occur" to "definitely will occur". Participants could click anywhere on the scale, yielding ratings ranging from 0–100. The order of presentation was randomised, with each single-cue stimulus presented once, and each compound twice, again with the order of presentation within each compound counterbalanced. The ratings test was self-paced. The last phase of the experiment (the inference test) aimed to extract a rational predictive judgement about the test stimuli by presenting all the relevant contingencies in the summation test at once on the screen. Participants were told that they would be viewing the results of the drugs again and could make another reasoned judgement which could be the same or different as before. Participants were shown that A1 led to outcome 1, A1X1 led to no outcome, B1 led to outcome 1 and C1 led to no outcome (A1+/A1X1−/B1+/C1−). They were then asked to rate how likely both outcomes 1 and 2 were to occur for the compounds B1X1 and B1C1 (the same compounds used in the summation test). These ratings were made in the same fashion as the predictive ratings, with all scores transformed to a scale of 0–100. This was then repeated for the corresponding stimuli with outcome 2 (A2+/A2X2−/B2+/C2−, test C2X2 and C2E2). All drug name allocations and drug-illness contingencies were the same as in training, with all writing presented in white on a black background. All analyses were performed with an alpha level of.05 and Greenhouse-Geisser adjusted p-values are reported where relevant. Figure 1 shows accuracy for each stimulus type across training, averaged in four equal blocks. Over all stimuli, the accuracy group were more accurate throughout all training blocks, lowest $F(1, 46)=5.09$, $p=.029$, and overall, $F(1, 46)=18.32$, p