IN PRESS: Journal of Experimental Psychology: Animal Learning & Cognition

The role of inhibition in the suboptimal choice task

Valeria V. González & Aaron P. Blaisdell

Author Notes

This research was conducted following the relevant ethics guidelines for research with animals and was approved by UCLA's institutional IACUC, and in compliance with the APA ethical standards in the treatment of animals.

Valeria V. González, Department of Psychology, University of California, Los Angeles

Aaron P. Blaisdell, Department of Psychology, University of California, Los Angeles

Valeria González was funded by the Portuguese Foundation for Science and Technology (FCT), Doctoral Grant PD/BD/114368/2016.

Correspondence concerning this article should be addressed to Aaron P. Blaisdell, UCLA Department of Psychology, 1285 Franz Hall, Los Angeles, CA 90095-1563. E-mail: blaisdell@psych.ucla.edu.

Abstract

Given a choice, pigeons prefer an initial-link stimulus that is followed by reliable signals that food will be delivered (S+) or not (S-) after a delay, over an alternative initial-link stimulus that is followed by unreliable signals of food; even when the former yields a lower overall probability of food. This suboptimal preference has been attributed to the combination of a biased attraction to the S+ and ignoring the S-. We evaluated the inhibitory properties of the S- in three experiments to investigate its role in suboptimal choice. In Experiment 1, pigeons were trained in an autoshaping procedure with the four terminal link stimuli of the suboptimal choice task; S+ was continuously reinforced, S3 and S4 were each partially reinforced on a 50% schedule, and S- was never reinforced. Summation tests showed that S- acquired inhibitory properties during training. Experiment 2 replicated the summation tests after training on the full suboptimal choice procedure and found that S- inhibition was not due to external inhibition. In Experiment 3, pigeons were trained on the suboptimal choice procedure and the development of inhibition was assessed throughout training. An analysis of individual differences across birds revealed that the response rates to S- were negatively correlated with the strength of suboptimal preference, both within subject as each bird acquired suboptimal preference, and across subjects once all birds had reached asymptotic levels of suboptimal preference. Thus, rather than ignoring the S-, we found evidence that birds attended to S- as an inhibitory stimulus. Future models explaining performance in the suboptimal choice task should consider inhibition to the S- as a factor in suboptimal choice.

Keywords: inhibition, suboptimal choice, response rate, pigeons

When hungry pigeons are confronted with a choice between two alternatives, one predicting a low probability of food but that is followed by predictable signals of food, and another predicting a higher probability of food but that is followed by unpredictable signals of food; they systematically choose the former. This preference has been called suboptimal because they fail to maximize food intake. For instance, Stagner and Zentall (2010) gave pigeons a choice between two initial-link options in a concurrent chain. If pigeons chose the suboptimal alternative, on 20% of the trials the choice led to a terminal-link stimulus (S+) that was always followed by food after 10s, whereas on the other 80% of the trials another terminal-link stimulus (S-) appeared for 10s ending always without food. If the optimal alternative was chosen, one of two terminal-link stimuli (S3 and S4) would appear, and after 10s was followed by food on half of the trials regardless of which terminal-link stimulus had been presented. Pigeons showed a strong preference for the suboptimal option, even though the overall probability of food for the optimal alternative was 2.5 times richer than for the suboptimal alternative (50% vs 20%, respectively).

This procedure was based on the observing response paradigm. An observing response is a response that produces a discriminative stimulus but that does not alter the probability of reinforcement (Browne & Dinsmoor, 1974). In his seminal work, Wyckoff (1952) found that pigeons were willing to press a pedal that produced stimuli- a red and a green light- *informing* which schedule was in effect in a given trial, despite that pressing the pedal did not change the current schedule of reinforcement. One of the earliest explanations for this finding emphasized the value of 'information', for instance knowing that when a red light is presented, a sugar pellet will be delivered. This idea, derived from information theory (Shannon & Weaver, 1949), holds that preference for information about the presence and absence of food should be equally valuable. For classic information theory, information can be understood as a resolution of uncertainty, which follows an inverse U curve from 0 to 1, in which 0 and 1 are the maximal reduction of uncertainty, and .5 is the point of minimum information or maximum uncertainty. The idea that information about the presence of food is valuable is straightforward, but information about the absence

of food can also be valuable because it allows the organism to devote resources elsewhere (Vasconcelos et al., 2015).

However, further research has found that information about the presence or absence of reinforcement is not equivalent (Dinsmoor, 1983). Indeed, it seems that animals prefer information yielding 'good news' (e.g., information about reinforcement) over information yielding 'bad news' (e.g., information about the absence of reinforcement). Although interesting, these experiments do not indicate whether animals learn that a "bad news" stimulus signals the absence of reinforcement, or instead that they learn to ignore the cue. Support for the account that "bad news" stimuli are learned about and not ignored comes from research showing that humans prefer bad news over no news (Fantino & Silberberg, 2010; Lieberman, Cathro, Nichol, & Watson, 1997), and that rats prefer a signaled over an unsignaled inescapable shock (Fanselow, 1980; Lockard, 1963; Miller, Marlin, & Berk, 1977).

Most studies involving predictors of food and no food have focused on the variables influencing the emission of the observing response rather than on the preference for informative of stimuli. Prokasy (1956) evaluated preference by placing rats in the middle of an E-maze. Rats could initially choose the left or the right arm. After the choice, the rat had to wait 30 s before gaining access to the goal box, which could be either baited or empty. The overall probability of reinforcement was equal in both arms, but one arm gave information: the walls were white when baited and black when not baited; the other arm had the same black and white walls on separate trials, but uncorrelated with the outcome. Under these contingencies, animals preferred the informative arm. It is important to note that the rats were not producing an observing response, but rather choosing the alternative that gave information about whether or not food would be delivered. This preference for the informative option was the beginning of the interesting question about the role of information in decision-making (Bower, Mclean, & Meacham, 1966; Roper & Zentall, 1999). Later, Gipson et al. (2009) and Stagner and Zentall (2010), found that the preference for the informative alternative occurred even when choosing that alternative also meant choosing less food–a suboptimal preference.

Most theoretical accounts of suboptimal preference assume that the preference is due to the combination of a biased attraction to the S+ and ignoring the S- (Zentall, 2016), with the result that the S- does not contribute to choice behavior between the initial link stimuli. For instance, the Reinforcement Rate Model (RRM; Vasconcelos, Monteiro, & Kacelnik, 2015), based on optimal foraging theory and consistent with the literature on the observing response, proposes that animals follow an information-seeking strategy, in which the suboptimal preference is due to the information embedded in the S+. The model also assumes that, because in nature animals are not attracted to stimuli that signal the absence of reinforcement, the S- plays no role in the decision process in the suboptimal choice task. (This contrasts with the observing response literature, where the signal for no-outcome is preferred over no signal.) Similarly, Cunningham and Shahan (2018) advanced an explanation based on the idea that animals learn the temporal relations between events. Applied to the suboptimal choice task, animals are choosing information about when reinforcement is delivered, and therefore also assumes that the S- is ignored because there is no reinforcer for a temporal relation to signal. (For other proposals that disregard the role of the S- see McDevitt, Dunn, Spetch, & Ludvig, 2016; Zentall, 2016.) By contrast, the Delta-Sigma Model (Gonzalez et al., 2020) proposes that two higher-order variables are responsible for suboptimal preference; a) the difference in the probability of reinforcement within the terminal links (Delta), and b) the ratio between the probability of reinforcement of each alternative initial link stimulus (Sigma). It assumes that animals need to pay attention to all probabilities, including a probability of 0 signaled by the S-, for behavior to become suboptimal. In summary, it is not clear if animals completely ignore the S-, or if the S- does affect suboptimal choice. If it does affect suboptimal choice performance, it is also not clear whether it does so through a perceptual, attentional, or learning process.

We hypothesize that the S- does contribute to suboptimal preference. Furthermore, we propose that it does so through the development of conditioned inhibition to the S-. That is, we propose that the S- becomes a Pavlovian conditioned inhibitor, and that its status as a conditioned inhibitor is related to the development of suboptimal preference in the suboptimal choice procedure. What is the evidence that the

S- terminal link stimulus becomes a conditioned inhibitor during the suboptimal choice task, and that its inhibitory status is related to suboptimal preference? A study showed that the S- developed inhibitory properties at the beginning of training, but that inhibition diminished with extended training (Laude et al., 2014). They used a variant of the typical suboptimal choice task in which the magnitude of the reinforcer following each stimulus, rather than the probability of reinforcement of each stimulus differed between alternatives (See Figure 1). Specifically, the S+ and S- terminal link stimuli that followed the suboptimal initial link stimulus led to 10 and 0 pellets, respectively; whereas S3 and S4 terminal link stimuli that followed the optimal initial link stimulus always resulted in 3 pellets. Thus, the overall average magnitude of reinforcement was 2 pellets for the suboptimal choice and 3 pellets for the optimal choice. Pigeons were trained in this task using color keys for the S+, S3, and S4 terminal links, and a vertical line for the S-. To assess the inhibitory properties of the vertical line S-, the rate of response to the compound S+S- was compared to S+ alone early and late in training using a within-subject design (Experiment 1) and between-groups design (Experiment 2). They found a significant reduction in inhibition as a function of extended training between the early and late tests in both experiments. This is a surprising result because most studies show that inhibitory properties accrue gradually during discrimination training (e.g., Thomas & Basbaum, 1972; see also Domjan, 2014), just the opposite of the findings of Laude et al. Moreover, they do not provide a theoretical reason why they expected the S- to start out as an inhibitory stimulus early in training, but to wane with extended training. It is possible that the S- initially reduced responding to the compound due to external inhibition, which waned as the S- became more familiar.

Despite the large effects in both experiments, the research designs have some peculiarities that warrant caution in the interpretation of the results of Laude et al. (2014). First, they varied the typical procedure of the task: each colored terminal link stimulus was followed by reinforcement 100% of the time. This is relevant especially for explanations that assumed animals are following information (Cunningham & Shahan, 2018; González, et al. 2020; Vasconcelos et al., 2015). In the procedure used by Laude et al. (2014), all color keys were informative. Until now, it is uncertain whether the same

mechanisms for the development of suboptimal preference underlie both types of procedure (Daniels & Sanabria, 2018). Second, the S- stimulus was a vertical line, whereas the remaining stimuli used for the initial and terminal links were colored keys. Color has been shown to be a more salient dimension to pigeons than stimulus shape (e.g., Blaisdell & Cook, 2005). Thus, failure to counterbalance the stimulus roles as S+ and S-, and choice of the less salient visual dimension for the S- may have contributed to the loss of stimulus control by the line orientation stimulus in their study. Third, the only assessment of inhibition to the S- was by measuring the difference in peck rate to the S+ versus the S+S- compound. Additional tests involving other stimulus compounds of the S+ and another stimulus (e.g., a novel stimulus or a different excitatory stimulus) are required to rule out external inhibition as an explanation for the reduction in response on compound test trials (Rescorla, 1969).

Two studies failed to find evidence that the S- plays any role in the decision process. In a conventional suboptimal choice task, Fortes, Vasconcelos, and Machado (2016) manipulated the probability of occurrence and the duration of the S-, finding that pigeons continued showing a suboptimal preference even when the S- was presented on 95% of the trials, or the delay of this stimulus was increased to 200 s. Nevertheless, neither of these manipulations would be expected to reduce the inhibitory properties of the S-, and on the contrary, they should strengthen them. An experiment using rats trained on a conventional suboptimal choice task found that inhibition to the S- increased with training (Trujano, López, Rojas-Leguizamón, & Orduña, 2016). Rats did not, however, develop a suboptimal preference. The authors concluded that the strong difference found between pigeons and rats in the suboptimal choice task is related to differences in the impact of conditioned inhibitors, such task that promote the development of conditioned inhibition do not allow for the development of suboptimal preference.

The current experiments had two goals. The first goal was to assess if the S- acquired inhibitory properties in a conventional procedure for suboptimal choice in pigeons, unlike the procedure used by Laude et al. (2014) which deviated in a number of ways from the conventional procedure (see above).

Thus, perhaps the conflicting results of Laude et al. (2014) could be attributed to external inhibition during initial summation tests of inhibition, with external inhibition waning with extended training. Likewise, Trujano et al. (2016) claimed that because rats showed evidence of Pavlovian inhibition to the S-, rats do not develop a suboptimal preference. Contrary to their position, recent evidence has found that rats choose suboptimally when the delays to reinforcement are extended (Cunningham & Shahan, 2018), or when the nature of the terminal link stimuli was changed from lights as used by Trujano et al. (2016) to tones used by Ojeda, Murphy, and Kacelnik (2018).

In Experiment 1, pigeons were trained in a Pavlovian autoshaping procedure involving only the terminal link stimuli (S+, S-, S3, and S4). Pigeons received 30-s presentations of each stimulus on separate trials. Each stimulus signaled a specific probability of reinforcement: p(food|S+)= 1; p(food|S-)= 0; p(food|S3)= .5; p(food|S4)= .5. Occasional nonreinforced probe trials with elements and compounds S+S-, S3S4, S3S- and S+S4 were delivered to assess the inhibitory properties of S-. If the S- had acquired inhibitory properties, we would expect a reduction in the responses to S+S- compared to S+ alone or to the S3S4 compound. Though not critical to our central hypothesis, we also included summation tests of compounds S3S- and S+S4 to compare the excitatory properties of a partially reinforced stimulus to a continuously reinforced stimulus. If the Pavlovian contingencies signaled by the terminal link stimuli are learned, the S- should develop into a conditioned inhibitor as a function of amount of training.
Experiment 2 extended the results of Experiment 1 by assessing inhibition to the S- after training on the suboptimal choice procedure. Pigeons were trained in the suboptimal choice procedure until stability of suboptimal preference was reached. After a reliable preference was acquired, they received a test session to assess Pavlovian inhibition using a summation test. During testing, animals were presented with probe trials of an S+S- compound but also S+ in compound with a new stimulus, to rule out external inhibition. If S- acquire inhibitory properties due to conditioned inhibition, we should observe a lower response rate to the compound of S+ and S- than to the compound of S+ and the non-trained stimulus.

A second aim was to explore the relationship between the development of inhibition to the S- on the one hand, and the strength of suboptimal preference on the other. To assess this, in Experiment 3, pigeons were trained on the typical suboptimal choice task involving both initial and terminal link stimuli. We measured the development of conditioned inhibition to the S- using summation tests with compound stimuli as in Experiments 1 and 2, and the development of preference for the suboptimal initial link stimulus (i.e., suboptimal preference) on choice trials during training. One hypothesis is that conditioned inhibition to the S- plays a causal role in the development of suboptimal preference, or at least is related to the causal process for the development of suboptimal preference. An alternative hypothesis is that conditioned inhibition to the S- plays neither a direct nor indirect causal role in the development of suboptimal preference. We took advantage of the large individual differences typical of pigeon experiments using choice procedures to find evidence to support one or the other hypothesis. If the development of suboptimal preference is directly or indirectly related to the development of conditioned inhibition to the S-, then we predict the strength of suboptimal preference and of conditioned inhibition to be positively correlated, such that individuals that show stronger conditioned inhibition will show stronger suboptimal preference. If the strength of suboptimal choice is uncorrelated with the strength of conditioned inhibition to the S-, this would fail to support the first hypothesis, and instead would suggest the two are not causally related, either directly or indirectly through other causal pathways. Finally, if suboptimal choice depends on ignoring the S-, as argued by Laude et al. (2014) as well as by many of the models discussed in the Introduction, then the strength of suboptimal preference should negatively correlate with the strength of conditioned inhibition to the S-. Tests of conditioned inhibition will use the negative summation test, in which the excitatory S+ and the putative inhibitory S- will be presented in compound and responses to these compound cue trials will be compared to responses on nonreinforced presentations of the excitatory S+ alone. Evidence for conditioned inhibition of S- would be shown by a significant reduction in responding on compound S+S- trials compared to S+ alone trials at test. If inhibition does not decay during training and is positively associated with the development of suboptimal preference, this result will encourage the incorporation of mechanisms of inhibition to the S- into models

of suboptimal choice. Although it is possible that learning inhibition to S- is independent from the development of a suboptimal preference, current and new models would need to integrate an explanation of how learning about a stimulus is not considered at the moment of choice.

Experiment 1

Before assessing the relationship between inhibition to the S- and suboptimal preference, we wished to determine whether terminal link stimuli, if trained on their own, acquire excitatory and inhibitory properties. Thus, pigeons received training with each of the terminal link stimulus elements (S+, S-, S3, and S4) on separate trials in a Pavlovian autoshaping procedure. S+ was always followed by food, S3 and S4 were followed by food on 50% of the trials in each session, and the S- was never followed by food. We measured the peck response rate to each stimulus. The inhibitory properties of the S- was evaluated using a summation test, in which the S- was presented in compound with a stimulus with excitatory properties (S+S- and S3S-). The response rate to the compounds was compared to the response rate to the elements and to other compounds such as S3S4 and S+S4. If S- acquired inhibitory properties as a function of training, we predicted a reduction in response rate to the compounds S+S- and S3S- compared to S+, S3S4, and S+S4.

Subjects

Five adult homing pigeons (*Columba livia*) from Double T farms, three males and two females, served as subjects. These pigeons had previously participated in a wide variety of behavioral experiments, including spatial overshadowing, object location encoding, response variability, and pattern learning, but were naïve with respect to the current procedures and stimuli which were selected to minimize transfer from prior experience. Subjects were individually housed in steel home cages with metal wire mesh floors in a vivarium. They were maintained at 80% of their free-feeding weight, but were given free access to water and grit while in their home cages. Testing occurred at approximately the midpoint of the light portion of the 12-hour light-dark cycle.

Materials and Methods

Apparatus

The experiment was conducted in a flat-black Plexiglas chamber (38 cm wide x 36 cm deep x 38 cm high). All stimuli were presented by computer on a color LCD monitor (NEC MultiSync LCD1550M). Stimuli were presented using the coding language Python (Python Software Foundation, https://www.python.org/) and the extension PsychoPy (Peirce, 2007). The bottom edge of the viewing window was 13 cm above the chamber floor. Pecks to the monitor were detected by an infrared touchscreen (Carroll Touch, Elotouch Systems, Fremont, CA) mounted on the front panel. A custom-built food hopper (Pololu, Robotics and Electronics, Las Vegas, NV) was located in the center of the front panel, its access hole flush with the floor. The hopper could deliver 3-s access to mixed grain as a food reward. All experimental events were controlled and recorded with a Pentium III-class computer (Dell, Austin, TX). A video card controlled the monitor in the SVGA graphics mode (800 x 600 pixels).

Stimuli

Each of the four stimuli, S+, S-, S3, and S4, were composed of two circles (either red, green, yellow, or blue) vertically (S+ and S3) or horizontally (S- and S4) aligned (Figure 2), with a size of 100 x 100 pixels. All the stimuli were presented in the center of the screen against a grey background.

Procedure

Pre-training. Pigeons were initially trained to peck each stimulus. A session consisted of 40 trials with an average duration of 10 minutes. Each stimulus was presented 10 times in random order within each session. On each trial, one pseudorandomly selected stimulus from the set of four was presented at the center of the grey screen. The stimulus remained on the screen until the required number of pecks was completed. Each trial ended with the delivery of the food reward followed by a 10-s Intertrial-Interval (ITI) during which the screen was black. Pigeons received two sessions of pre-training, the first under a

continuous schedule of reinforcement (CRF), the second under a fixed-ratio (FR) 10 schedule of reinforcement.

Training. Pigeons received a total of 25 daily sessions of training on an autoshaping procedure, six days per week. Each session consisted of 80 trials. In each session, the S+ was presented on eight trials each ending in food reward, the S- was presented 32 times and was never followed by food, S3 was presented eight times, four trials followed by food reward and four nonrewarded, and S4 was presented on 32 trials, in which 16 were followed by food reward and the remainder were nonrewarded. On each trial, a stimulus appeared in the center of the screen for either 10 s (during the first 10 sessions), or 30 s (during the remainder of training). The order of the trials was randomized. An ITI of 10 s with a black screen separated trials. Stimulus role was pseudorandomly assigned across subjects, in a way that any given combination did not repeat for another subject. During the first 10 sessions of training we observed low rates of pecking by some pigeons, and thus increased presentation duration to 30 s to allow more time to accumulate pecks. Pigeons received 15 sessions of training with stimuli of 30-s duration before experiencing the first test.

Test. Following the 15 training sessions with the longer duration stimuli, pigeons received test sessions in blocks of five sessions interspersed with blocks of 5 training-only sessions (e.g., sessions 16-20 for testing, sessions 21-25 for training, sessions 26-30 for testing, etc.). Pigeons received four blocks of five test sessions. Each test session of 56 trials started with ten presentations of the training stimuli, those were randomly selected from the pool with the same probabilities used during training. After the 10^{th} trial of each test session, 9 nonreinforced test trials were randomly interspersed among training trials for the remainder of the session. Test trials consisted of 2 presentations each of the following compounds: S+S-, S3S4, S3S-, and S+S4 (Figure 1, right panel), and one nonreinforced presentation of the S+. Each test trial was 30 s in duration and, as with training trials, the number of pecks during each test trial were recorded.

Data analysis

The peck rate for each stimulus was calculated for each session. Session peck rates were calculated for each stimulus and then averaged across blocks of five sessions. For the training phase, we analyzed the sessions with the 30 s duration, obtaining six blocks in which the first three corresponded to the training before the testing phase begun, and the following three corresponded to the blocks alternating test sessions. For the Test phase, we averaged pecks for all stimuli in each session in four blocks of five sessions. Given that the rate of response (RR) was highly variable between pigeons, the data were normalized. For each block, each stimulus RR was divided by the total RR across stimuli (S_{target} RR/Total RR).

Repeated-Measures (RM) Factorial Analysis of Variance (ANOVA) were implemented to analyze the blocks of training and test, using Holm corrections for the post-hoc analysis when main effects were found.

Results

The upper panel of Figure 3 shows the normalized mean peck rate for each stimulus as a function of training blocks. S+, S3, and S4 maintained similar RRs, the RR for S- was lower than the RR of any other stimuli. Supporting these observations, a RM Factorial ANOVA conducted on normalized peck rates with Block and Stimulus as repeated measures factors revealed a main effect of Stimulus, $F(3, 12) = 19.641$, $MSE = 0.031$ $p < .001$, $\eta^2 = .675$, but no effect of Block, $F(5, 20) = 0.053$, $MSE = 0.0002$, $p = .998$, $\eta^2 = .002$ 1.0, nor interaction, $F(15, 60) = 1.289$, $MSE = 0.006$, $p = .238$, $\eta^2 = .046$. Post-hoc analyses for the Stimulus factor using Holm correction indicated that S- RR was lower than S+, S3, and S4 RR ($p_{Holm} < .002$), whereas RRs across the remaining stimuli did not differ.

The bottom panel of Figure 3 shows the normalized response rate across test blocks for each stimulus. A RM factorial ANOVA with Block and Stimulus as within-subjects factors found a significant effect of Block [$F(3, 12) = 10.901$, $MSE = 0.002$, $p < .001$, $\eta^2 = .083$] and Stimulus [$F(3, 12) = 29.409$, $MSE = 0.005$, $p < .001$, $\eta^2 = .623$] but no interaction [$F(9, 36) = 1.697$, $MSE = 0.002$, $p = .126$, $\eta^2 = .053$].

Post-hoc analysis for Block found that Block 1 was different of Block 4 ($p_{holm} < .001$) and Block 2 from Block 4 ($p_{holm} = .043$). The analysis of the different stimuli indicated that the RR of S+ was significantly different from S- ($p_{holm} < .001$), S3 ($p_{holm} = .025$) and S4 ($p_{holm} = .043$), and that RR of S- was also different from S3 ($p_{holm} < .001$) and S4 ($p_{holm} < .001$); this last result replicates what was found during training blocks.

Figure 4 shows the mean normalized RR across all blocks of testing for each probe stimulus. We observed higher responses for all stimuli that did not include S- by itself or in compound with another stimulus (S+S-, S3S-). A RM factorial ANOVA conducted on normalized response rates for all stimuli in the test phase, with Block and Stimulus as factors, found a main effect of Stimulus, $F(7, 28) = 12.979$, $MSE = 0.008$, $p < .001$, $\eta^2 = .503$ but no effect of Block, $F(3, 12) = 1.077$, $MSE = 1.45e^{-19}$, $p = .395$, $\eta^2 = 3.292e^{-19}$, nor interaction, $F(21, 84) = 2.358$, $MSE = 0.004$, $p = .127$, $\eta^2 = .127$ (see Figure 4). Post-hoc analysis for the Stimulus factor using Holm correction found significantly lower RR to S- than to the other three training stimuli ($p <. 001$), evidencing that pigeons did not peck the stimulus associated with absence of food. No differences were found between S+, S3, and S4, as observed during training. Of interest to evaluate inhibition to the S-, we found that S+ RR was significantly higher than the RR to the S+S- compound ($p <. 001$) and to the S3S- compound ($p <. 001$). RRs to S+ and S+S4 did not differ ($p =.731$), this is an important comparison as it shows that merely presenting two stimuli together in a novel compound did not produce external inhibition or generalization decrement.

We predicted RR to compound test stimuli to show a linear relationship to the average probability of reinforcement signaled by the elements, with highest RR to the S+S4 compound signaling an average probability of reinforcement of .75 (i.e. p(food|S+)= 1.0 and (food|S4)= .5), followed by S3S4 with an overall probability of reinforcement of .5 (i.e. p(food|S3)= .5 and (food|S4)= .5), then S+S- (also overall probability of reinforcement of .5, but with S- serving as a conditioned inhibitor), and finally S3S- (overall probability of .25; from (food|S3)= .5 and p(food|S-)= 0, with S- serving as a conditioned inhibitor). As anticipated, a within-subjects contrast for the compound stimuli revealed a significant linear

trend, c= -.113, $t(12)$ = -4.766, $p < .001$. If RR to each compound reflected only the average probability of reinforcement signaled by the elements, we would predict equivalent RRs to the S3S4 compound and the S+S- compound, given that they both signal an average probability of reinforcement of .5. Given that the RR to the S+S- compound was significantly lower than to the S3S4 compound, this suggests the additional operation of conditioned inhibition to S- as a major factor in determining RR to the compound. That is, S- signals a 0 probability of reward, that translates in a negative rather than a neutral value for that stimulus, or increased probability of reward omission. It is in this way that conditioned inhibitors exert their modulating effect over conditioned excitors, such as to reduce the response elicited by the CS and by withdrawing from or avoiding the inhibitory CS (Hearst, Bottjer, & Walker, 1980; Wasserman et al., 1974).

Experiment 2

In the previous experiment, we used a Pavlovian autoshaping procedure, in which only the terminal links of the suboptimal choice task were presented, maintaining the frequency with which each stimulus is presented in the typical suboptimal choice task. The objective was to assess inhibition to the S- in a summation test. The results suggested that the S- passed a summation test of conditioned inhibition, thereby establishing that the S- acquired inhibitory properties during training. Interestingly, these results also suggest that the pigeons were not ignoring the S-. If so, the S- would not have affected responding controlled by the excitatory stimuli on the compound test trials. Indeed, the negative summation test is considered one of the preferred tests of conditioned inhibition precisely because it rules out inattention to the inhibitor as an alternative explanation (Cole, Barnet, & Miller, 1997; Rescorla, 1969). There were two goals for Experiment 2. First, we wished to replicate the summation test for conditioned inhibition after pigeons had received training on the full suboptimal choice procedure. Second, we included a control test of a compound of S+ and a novel stimulus to control for external inhibition. External inhibition is the negative summation effect that has sometimes been observed to be produced by an untrained cue presented in compound with a trained excitor (e.g., S+; Pavlov, 1927). In

the previous experiment we used S4, which was not trained with S+ but it had a history of reinforcement that could influence the response rate observed during test sessions. By contrast, if responses to the compound of the S+ and the untrained cue fails to produce negative summation in our test procedure, we can rule out external inhibition as an explanation for the negative summation observed on the S+S- test compound trials, thereby supporting true conditioned inhibition to the S-.

Subjects

Eight adult homing pigeons (*Columba livia*) from Double T farms were used. The pigeons had experience with an intelligence battery task but were naïve with respect to the current procedures and stimuli which were selected to minimize transfer from prior experience. Subjects were individually housed in steel home cages with metal wire mesh floors in a vivarium. They were maintained as described in Experiment 1.

Materials and Methods

Apparatus

The apparatus was the same as that used in Experiment 1.

Stimuli

Seven visual stimuli were used. Two of them were colored circles with geometric patterns that were used as the initial link stimuli (IL; Figure 7 top left panel). From the remaining five stimulus pool, four were randomly selected for each pigeon to be used as terminal link stimuli (TL). TL stimuli were composed of two same-colored circles (red, green, yellow, blue or orange) vertically or horizontally aligned during training, and the combination of two pairs (one vertical and one horizontal pair) of circles used during the summation test phase, both presented in the same fashion that Experiment 1 (Figure 7 top right panel). All stimuli were 100 x 100 pixels. Stimuli were presented against a grey background and

appeared on the left and right side of the screen counterbalanced throughout each session. Assignment of stimuli to function was pseudorandomly assigned across subjects.

Procedure

Training. There were two types of trials in each training session of 90 trials, free-choice (30 trials) and forced-choice (60 trials). In a free-choice trial, pigeons were presented with a choice between IL1 and IL2 presented on the left and right sides of the screen, counterbalanced. When the pigeon pecked one of the initial link stimuli, both stimuli disappeared and the chosen one was replaced by a terminal link stimulus presented for 30s on the same side of the screen as the selected initial link stimulus. Choice of IL1, the suboptimal alternative, was followed on 20% of the trials by S+, and on 80% of the trials by S-. S+ was always followed by food upon its termination, whereas S- was never followed by food. Choice of IL2, the optimal alternative, was always followed by S3 reinforced on 50% of trials in each session. A black screen presented for 10 s served as the ITI. On forced-choice trials, pigeons were presented with only one of the two alternative initial link stimuli, each appearing on 30 of the forced-choice trials. Notice that we used only one terminal link for the optimal alternative. Previous research has shown that, if other parameters are constant, the number of terminal links in the optimal alternative does not impact preference (Macias et al., 2020).

The IL stimulus alternatives for both forced-choice and free-choice trials appeared half of the time on each side of the screen (i.e., left/right counterbalanced). The duration of the session was set to a maximum of 120 minutes. Training started with the terminal link stimuli lasting 10 s. When pigeons were completing the entire training sessions, the delay was extended to 30 s for a minimum of 4 sessions before finishing the training phase. The first session of training also included two 10-s presentations of Cue E, which was later used in the summation test to assess external inhibition. One pigeon, Waluigi, failed to complete more than 10 trials in any session with terminal link stimuli of 30 s delay, and thus the delay was set to 10 s for the entire experiment. The training phase consisted of a minimum of 15 sessions and continued until performance stabilized. Stability was determined when: (a) there was no increasing or

decreasing trend in the proportion of choices over three consecutive sessions, and (b) the difference between the highest and lowest preference in the preceding three sessions was no greater than 15%.

Summation Test. The pigeons had two sessions in which inhibition to S- was assessed with a summation test. Each test session consisted of 20 trials, 10 free-choice trials of IL 1 and IL 2, and 10 test trials. Test trials consisted of 4 presentations each of the following compounds: S+S- and S+E- (E was the novel cue pre-exposed in the first session of Training), and two nonreinforced presentation of S+. Each test trial stimulus was 30 s in duration. All trials were followed by a 10 s ITI.

Data analysis

The proportion of preference for the suboptimal and informative alternative IL stimuli was calculated for each session for every pigeon in all phases of the experiment. The peck rate for each terminal link stimulus was calculated for each session in a similar fashion to Experiment 1. Session peck RRs were calculated for each stimulus and then averaged across sessions. For the Test phase, RRs were averaged for the probe stimuli in each session. Each stimulus RR was divided by the total RR across stimuli (S_{target} RR/Total RR). RM factorial ANOVAs were used to analyze training and test data, using Holm corrections for the post-hoc analysis when main effects were found.

Results

Figure 5 shows the acquisition curves for all pigeons, with circles representing data from trials with 10-s terminal link stimuli, and triangles representing data from trials with 30-s terminal link stimuli. All pigeons showed a suboptimal preference, with five of them showing a strong preference for the suboptimal alternative. A RM ANOVA found a successful acquisition of preference, $F(18,36) = 5.129$, $MSE = 0.016$, $p < .001$, $\eta^2 = 0.719$. A comparison between the last session of 10 s delay and the last session with a 30 s delay found no differences in preference, $t(6) = -0.234$, $p = .823$, $d = -0.088$. A one sample T-test using the last session of training for each pigeon found that preference for the suboptimal alternative was significantly above .5, $t(7) = 4.946$, $p < .001$, $d = 4.882$.

Figure 6 shows the mean normalized RR during the probe trials for S+ and compounds S+S- and S+E. A RM factorial ANOVA using Test Session and Stimulus Type as factors. Mauchly's test indicated that the assumption of sphericity was violated for the stimulus factor, $\chi^2(2) = 6.464$, $p = .039$, therefore degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity ($\varepsilon = .603$). We found a main effect of Stimulus Type, $F(1.205, 8.436) = 18.939$, $MSE = 0.035$, $p = .021$, $\eta^2 = 0.547$; but no effect of Test Session, $F(1,7) < 1.0$, $MSE = 9.706e^{-19}$, $p = .893$, $\eta^2 = 0.001$; nor Test Session x Stimulus Type interaction, $F(1.568, 10.975) = 4.129$, $MSE = 0.021$, $p = .052$, $\eta^2 = 0.094$. The post-hoc analysis of Stimulus Type using Holm correction showed lower RR to S+S- than to S+E ($p = .005$), RR to S+S- was also lower than to S+ ($p < .001$), and RR to S+E was lower than to S+ ($p = .032$). Thus, whereas there was evidence of a small amount of external inhibition produced by novel stimulus E, there was a much stronger effect of conditioned inhibition produced by S-. This establishes that S- becomes a conditioned inhibitor as evidence through a negative summation test.

Experiment 3

Experiment 2 replicated the typical suboptimal preference that develops during a conventional suboptimal choice task. More importantly, it found evidence of inhibitory properties to the S- in the summation test. Having established inhibition to the S- using the proper control and embedded in the suboptimal choice task, we can now test the hypothesis that development of preference for the suboptimal initial link stimulus is related to the acquisition of inhibition to the S- signaled by the suboptimal initial link stimulus. To test this, we trained pigeons on the suboptimal choice task. As in Experiment 1, test sessions were periodically introduced later in training, in which we presented the nonreinforced compounds S+S4, S3S4, S+S- and S3S-. We expected a reduction in the number of responses (RR) with compound S+S- compared to S+. Furthermore, we expected suboptimal preference to develop while the response rate to the S- decreased. Thus, we predicted that the strength of the suboptimal preference would correlate with the strength of inhibition, with birds showing stronger suboptimal preference also showing suppression of responding to the S- terminal link stimulus.

Subjects

The same pigeons were used and maintained as in Experiment 1.

Materials and Methods

Apparatus

The apparatus was the same as that used in Experiment 1 and 2.

Stimuli

A new set of terminal link stimuli were created for purpose of Experiment 3 because the same pigeons from Experiment 1 were used. The same two circles with a colorful pattern from Experiment 2 served as initial link stimuli, and four pairs of geometric shapes with a black and white pattern served as terminal link stimuli (

Figure 7, bottom panels). As in the previous experiments, the S+ and S3 pairs were arranged vertically, whereas the S- and S4 pairs were arranged horizontally. All stimuli were 100 x 100 pixels. Stimuli were presented against a grey background and appeared on the left and right side of the screen counterbalanced throughout each session. Assignment of stimuli to function was pseudorandomly assigned across subjects.

Procedure

Pre-training. All stimuli (two initial links and four terminal links) were presented individually in a similar fashion as in Experiments 1 and 2.

Training. There were two types of trials in each training session of 80 trials, free-choice (20 trials) and forced-choice (60 trials). The trials were presented as described in Experiment 2 (Figure 7, left panel). The duration of the session was set to a maximum of 90 minutes. Pigeons initially received 15 sessions of training, after which blocks of 5 training sessions were alternated with blocks of 5 test sessions, with a total of 30 sessions as described in this phase.

Test. Pigeons received four blocks of five test sessions. A test session had 96 trials, 80 training trials as described above, and 16 nonreinforced compound test trials. The first 10 trials in each test session always consisted of forced-choice training trials. After the 10^{th} trial of the session, 50 forced-choice, 20 free-choice, and 16 test trials were randomly interspersed throughout the remainder of the session. Test trials entailed the presentation of one of four compound stimuli (Figure 7, right panel) presented for 30 s on half of the trials on the left side and the other half on the right side of the screen.

Sessions were included in data analysis only if at least half of the training trials and all of the test trials had been presented. In consequence, two pigeons, Cousteau and Darwin, repeated some training sessions (10 and 5, respectively) and some test sessions (11 and 3, respectively). Response rates were averaged in blocks of five sessions, obtaining six blocks (seven for Herriot) for the training phase and four (five for Herriot) blocks for the test phase.

Data analysis

Preference for the suboptimal alternative was defined as the number of choices to the suboptimal alternative divided by the total number of choice trials completed. RR to each terminal link stimulus or compound test stimulus was also recorded on each trial. Given that the RRs were highly variable across pigeons, the data were normalized to allow comparisons as it was done in previous experiments (RR stim/ total RR). As with choice preference, the preference was collapsed in blocks of five sessions. RM factorial ANOVAs were implemented to analyze training and test data, using Holm corrections for the post-hoc analysis when main effects were found. Because we wanted to look at differences in preference and RR, many of the test phase data were analyzed separately for each pigeon.

Results

The left panel of Figure 8 shows choice preference across training blocks for each pigeon. The rate at which suboptimal preference developed varied across pigeons. Hawthorne acquired an almost exclusive preference for the suboptimal alternative by Block 3 (before the first test block was introduced). Goodall and Cousteau reached a similarly high preference for the suboptimal alternative by Block 4. Although not as strong as in the three birds discussed above, Darwin reached an asymptotic level of preference for the suboptimal alternative of around 0.6 by the second block of training. Finally, Herriot began with a strong preference for the optimal alternative, and only by Block 5 had reached 50% preference for the suboptimal alternative, reaching a suboptimal preference close to .7 by the 6th block of training. An RM ANOVA conducted on percentage of suboptimal preference with Block as a factor found a main effect of Block, $F(5, 20) = 15.72$, $MSE = 0.010$, $p < .001$, $\eta^2 = .797$. Post-hoc analyses using Holm correction for the Block factor found a significant difference between Block 1 and Blocks 5 (p = .034) and 6 ($p = .006$), and Block 2 against Blocks 4 to 6 ($p_{2vs4} = .11$, $p_{2vs5} = .004$ and $p_{2vs6} = .025$), suggesting that all pigeons developed a suboptimal preference. Whereas four out of five pigeons showed a strong preference for the suboptimal alternative (>.7), Darwin showed a preference near .6, that could suggest

indifference between alternatives, but it is also considered a suboptimal preference because the animal is losing reinforcement.

Figure 8 right panel shows choice preference for the suboptimal choice during test sessions. Given that one of the pigeons (Herriot) had an extra block of testing, the first two blocks (M_1= .64 and M_2 = .64) were averaged to facilitate comparison with the other four pigeons. The graph suggests that preference during test blocks did not differ from the preference shown in later blocks of training. An RM ANOVA was conducted on preference for suboptimal choice during test blocks using Block as a factor. Mauchly's test indicated that the assumption of sphericity was violated, $\chi^2(5) = 13.792$, $p = .024$, therefore degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity ($\varepsilon = .513$). The results found no main effect of Block, $F(1.539, 6.158) = 1.117$, $MSE = 0.024$, $p = .366$, $\eta^2 = .218$, which confirmed the visual inspection that preference did not change through test blocks.

Figure 9 shows the normalized average RR to each stimulus across all blocks of training. An overall similar RR to S+, S3, and S4 was maintained across all blocks of training, whereas the RR to S- dropped considerably from Block 1 to Block 3 after which it remained close to 0. An RM factorial ANOVA conducted on normalized RR during training with Block and Stimulus as factors found a main effect of Stimulus, $F(3, 12) = 40.569$, $MSE = 0.012$, $p < .001$, $\eta^2 = .706$, and a significant Block X Stimulus interaction, $F(15, 60) = 4.369$, $MSE = 0.004$, $p < .001$, $\eta^2 = .117$, but no effect of Block, $F(5,20) = 0.609$, $MSE = 1{,}780e^{-19}$, $p = .694$, $\eta^2 = .002$. Post-hoc analyses of Stimulus using Holm correction found a significant difference between S- and the other three stimuli ($p < .001$), there is also a difference between S+ and S4 ($p = .012$). This result and the significant interaction suggest that RR changed differently over time for the different stimuli, with RR to S- dropping rapidly.

Figure 10 shows the mean normalized response rate across all blocks of testing per stimulus for each subject. The left four bars correspond to the training stimuli whereas the right four bars correspond to the compound stimuli. S+ had the highest RR, followed by S3 and S4 which had a similar RR, and finally S-, the stimulus that had never been reinforced, had an RR close to zero. Overall, RRs to

compounds were lower than to elements, suggesting either generalization decrement or that animals were learning that the compounds were never reinforced due to the repeated nonreinforced testing. An RM factorial ANOVA conducted on normalized RR during test sessions with Block and Stimulus as factors found a main effect of Stimulus, $F(7, 28) = 17.637$, $MSE = 0.010$, $p < .001$, $\eta^2 = .703$, and Block X Stimulus interaction, $F(21,84)=1.975$, $MSE = 0.002$, $p= .015$, $\eta^2 = 0.46$; but no effect of Block, $F(3, 12) = 0.442$, $MSE = 3.532e^{-19}$, $p = .727$, $\eta^2 = .002$. Therefore, the test data were pooled across blocks to test specific predictions. Planned-comparisons found that RR to S+ was higher than to S+S- ($p < .001$). S+ RR was also higher than S+S4 RR ($p <.001$), suggesting that the repeated testing across multiple training/testing cycles could have resulted in pigeons learning that compound trials were never reinforced, thus producing generalization decrement on compound trials. Because we observed generalization decrement to all the compound test types, the analyses below utilized RR to S+ and S- alone as measures of excitation and inhibition, following the validity of S- responses as an indicator of conditioned inhibition initially established by Wasserman et al. (1974; see also Hearst & Franklin, 1977; Hearst et al., 1980).

Figure 11 shows the correlations between suboptimal preference and S+ RR (left panels) and S- RR (right panels) for each pigeon. We hypothesized that the development of inhibitory properties of S- should correlate with the level of suboptimal preference observed by each pigeon. It can also be the case, however, that learning to S+ also tracks suboptimal preference, in which case responses to S+ should also have a strong correlation with preference. Given that pigeons developed a suboptimal preference at different rates, we performed separate Pearson's R correlations between S+ RR and suboptimal preference, and between S- RR and suboptimal preference, by bird. The comparison between suboptimal preference and S+ RR across training showed a strong positive correlation for Goodall (r= .89, $p = .017$) and Cousteau (r = .85, $p = .032$), a moderate but nonsignificant positive correlation for Hawthorne (r= .65, $p = .369$) and Herriot (r = .45, $p = .158$), and a very low, nonsignificant positive correlation for Darwin (r = .16, p = .746). For the comparison between suboptimal preference and S- RR across training,

Hawthorne, Goodall, and Cousteau each showed a strong negative correlation (r = -.95, r = -.85, and r = -.72, respectively; ps < .05). Herriot and Darwin also showed a negative but nonsignificant correlation (r = -.58, p = .227; r = -.58, p = .222, respectively). Comparing both cues, we observed that neither S+ nor S- RR correlates with preference for Darwin or Herriot. For Hawthorne, S- RR correlates with suboptimal preference; and for Goodall and Cousteau S+ and S- RR both correlate to a similar level with preference. A one-sample t-test using the Pearson's r values of S- RR collapsed across all birds suggests that the mean Pearson's r was significantly more negative than zero, $t(4)$= -10.22, p < .001, d = -4.568, confirming a general negative relationship between S- RR and the level of preference for the suboptimal alternative across subjects. A one-sample t-test using the Pearson's r values of S+ RR collapsed across pigeons found a significant general positive relationship between S+ RR and the level of suboptimal preference between subjects, $t(4)$= 4.475, p = .006, d = 2.001 These results support the hypothesis that inhibition to S- is related to the development of suboptimal preference.

The above analysis tells us about how preference and rate of response to S- changed for each individual subject across training sessions. Because suboptimal choice and response rate increases to S+ or decreases to S- across training are both confounded with the amount of training, we followed up with an analysis of these correlations across pigeons at asymptote. The correlation between asymptotic levels of suboptimal preference and S+ RR or S- RR are shown in Figure 12. The small n renders these analyses very underpowered. Nevertheless, we observed a moderate linear correlation between S+ RR and asymptotic level of suboptimal preference (r = .36, p = .275) and a moderate negative linear correlation between S- RR and suboptimal preference (r = -.42, p = .271). Although these results are non-significant, they suggest a trend between the degree of suboptimal preference and what is learned about S+ and S- terminal link stimuli. Presumably with a higher-powered study these analyses, especially that involving the S-, would yield significant correlations. Whether this relationship reflects a direct or indirect causal link remains an empirical question for further study. Despite the strong generalization decrement on all compound test trial types, we include an analysis of the relationship between the mean asymptotic

preference that was reached during test blocks and the mean inhibitory properties of S- expressed as the difference in RR between S+ and S+S- (i.e., the summation test of conditioned inhibition). This analysis revealed a moderate but non-significant correlation ($r = .327$, $p = .295$), though this analysis should be interpreted with caution given the strong generalization decrement observed even on compound test trials that did not involve the S-.

General discussion

The preference observed in the suboptimal choice task is an intriguing phenomenon. The mechanisms that underlie this preference are still open to debate. The present studies tested the hypothesis that the S- develops inhibitory properties, and the development of S- inhibition is related to the development of suboptimal preference.

In Experiment 1, pigeons were trained only with the terminal link stimuli. Using this procedure, we found conditioned inhibition developed to S- with training, as assessed in a summation test. Moreover, we observed that the relative rate of response to each test compound correlated with the combined excitatory and inhibitory values of the elements. The compounds that contained only continuously or partially reinforced stimuli elicited higher peck rates than compounds that contained the nonreinforced S-. These results suggest that S- had acquired inhibitory properties such that it passed a summation test of conditioned inhibition. Inhibition on summation tests had a stronger suppressive effect on a partially reinforced (S3S-) compared to a continuously reinforced (S+S-) stimulus.

In Experiment 2, pigeons were trained on the conventional suboptimal choice task and then received two summation tests in which the RR to S+ was compared to compounds S+S- and S+E, where E was a novel cue with no prior training history. Experiment 2 replicated the pattern of responding to the compound test stimuli observed in Experiment 1, but following training on a full suboptimal choice procedure. Moreover, we found evidence of conditioned inhibition independent from, or in addition to, any effect of external inhibition. That is, response rates to S+S- were significantly lower than to S+E.

Thus, the terminal link S- stimulus shows evidence of becoming a true conditioned inhibitor. Finally, in Experiment 3, we again replicated the suboptimal choice procedure, and included periodic summation tests throughout the latter half of training. Pigeons developed suboptimal preference over training. We found a negative correlation between the strength of suboptimal preference and the response rate to the S-, suggesting that as each bird acquired a suboptimal preference, the S- became more inhibitory. We chose to use response rate to the S- as a measure of inhibition following its validity established by Wasserman et al. (1974; see also, Hearst & Franklin, 1977; Hearst et al., 1980) because, unlike in Experiments 1 and 2, in Experiment 3 we observed strong generalization decrement on all compound test trials, even those that only contained trained excitatory stimuli as elements.

While we also found, as expected, a positive correlation between response rates to S+ and strength of suboptimal preference by the end of training, more importantly, we observed a moderate negative correlation with RR to the S- and the strength of suboptimal preference by the end of training, suggesting that the stronger the inhibition to S-, the stronger the asymptotic level of suboptimal preference. Unfortunately, this analysis was too underpowered to detect a significant moderate effect.

We should point out that these data are only correlational and may or may not reflect a causal relationship between suboptimal preference and learning about S-. One possible explanation for the correlation is that the development of preference for the suboptimal alternative initial link stimulus is dependent on development of inhibition to the S- terminal link stimulus. Alternatively, the development of inhibition to the S- and of suboptimal preference could each be the result of another, yet-to-be-identified processes. We also found that for some pigeons, responses to S+ also correlated with preference. A third alternative could be that learning excitatory properties of S+ and inhibitory properties of S- interact to make possible the development of suboptimal preference. Another issue with interpreting the correlation is that changes in response rate to the S+ and S- cues are confounded with the amount of training the pigeons had received, thus it is not impossible to rule out other learning processes that could account for the observed preference, though what those learning processes are needs to be identified.

Further empirical work is needed to better understand the source of the correlation, with the aid of more theoretical considerations.

Despite the open questions regarding the source of the correlation between inhibition to the S- and development of suboptimal preference, our results contradict some previous studies. Previously, Laude et al. (2014) failed to find evidence supporting any relationship between S- inhibition and suboptimal choice, in fact showing that inhibition to S- waned with training compared with the inhibition showed early in training. Their result could indicate a de-correlation between inhibition and preference; however, we can only speculate because they did not report the changes in response to S- across the training history. Also, we identified many shortcomings with their study that precludes strong support of their conclusions (see Introduction). Nevertheless, as we discussed in the introduction, the procedural differences between their study and ours are considerable. Our study addresses some of these potentially problematic issues. First, we used the traditional probabilistic version of the suboptimal task, in which the choice is between an informative and a non-informative alternative; instead of the magnitude version of the task, in which the choice is between two informative alternatives that signal different magnitudes of reward. Second, the stimulus that served as the S- in our procedure was from the same dimension as the other terminal link stimuli, thereby avoiding potential confounds due to generalization decrement or differences in associability of the S- relative to the other stimuli. Third, the stimuli that served as the terminal links, including the S-, were counterbalanced across pigeons. Fourth, we assessed inhibition via summation tests with compounds of terminal link stimuli, including a key control in which the S+ was presented with a familiar but non trained cue to rule out external inhibition as an explanation for the response decrement with the compound. Fifth, we assessed inhibition at various time points during training so that we could track the development of conditioned inhibition. Finally, we analyzed individual subject data to assess individual differences in learning and performance as a more sensitive test of the relationship between inhibition to the S- and suboptimal choice.

As mentioned before, Trujano et al. (2016) report an experiment evaluating inhibition in the suboptimal choice procedure with rats. They found evidence of inhibition to the S- terminal link stimulus, but in their study, rats never developed suboptimal preference. They suggested that there is a difference between how rats and pigeons learn the suboptimal choice task; and that the lack of inhibition in previous experiments with pigeons indicates that pigeons did not encode the task in the same way. Instead, rats that showed inhibition to the S- chose the optimal alternative, suggesting they were sensitive to the signal values of the terminal link stimuli. We report contrasting results, however. Moreover, other researchers have found rats to be suboptimal when the difference in overall probability of reinforcement between alternatives is reduced (Ojeda, Murphy & Kacelnik, 2018), the when delay to reinforcement is increased (Cunningham & Shahan, 2019), or when levers are used as initial link stimuli rather than lights or other visual cues (Chow et al., 2017). It would be interesting to evaluate the role of inhibition to the S- in a task in which rats develop a preference for the suboptimal alternative. It is possible that the results of Trujano et al. (2016) could be explained by a difference in parameters, rather than an intrinsic species difference between rats and pigeons.

The results of our experiments might inform on the role of the S- terminal link stimulus in models of suboptimal choice. Current versions of the RRM (Vasconcelos, Monteiro & Kacelnik, 2015) developed from an ecological perspective; or the temporal information hypothesis (Cunningham & Shahan, 2018) among others; assume that the S- is ignored by the subject in the sense that it plays no role in the subjects' choice. The present results at least challenge the mechanisms assumed by these models. If an animal is learning about a stimulus, this should somehow contribute to the choice. Indeed, it is possible that the inhibitory properties of the S- indicate that the subject is learning that a stimulus (the S-) is never reinforced.

In comparison, the Delta-Sigma model (Gonzalez et al., 2020) treats S- as just another value that is 'contrasted' to assign value to an alternative. This model claims that the contrast between the two terminal-link stimuli that follow the same initial link stimulus is a primary factor in the decision process:

the greater the contrast, the greater the preference for that alternative. In this case, it is possible that inhibition to S- increases the appeal of S+, which could result in overweighing the latter in the decision process. This is specially interesting if we consider the early literature on the observing response that suggested that animals prefer information over no information. As mentioned previously, classic information theory suggests (Shannon & Weaver, 1949) a signal followed by food with probability 1 and 0 are equally informative. Nevertheless, some experiments studying the observing response suggested that animals prefer information about food (good news) over information about absence of food (bad news, Dinsmoor, 1983). However, in the present situation, animals are faced with a choice between an initial link stimulus that is followed by informative cues vs an initial link stimulus that is followed by noninformative cues. The informative alternative is followed on some trials by good news, and on other trials by bad news. Our evidence that animals learn that S- becomes inhibitory, and that the strength of inhibition tracks the strength of suboptimal preference, suggests that animals prefer information over no information, even when this information is about bad news. How the two sources of information are combined to assign value is an open question that requires further research.

Here, we propose that by acquiring properties of conditioned inhibition, the S- predicts the explicit absence of food. We further suggest that the S- becoming a conditioned inhibitor might be causally related to the development of suboptimal preference, though this stronger claim has yet to be empirically tested. This adds to the explanation assumed for some models of suboptimal choice that it is the combine effect of S+ and S- that accounts for the development of suboptimal preference. Indeed, it is possible that the learning to S- increases the value of S+, biasing the preference for an alternative that overall gives less food. This is in line with other experiments that did not find a big impact of overall reinforcement rate when both alternatives were informative (Experiment 2, Gonzalez et al., 2020; Zentall, & Stagner, 2011). Future assessments of the relationship between conditioned inhibition to the S- and development of suboptimal preference could inform the development of alternative models of suboptimal choice that include a role for processes of inhibition.

References

Blaisdell, A. P., & Cook, R. G. (2005). Two-item same-different concept learning in pigeons. *Animal Learning & Behavior, 33*(1), 67-77.

Bower, G., Mclean, J., & Meacham, J. (1966). Value of Knowing When Reinforcement Is Due. *Journal of Comparative and Physiological Psychology, 62*(2), 184–192. https://doi.org/10.1037/h0023682

Browne, M. P., & Dinsmoor, J. A. (1974). Wyckoff's observing response: Pigeons learn to observe stimuli for free food but not stimuli for extinction. *Learning and Motivation, 5*(2), 165–173. https://doi.org/10.1016/0023-9690(74)90023-X

Chow, J. J., Smith, A. P., Wilson, A. G., Zentall, T. R., & Beckmann, J. S. (2017). Suboptimal choice in rats: Incentive salience attribution promotes maladaptive decision-making. *Behavioural Brain Research, 320*, 244-254.

Cunningham, P. J., & Shahan, T. A. (2018). Suboptimal choice, reward-predictive signals, and temporal information. *Journal of Experimental Psychology: Animal Learning and Cognition*, *44*(1), 1–22. https://doi.org/10.1037/xan0000160

Cunningham, P. J., & Shahan, T. A. (2019). Rats Engage in Suboptimal Choice When the Delay to Food Is Sufficiently Long. *Journal of Experimental Psychology: Animal Learning and Cognition, 45*, 301-310. https://doi.org/10.1037/xan0000211

Cole, R. P., Barnet, R. C., & Miller, R. R. (1997). An evaluation of conditioned inhibition as defined by Rescorla's two-test strategy. *Learning and Motivation, 28*(3), 323-341.

Daniels, C. W., & Sanabria, F. (2018). An Associability Decay Model of Paradoxical Choice. *Journal of Experimental Psychology: Animal Learning and Cognition*, *44*(3), 258–271. http://dx.doi.org/10.1037/xan0000179

Dinsmoor, J. A. (1983). Observing and conditioned reinforcement. *Behavioral and Brain Sciences, 6*(4), 693–704. https://doi.org/10.1017/S0140525X00017969

Domjan, M. (2014). *The principles of learning and behavior*. Nelson Education.

Fanselow, M. S. (1980). Signaled shock-free periods and preference for signaled shock. *Journal of Experimental Psychology: Animal Behavior Processes, 6*(1), 65.

Fantino, E., & Silberberg, A. (2010). Revisiting the role of bad news in maintaining human observing behavior. J*ournal of the Experimental Analysis of Behavior, 93*(2), 157–170. https://doi.org/10.1901/jeab.2010.93-157

Gipson, C. D., Alessandri, J. J. D., Miller, H. C., & Zentall, T. R. (2009). Preference for 50% reinforcement over 75% reinforcement by pigeons. *Learning & Behavior*, *37*(4), 289–298. http://doi.org/10.3758/LB.37.4.289

González, V. V., Macías, A., Machado, A., & Vasconcelos, M. (2020). The Δ–∑ hypothesis: How contrast and reinforcement rate combine to generate suboptimal choice. *Journal of the Experimental Analysis of Behavior*, *113*(3), 591–608. https://doi.org/10.1002/jeab.595

Hearst, E., & Franklin, S. R. (1977). Positive and negative relations between a signal and food: Approach-withdrawal behavior to the signal. *Journal of Experimental Psychology: Animal Behavior Processes, 3*(1), 37–52. https://doi.org/10.1037/0097-7403.3.1.37Hearst, E., Bottjer, S. W., & Walker, E. (1980). Conditioned approach-withdrawal behavior and some signal-food relations in pigeons: Performance and positive vs. negative "associative strength". *Bulletin of the Psychonomic Society, 16*(3), 183-186.

Laude, J. R., Stagner, J. P., & Zentall, T. R. (2014). Suboptimal choice by pigeons may result from the diminishing effect of nonreinforcement. *Journal of Experimental Psychology: Animal Learning and Cognition*, *40*(1), 12–21. http://doi.org/10.1037/xan0000010

Lieberman, D. A., Cathro, J. S., Nichol, K., & Watson, E. (1997). The role of S-in human observing behavior: Bad news is sometimes better than no news. *Learning and Motivation, 28*(1), 20–42. https://doi.org/10.1006/lmot.1997.0951

Lockard, J. S. (1963). Choice of a warning signal or no warning signal in an unavoidable shock situation. *Journal of Comparative and Physiological Psychology, 56*(3), 526.

Macías, A., González, V. V., Machado, A., & Vasconcelos, M. (2021). The functional equivalence of two variants of the suboptimal choice task: choice proportion and response latency as measures of value. *Animal cognition, 24*(1), 85–98. https://doi.org/10.1007/s10071-020-01418-8

McDevitt, M. A., Dunn, R. M., Spetch, M. L., & Ludvig, E. A. (2016). When good news leads to bad choices. *Journal of the Experimental Analysis of Behavior*. https://doi.org/10.1002/jeab.192

Miller, R. R., Marlin, N. A., & Berk, A. M. (1977). Reliability and sources of control of preference for signaled shock. Animal Learning & Behavior, 5(3), 303-308.

Ojeda, A., Murphy, R. A., & Kacelnik, A. (2018). Paradoxical choice in rats: Subjective valuation and mechanism of choice. *Behavioural Processes*, *152*(March), 73–80. https://doi.org/10.1016/j.beproc.2018.03.024

Pavlov, I. P. (1927). *Conditioned reflexes: an investigation of the physiological activity of the cerebral cortex.* Oxford Univ. Press.

Prokasy, W. F. (1956). The acquisition of observing responses in the absence of differential external reinforcement. *Journal of Comparative and Physiological Psychology, 49*(2), 131–134. https://doi.org/10.1037/h0046740

Rescorla, R. A. (1969). Pavlovian conditioned inhibition. *Psychological Bulletin*, *72*(2), 77–94. http://doi.org/10.1037/h0027760

Roper, K. L., & Zentall, T. R. (1999). Observing Behavior in Pigeons: The Effect of Reinforcement Probability and Response Cost Using a Symmetrical Choice Procedure. *Learning and Motivation, 30*(3), 201–220. https://doi.org/10.1006/lmot.1999.1030

Shannon, C. E., & Weaver, W. (1949). *The mathematical theory of communication*. The University of Illinois Press (Vol. 25). Chicago, IL. https://doi.org/10.4992/jjpsy.25.110

Stagner, J. P., & Zentall, T. R. (2010). Suboptimal choice behavior by pigeons. *Psychonomic Bulletin & Review*, *17*(3), 412–6. http://doi.org/10.3758/PBR.17.3.412

Trujano, R. E., López, P., Rojas-Leguizamón, M., & Orduña, V. (2016). Optimal behavior by rats in a choice task is associated to a persistent conditioned inhibition effect. *Behavioural Processes*, *130*, 65–70. https://doi.org/10.1016/j.beproc.2016.07.005

Vasconcelos, M., Monteiro, T., & Kacelnik, A. (2015). Irrational choice and the value of information. *Scientific Reports*, *5*, 13874. https://doi.org/10.1038/srep13874

Wyckoff, L. B. (1952). The role of observing responses in discrimination learning. Part I. P*sychological Review, 59*(6), 431–442. https://doi.org/10.1037/h0053932

Wasserman, E. A., Franklin, S. R., & Hearst, E. (1974). Pavlovian appetitive contingencies and approach versus withdrawal to conditioned stimuli in pigeons. *Journal of Comparative and Physiological Psychology, 86*(4), 616–627. https://doi.org/10.1037/h0036171

Zentall, T. R. (2016). Resolving the paradox of suboptimal choice. *Journal of Experimental Psychology: Animal Learning and Cognition*, *42*(1), 1–14. https://doi.org/10.1037/xan0000085

Figure 1.

Procedure used by Laude, Stagner and Zentall (2014)

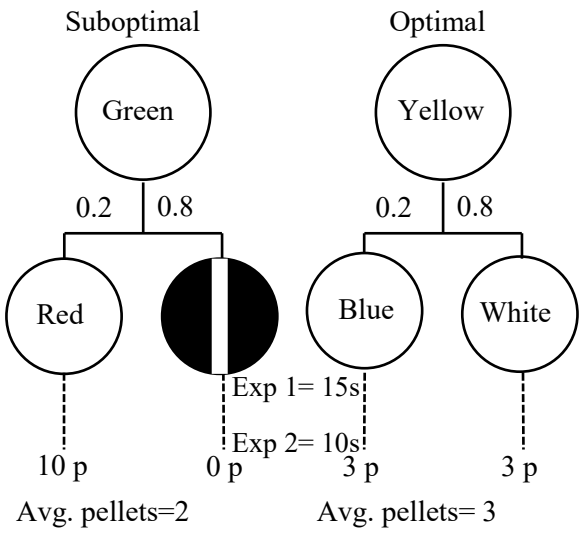

Note. They evaluated the inhibitory properties of the stimulus never reinforced (S-, black circle with a white vertical line). After the choice, each signal was followed by a different magnitude of reinforcement.

Figure 2.

Stimuli used in Experiment 1.

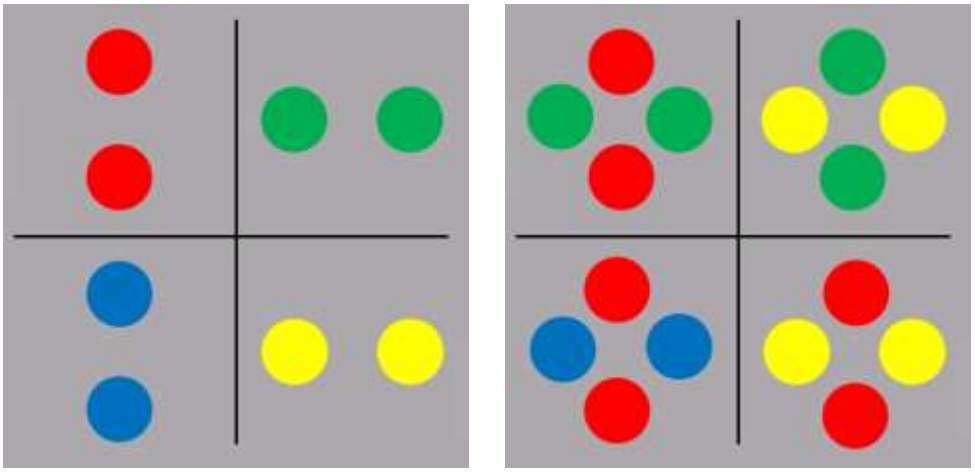

Note. Left panel: Examples of stimuli presented during the Pre-training and Training phases of Experiment 1. Right panel: Examples of compound stimuli presented on probe trials during the Testing phase of Experiment 1. Color and orientation were counterbalanced across pigeons with the restriction that S+ and S3, and S- and S4 should have the same orientation to be able to create the compound stimuli.

Figure 3.

Mean normalized RR for each block of five sessions of training (Upper panel) and testing (Lower panel) for each stimulus in Experiment 1.

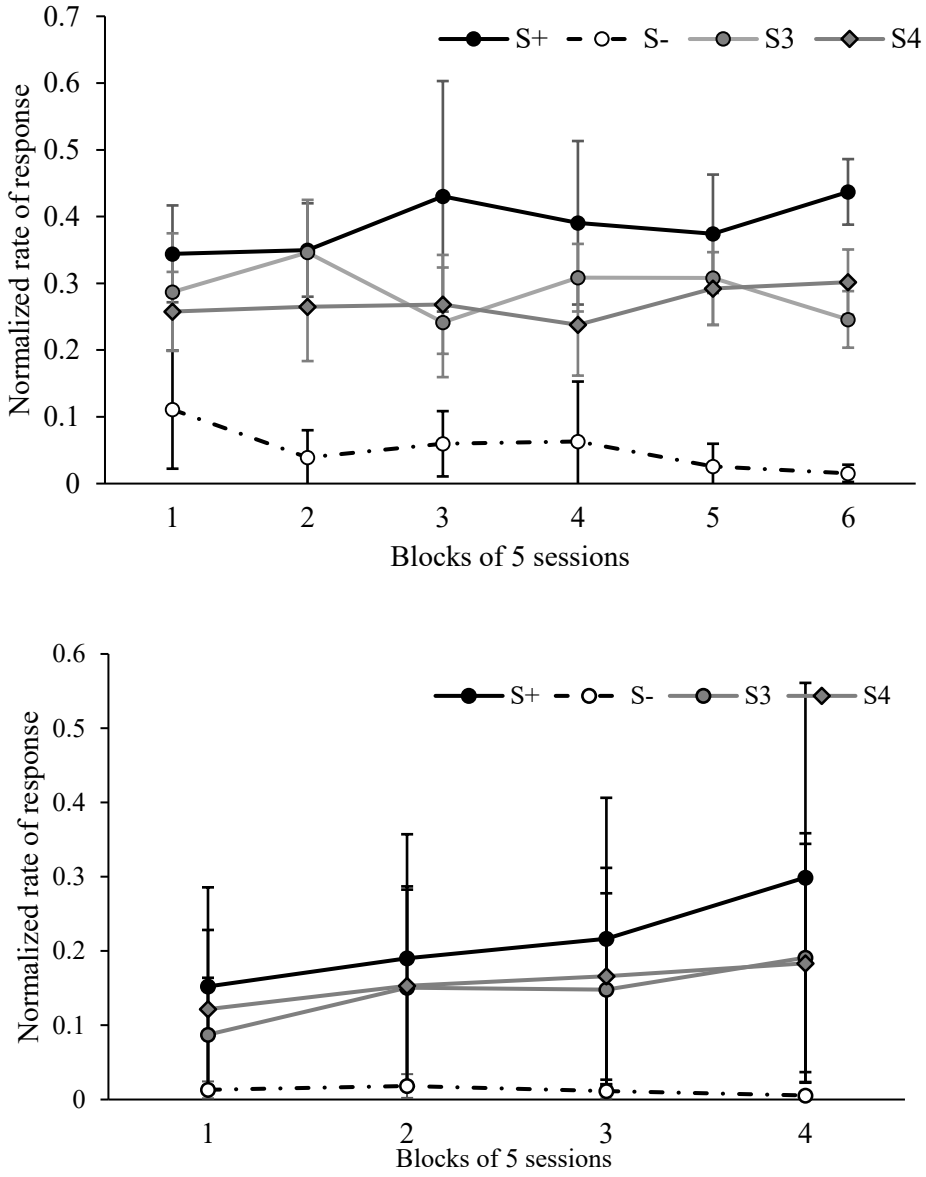

Note. Error bars represent 95% confidence interval.

Figure 4.

Mean normalized RR across the last two blocks of testing for each probe stimulus in Experiment 1.

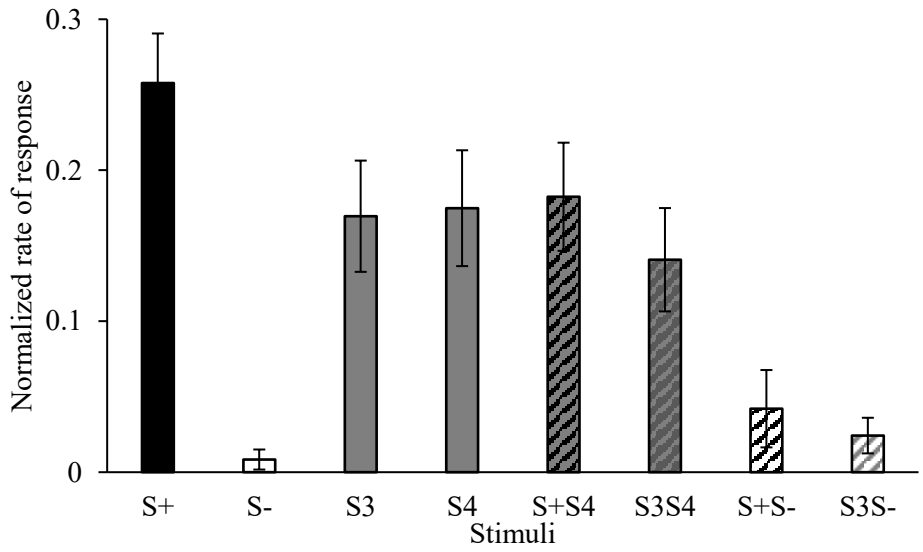

Note. Error bars correspond to the 95% confidence interval.

Figure 5.

Suboptimal preference for each pigeon across session for Experiment 2.

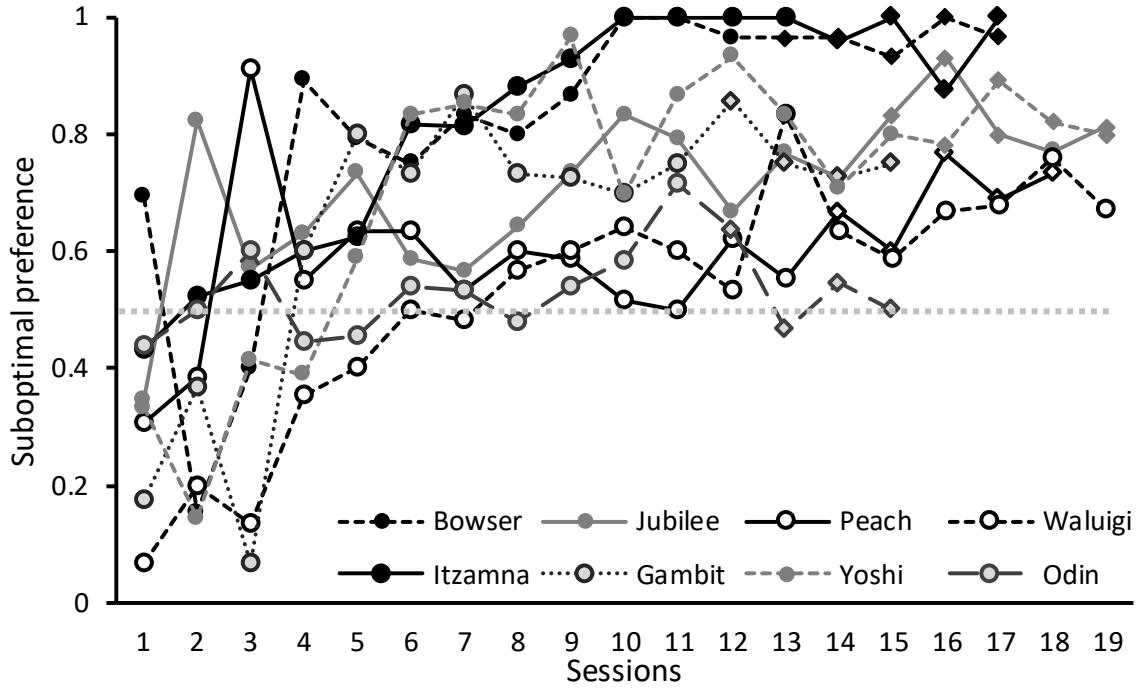

Note. The circles correspond to 10 s and the diamonds to 30 s terminal link duration.

Figure 6.

Mean normalized RR during testing for each stimulus for Experiment 2.

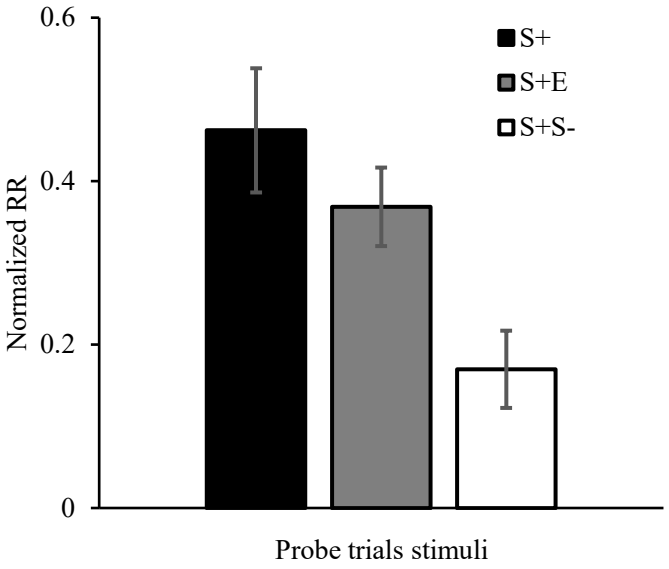

Note. Error bars correspond to the 95% confidence interval.

Figure 7.

Suboptimal task and stimuli examples in Experiment 2 (top panel) and Experiment 3 (bottom panel).

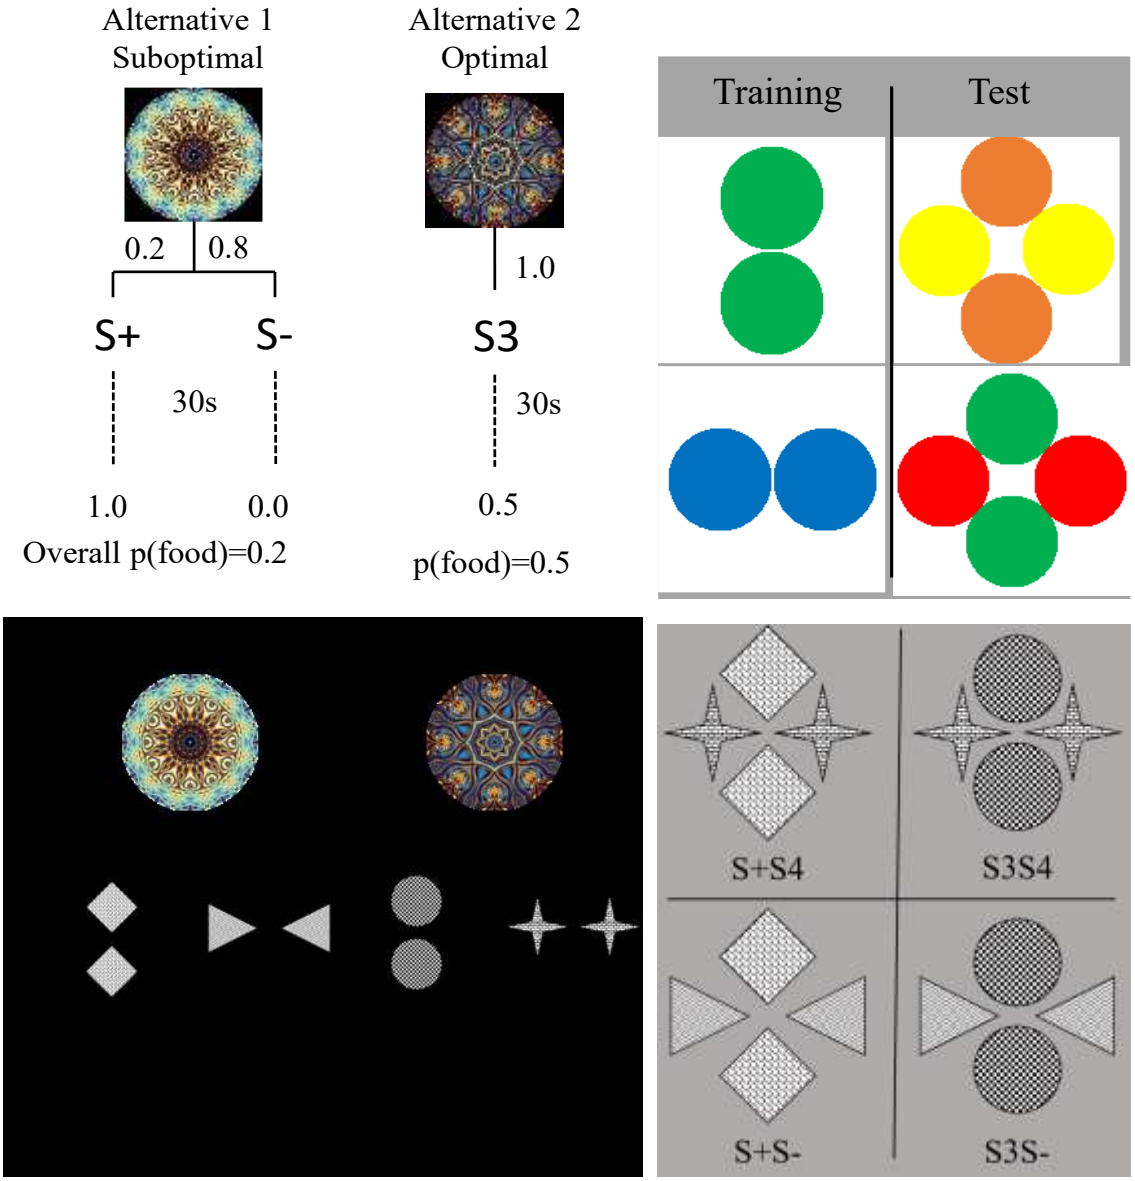

Note. Left top and bottom panel: Structure of the suboptimal choice task used on choice trials during training in Experiment 2 and 3. The initial link stimuli were counterbalanced across pigeons. Terminal link stimuli could be two pairs of colors (green, red, blue, yellow or orange) for Experiment 2 or figures (diamonds, triangles, circles, or 4-point stars) for Experiment 3, presented vertically (S+ and S3) or horizontally (S- and S4). For instance in the bottom left figure, the vertical pair of diamonds correspond

to the S+, the horizontal pair of triangles correspond to the S-, the vertical pair of circles correspond to the S3, and the horizontal pair of stars correspond to the S4. *Right top and bottom panel*: Examples of four possible training and compound stimuli used during testing in Experiment 2 (top) and compound stimuli used in Experiment 3 (bottom). Stimuli were presented half of the time in each side of the screen within each test session.

Figure 8.

Proportion of suboptimal preference during training (left panel) and testing (right panel) in blocks of five sessions by pigeon in Experiment 3.

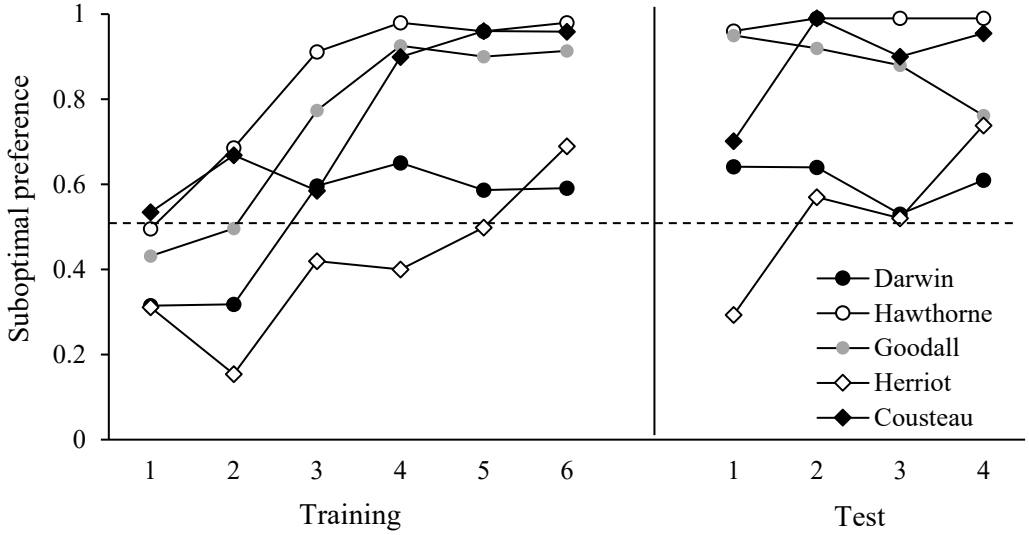

Note. As with the data analysis, for Herriot blocks 2 and 3 of training, and 1 and 2 for test were averaged for a better comparison.

Figure 9.

Mean normalized RR to each stimulus by 5-session block of training in Experiment 3.

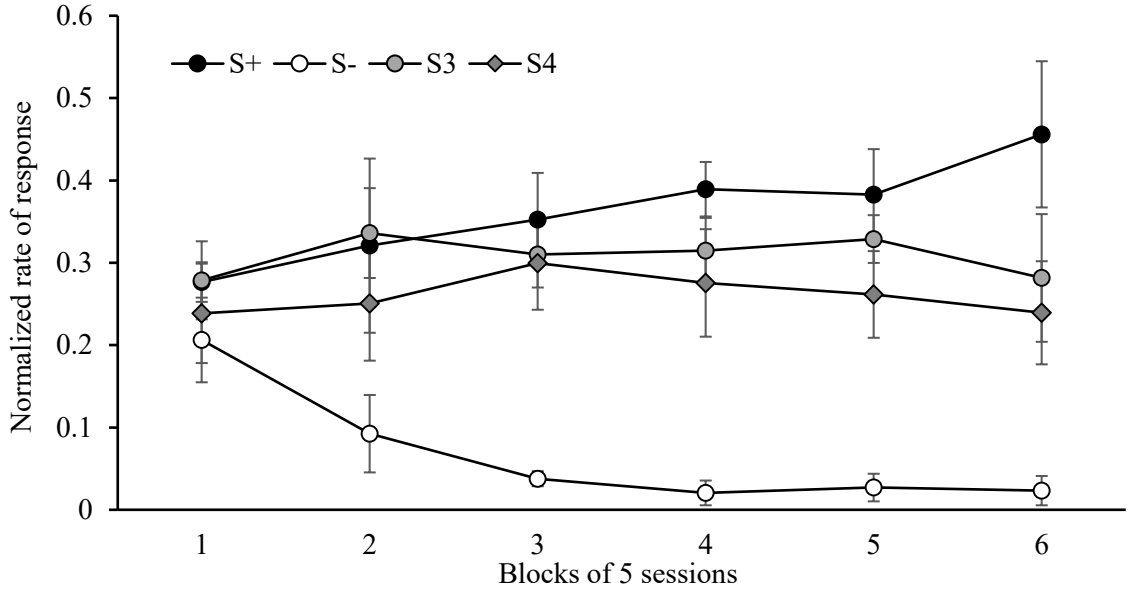

Note. For Herriot, blocks 2 and 3 of training were averaged. Error bars correspond to the 95% confidence interval.

Figure 10.

Mean normalized RR across all blocks of testing for each stimulus in Experiment 3.

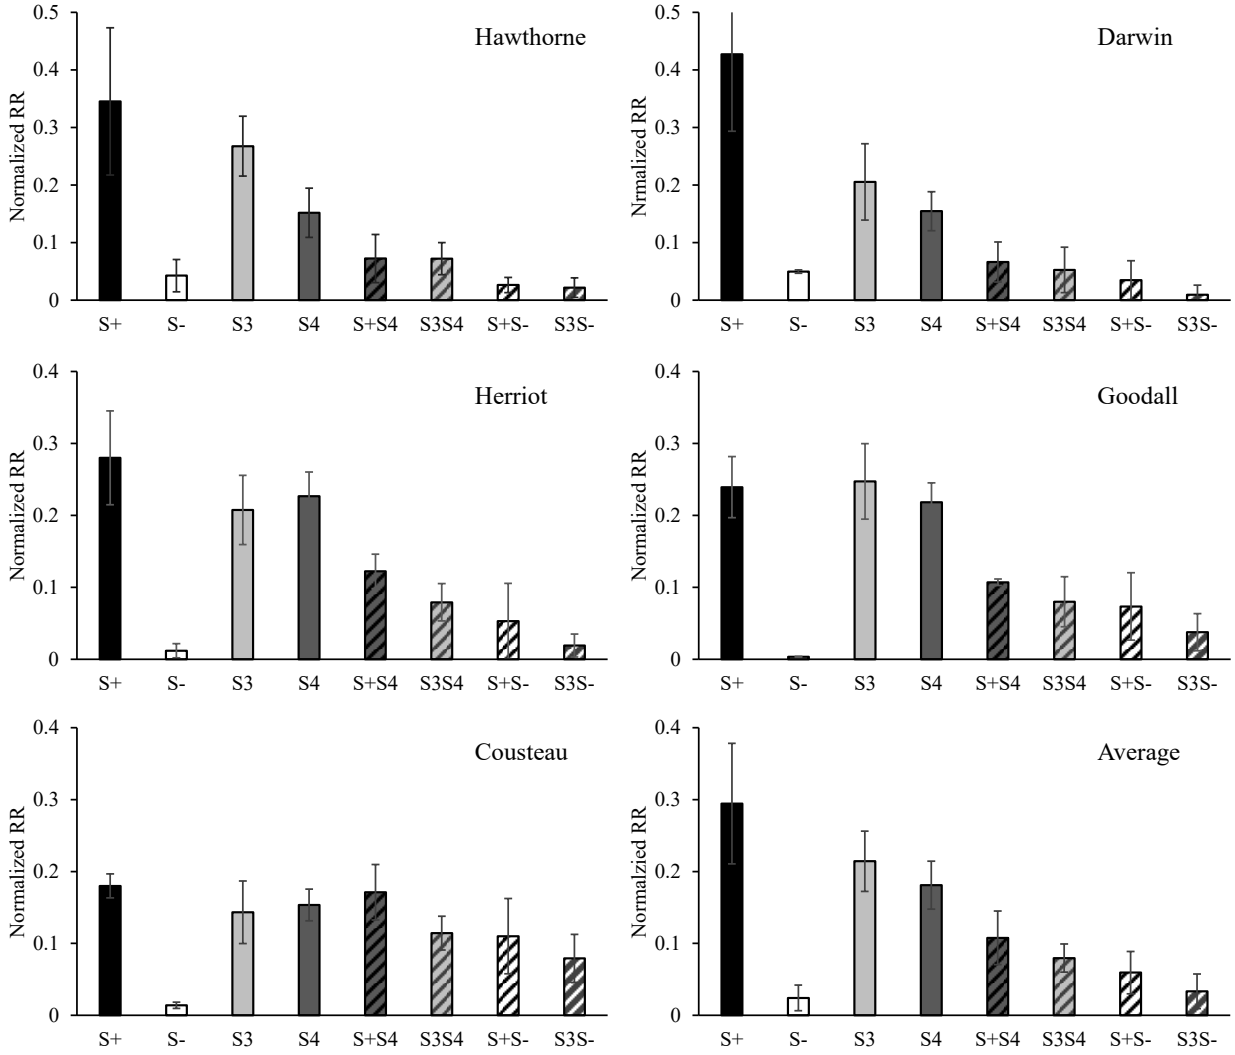

Note. Error bars correspond to the 95% confidence interval.

Figure 11.

The correlation between suboptimal preference and normalized RR to S+ (left panels) and S- (right panel) during the 6 blocks of training for each pigeon in Experiment 3.

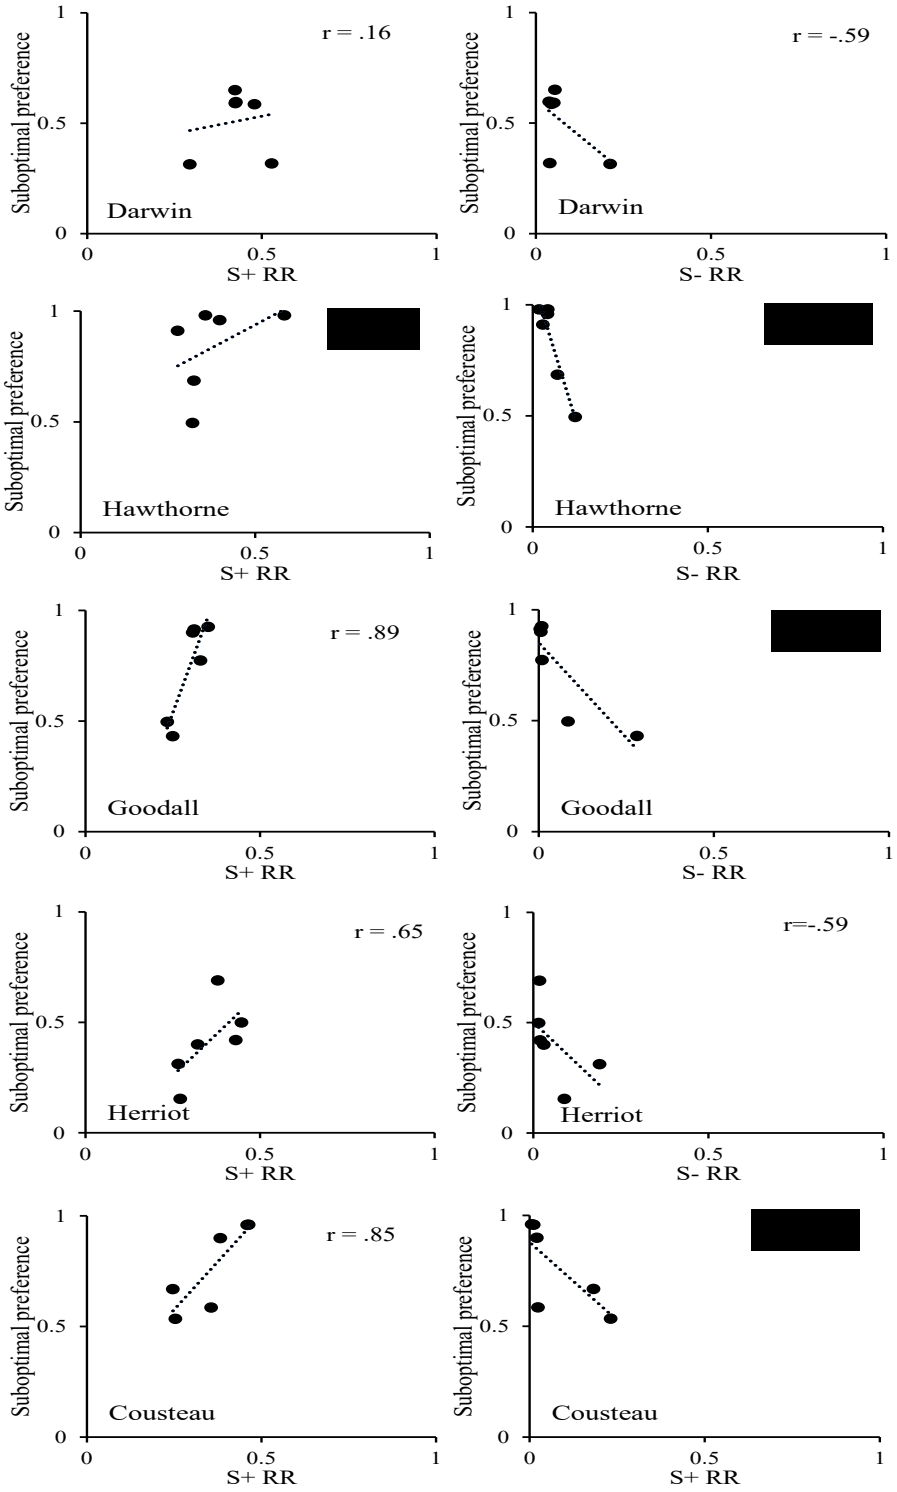

Note. Dashed lines represent the best linear fitting. r = Pearson's correlation.

Figure 12.

The correlation between suboptimal preference and normalized RR to S+ (left panel) and S- (right panel) during the last block of training in Experiment 3.

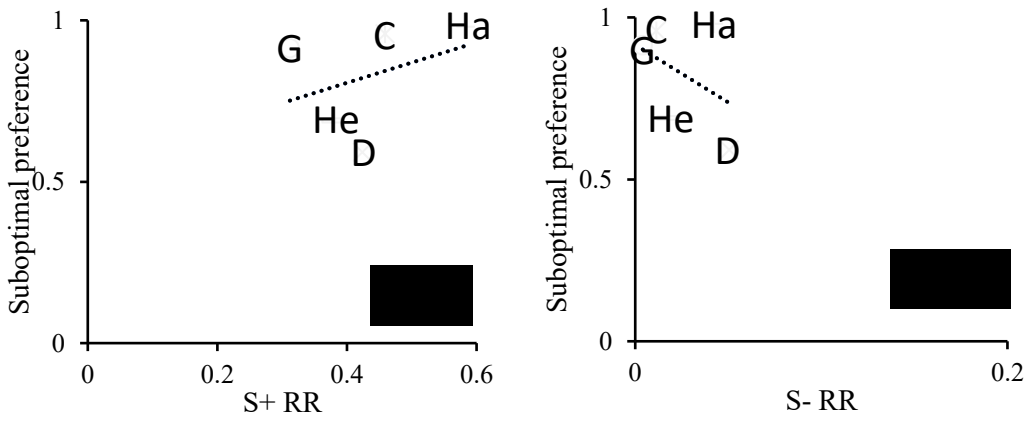

Note. Dashed lines represent the best linear fitting. r = Pearson's correlation. Symbols represent individual pigeon names: C = Cousteau, D = Darwin, G = Goodall, Ha = Hawthorne, He = Herriot.