Published in final edited form as:

Cogn Affect Behav Neurosci. 2014 June ; 14(2): 729–741. doi:10.3758/s13415-013-0220-4.

Training Attention Improves Decision Making in Individuals with Elevated Self-Reported Depressive Symptoms

Jessica A. Cooper^{a,c}, Marissa A. Gorlick^{a,c}, Taylor Denny^a, Darrell A. Worthy^b, Christopher G. Beevers^{a,c}, and W. Todd Maddox^{a,c,d,e,*},1

^aUniversity of Texas at Austin Psychology Department

^bTexas A&M Psychology Department

^cUniversity of Texas at Austin Institute for Mental Health Research

^dUniversity of Texas at Austin Institute for Neuroscience

^eUniversity of Texas at Austin Center for Perceptual Systems

Abstract

Depression is often characterized by attentional biases toward negative items and away from positive items, which likely affects reward and punishment processing. Recent work reported that training attention away from negative stimuli reduced this bias and reduced depressive symptoms. However, the effect of attention training on subsequent learning has yet to be explored. In the current study, participants were required to learn to maximize reward during decision-making. Undergraduates with elevated self-reported depressive symptoms received attention training toward positive stimuli prior to performing the decision-making task ($n=20$; active training). The active training group was compared to two groups: undergraduates with elevated self-reported depressive symptoms who received placebo training ($n=22$; placebo training) and control subjects with low levels of depressive symptoms ($n=33$; non-depressive control). The placebo-training depressive group performed worse and switched between options more than non-depressive controls on the reward maximization task. However, depressives that received active training performed as well as non-depressive controls. Computational modeling indicated that the placebo-trained group learned more from negative than from positive prediction errors, leading to more frequent switching. The non-depressive control and active training depressive groups showed similar learning from positive and negative prediction errors, leading to less frequent switching and better performance. Our results indicate that individuals with elevated depressive symptoms are impaired at reward maximization, but that the deficit can be improved with attention training toward positive stimuli.

Keywords

depression; decision-making; computational modeling; reflexive processing; reward; punishment

*Correspondence concerning this article should be addressed to W. Todd Maddox, Department of Psychology, The University of Texas, 108 E. Dean Keeton Stop A8000, Austin, TX 78712. Maddox@psy.utexas.edu.

¹This research was funded by NIDA grant DA032457 to WTM and CGB, and AFOSR grant FA9550-11-C-0028 to WTM. The first author was supported by a National Defense Science and Engineering Graduate (NDSEG) Fellowship. We thank Seth Koslov for help with data collection.

Depression is a common, recurrent, and impairing condition that predicts negative life events including future suicide attempts, interpersonal problems, unemployment, and substance abuse (Kessler & Walters, 1998; Kessler et al., 2003). The World Health Organization reports that approximately 120 million people currently suffer from depression and many more have elevated depressive symptoms. Further, adults with elevated depressive symptoms, even in the absence of Major Depressive Disorder, have poor physical, social, and role functioning compared to a demographically similar group without a chronic health condition. The well-being and psychosocial functioning of individuals with elevated depressive symptoms is comparable to that of people with major chronic medical conditions, such as hypertension, diabetes, and arthritis (Wells & Trust, 1989).

Cognitive theories of depression (e.g. Beck, 1976; Teasdale, 1988) argue that a contributing factor to depression is an attentional bias for depression-relevant themes. Depressed individuals focus attention on negative self-referent thoughts and exhibit enhanced effortful recall of negatively valenced material (Mathews & MacLeod, 2005). Supporting this, several studies have documented a negative attentional bias in depression (Mogg & Bradley, 2005). In addition, depression is also associated with the absence of a positive attentional bias. That is, non-depressed individuals typically have an attentional bias toward positive stimuli relative to neutral stimuli, which is often lacking among depressives (Ellis, Beevers, & Wells, 2011; Gotlib & Krasnoperova, 1998; Sears, Thomas, LeHuquet, & Johnson, 2010). Thus, depression involves biased attention toward negative stimuli and the absence of a bias toward positive stimuli.

Depression is also associated with decreased sensitivity to reward (i.e. Henriques, Glowacki, & Davidson, 1994; Pizzagalli et al., 2009) and decision making deficits (Beevers et al., 2013; Maddox, Gorlick, Worthy, & Beevers, 2012; Pizzagalli, Iosifescu, Hallett, Ratner, & Fava, 2008). We hypothesize that attentional biases and decision-making deficits are not independent, and that depressives' deficits in decision-making may be a consequence of an attentional bias toward negative stimuli (Gotlib & Joormann, 2010). That is, depressives' hypersensitivity to punishment and biased attention toward negative stimuli may cause them to respond sub-optimally in decision-making tasks (Eshel & Roiser, 2010). Thus, if depressives' attentional bias can be alleviated, decision-making should also improve. The goal of the current study is to test this possibility by experimentally manipulating attention bias toward positive items before a decision making task. We predict training will attenuate deficits in reward-based decision making that have been observed in individuals with elevated depressive symptoms.

Reward Processing and Decision-Making Deficits in Depressives

As noted above, depression is associated with decreased sensitivity to rewarding stimuli (Berenbaum & Oltmanns, 1992; Gotlib & Joormann, 2010; Henriques et al., 1994; Henriques & Davidson, 2000). Individuals with elevated depressive symptoms exhibit attenuated behavioral responsiveness to monetary gains (Henriques & Davidson, 2000), but hypersensitivity to negative feedback and punishment (Eshel & Roiser, 2010). In addition, those with depression show significantly weaker functional responses to monetary rewards

in the bilateral caudate and left nucleus accumbens, areas of the brain involved in immediate feedback processing, but these differences were not seen in the neutral or punishment conditions (Pizzagalli et al., 2009).

Depressives' deficits in reward processing may contribute to deficits in decision-making tasks. In two recent studies, we found that individuals with elevated depressive symptoms showed impaired decision-making when asked to maximize gains, but showed enhanced decision-making when asked to minimize losses (Beevers et al., 2013; Maddox et al., 2012). Other studies have found similar results. For example, Kunisato et al. (2012) found that depressives have reward-based decision-making deficits and respond with more variable action selection than non-depressed participants. Using a probabilistic reward task, Pizzagalli et al. (2008) found that depressives were responsive to delivery of single rewards, but were impaired at integrating reinforcement history over time.

We propose that observed decision-making deficits are due to depressive's negative attentional biases that may undermine learning from rewards. Decision-making is a complex process, but one critical aspect of decision-making, in healthy individuals and clinical populations alike, is reward-based learning (i.e. Ridderinkhof & van den Wildenberg, 2004). One necessary aspect of reward-based learning is acquisition of knowledge, implicit or explicit, about the relationships between stimuli and actions (Berridge & Robinson, 2003). Thus, individuals must learn from available rewards and use this information to make decisions.

We hypothesize that because depressed individuals pay more attention to negative stimuli than positive stimuli, they will learn more from negative information than positive information. Non-depressed individuals, in contrast, should show the opposite pattern. This explanation is consistent with our prior work where we found enhanced punishment processing and reduced reward processing in individuals with elevated depressive symptoms (Beevers et al., 2013; Maddox et al., 2012). It follows that attenuating these individuals' negative attentional bias should, in turn, enhance reward-based decision-making performance. We test this by manipulating attention to positive stimuli.

Attention Training

Recent research suggests that biases in attention toward emotion stimuli is malleable and can be altered with targeted training. A seminal study by MacLeod and colleagues (2002) found that attention training could create a negative bias in healthy individuals who did not initially possess such a bias and that such training could lead to a greater negative response to a laboratory stressor. Similarly, healthy individuals trained to direct attention toward positive stimuli spent significantly less time viewing negative images during a subsequent visual stress task, thus demonstrating a learned aversion to negative stimuli (Wadlinger & Isaacowitz, 2008).

This work with healthy individuals suggests that attention training can causally mediate emotional vulnerability and that modifying selective information processing may have potential therapeutic value (Hakamata et al., 2010; Hallion & Ruscio, 2011). Building upon this work, Wells and Beevers (2010) used a variant of the dot probe task (MacLeod,

Mathews, & Tata, 1986) to train depressed individuals to shift attention away from negative images. During training, participants viewed negative and neutral image pairs followed by a dot probe. In the active training condition, neutral images predicted the probe location 85% of the time, whereas in the placebo training control condition, neither image type (neutral or negative) predicted the probe location. Depressed individuals in the active attention training condition showed a reduced bias toward negative items and reported significant reductions in depressive symptoms two weeks post-training compared to control participants (see also Baert, De Raedt, Schacht & Koster, 2010)

Present Research

In the present research we use attention training to direct the attention of individuals with elevated depressive symptoms toward positive information in order to improve reward-based decision-making. During training, participants viewed positive and neutral word pairs followed by a dot probe. In the active attention training condition, positive words predicted the probe location 85% of the time, whereas in the placebo training condition, neither word type (neutral or positive) predicted the probe location. Thus, the training paradigm implicitly trained participants to shift their attention toward positive information.

Following training, participants completed a decision-making task similar to that used in Beevers et al (2013). One concern about the decision-making task used in Beevers et al. was that the reward values for each option were too dramatically different, making the task too easy and subsequently diminishing performance differences between individuals with elevated depressive symptoms and non-depressives. Thus, the present study increased the variability of each option. This increased variability in the reward structure makes the task more challenging by requiring participants to learn the value of each option by taking rewards over several trials into consideration, as opposed to responding only to the most recently received reward. This is important because previous research suggests that depressed individuals are impaired at integrating reward history over time (Pizzagalli et al., 2008). Supporting this idea, our prior work (Beevers et al., 2013) found that individuals with elevated depressive symptoms were more likely to alter their expected values for each option based on *recently* received rewards, whereas non-depressed control participants relied on a longer sequence of previous rewards in determining expected reward value.

Taken together, we make the following predictions. First, we predict that individuals with elevated depressive symptoms who received placebo attention training will show a decision-making performance deficit relative to non-depressive controls. Second, we predict that individuals with elevated depressive symptoms who received active attention training toward positive stimuli will show enhanced performance compared to those in the placebo training condition and comparable performance to non-depressive controls.

Method

Participants

Participants were 92 newly recruited undergraduate students who completed the study as a part of a research requirement for an introduction to psychology course. Participants

completed the short form of the Beck Depression Inventory (BDI; Beck & Steer, 1993) during a pre-testing survey battery. Participants whose scores were above 7 on the BDI-SF were contacted about participating in the depressive groups, while participants whose scores were below 7 on the short form were contacted about participating in the non-depressive control group.

Procedure

At the beginning of the experimental session all participants completed a demographic form and a series of computer-based questionnaires that included the Center for Epidemiological Studies Depression Scale (CES-D; Radloff, 1977). The CES-D scores were used to verify that participants were still experiencing depressive symptoms at the time of testing and to validate previously recorded BDI scores. Participants with elevated levels of self-reported depressive symptoms were randomly assigned to the active training or placebo training groups. Participants were not told that depression was a measure of interest and were not told anything about their group membership. Research assistants were also blind to participant assignment to placebo or active training conditions. It is also important to note that the term “training” was not mentioned to the participants—Participants were not given any information that they might be undergoing attention training until they were debriefed at the end of the experiment. This ensured that the participants were unaware of the study’s purpose as well as experimental condition.

For the placebo training and active training groups, participants completed three blocks (all in the same session) of the attention-training task, approximately fifteen minutes each. Each block consisted of 168 trials for a total of 504 trials. Before and after the three-block training session participants completed a two-item questionnaire about their current mood. Participants were given two minutes between each block to relax before starting the next block. Immediately following training participants completed the decision making task. The low-depressive symptom control group did not undergo training and completed the decision-making task immediately following completion of demographic information and questionnaires.

Depression Classification—Following convention (Weissman & Sholomskas, 1977), participants who scored 15 or less on the CES-D were classified as having low depressive symptoms, and participants who obtained a 16 or greater were classified as having elevated depressive symptoms. CES-D scores of 16 or greater reflect moderate or greater symptoms of depression (Radloff, 1977). A cut-point of 16 on the CES-D has very good sensitivity and specificity for the prediction of current major depressive disorder (Beekman et al., 1997). Participants were only included in the analysis if their CES-D score (16 or higher for the elevated depressive symptom groups and 15 or lower for the non-depressive group), was consistent with their classification from the previously recorded BDI-SF score, resulting in 75 subjects: $n=20$ elevated depressive symptom participants in the active attention training group, $n=22$ elevated depressive symptom participants in the placebo training group, and $n=33$ in the non-depressive control group.

Attention Training—This task was designed to train participants' attention toward positive information (i.e., words) using a modified dot-probe paradigm. The task used neutral and positive valence words from the Affective Norms for English Words list (ANEW; M. M. Bradley & Lang, 1999). Words were matched for letter length and frequency use in the English language. Therefore, the only differences between the list of positive and neutral words were valence and arousal.

The attention-training task is displayed in Figure 1a. Each trial of the task began with a 500ms presentation of a fixation-cross. Following the cross, a pair of stimuli, a positive word and a neutral word, were randomly presented to the right and the left side of the computer screen for 1000ms. The words then disappeared and a dot-probe (i.e., * or **) appeared behind one of the previously displayed words. This probe appeared on the screen until the participant pressed one of two response buttons to indicate the identity of the probe (1 or 2 asterisks'). In the active training condition, the probe was presented in the location associated with the positive word on 85% of the trials and the neutral word on 15% of the trials. In the placebo training condition, the probe was presented in the location of the neutral word on 50% of the trials and in the location of the positive word on 50% of the trials. In both conditions, the positive word appeared randomly and equally on either side of the screen.

We selected the stimulus duration of 1000ms based on previous research (B. P. Bradley, Mogg, & Lee, 1997) that found that levels of depressive symptoms on the BDI were strongly correlated with attentional bias score for valenced words presented for 1000ms. Consistent with previous research (T. T. Wells & Beevers, 2010) we used 85% positive rather than 100% in the training condition to keep the intent of the study from being transparent.

Decision-Making Task—The decision making task was performed on PC computers using Matlab software with Psychtoolbox 2.54 (Brainard, 1997; Pelli, 1997). Participants were given a hypothetical scenario that they would be testing two oxygen extraction systems on Mars with the goal of collecting enough oxygen to sustain life. Participants were told that on each trial they would select one of the two systems. After each selection a bar, representing a small oxygen tank, would show the amount of oxygen that they had just extracted. The oxygen would then be moved into the larger tank and the next trial would begin. A line on the larger tank corresponded to the amount of oxygen needed to sustain life on Mars. Participants were given the goal of trying to collect this amount of oxygen over the course of the experiment. Figure 1b shows a sample screen shot from the decision making experiment. The goal line was set at the equivalent of selecting the optimal option on approximately 80% of trials. Participants performed a total of 150 trials, and were told nothing about the nature of the reward structure.

The reward structure for the decision-making task is shown in Figure 2. In this two-option task, the average reward for the sub-optimal option is 55 points (units of oxygen), while the average reward for the optimal option is 65 points. However, because the standard deviation around the mean reward for each option is 15 units it is not the case that the optimal option yields a larger reward on every trial.

Results

The mean CES-D scores for the elevated depressive symptom active training and placebo training groups were 35.20 and 33.23, respectively. These did not differ significantly, $t(40)=.800$, $p=.429$, Cohen's $d=.252$. The mean CES-D score for the non-depressive control group was 7.70, significantly lower than both the active training and placebo training groups ($p's<.001$).

Performance in the decision-making task was measured by analyzing the proportion of trials on which participants selected the optimal choice (the higher average reward) throughout the experiment. We used one-sample t-tests to determine whether performance significantly exceeded chance. Performance exceeded chance in the active attention training group ($M=.69$), $t(19)=6.527$, $p<.001$, the placebo training group ($M=.61$), $t(21)=3.118$, $p=.005$, and the non-depressive control group ($M=.72$), $t(32)=7.595$, $p<.001$. Thus, all groups learned to select the optimal choice.

We next used an ANOVA to examine the effect of group (placebo training, active training, non-depressive control) on the proportion of optimal selections (Figure 3a). We observed an overall effect of group, $F(2,72)=3.786$, $p=.027$, partial- $\eta^2=.095$. There was also a significant linear contrast, $F(1,72)=7.472$, $p=.008$, partial- $\eta^2=.094$. As predicted, the participants who received placebo training selected the optimal choice less often than the non-depressive control group, $t(53)=-2.573$, $p=.013$, Cohen's $d=.707$. Also as predicted, the participants who received placebo training selected the optimal choice less often than the active attention training group $t(40)=-1.828$, $p=.06$, Cohen's $d=.578$. There was not a statistically significant difference in performance across the attention trained participants with elevated depressive symptoms and the non-depressive control group, $t(51)=-.786$, $p=.436$, Cohen's $d=.220$. A post hoc power analysis revealed that on the basis of the between-groups comparison effect size observed in the present study ($d=.220$), an N of approximately 648 would be needed to obtain statistical power at the recommended .80 level (Cohen, 1988) between the active attention training and non-depressive groups.

Optimal performance in this task relies on exploitation of the choice option that awards the highest average reward. We hypothesized that the observed performance deficit in depressives may be attributable to sub-optimal switching between options based on fluctuations in the reward environment. To explore this possibility, we examined the number of trials in which participants switched between reward options within each of the three participant groups (Figure 4a). Using an ANOVA, we found a significant effect of group membership on the number of switch trials $F(2,72)=3.181$, $p=.047$, partial $\eta^2=.081$. The placebo training group ($M=48.45$) switched between options significantly more often than the non-depressive control group ($M=33.00$), $t(53)=2.515$, $p=.015$, Cohen's $d=.690$. The active attention training group ($M=39.15$) was intermediate between the placebo-trained group and the non-depressive group. Consistent with these findings, there was a significant linear contrast for depression group, $F(1,72)=6.362$, $p=.014$, partial- $\eta^2=.081$, suggesting that switching decreased monotonically across the placebo training, active training, and non-depressive control groups. Across all participants, number of switches was negatively correlated with performance on the task $r(74)=-.453$, $p<.001$.

We also examined ‘delta reward values’ on trials that preceded switch trials. To calculate a delta reward value for each trial we compared the reward received on that trial to the most recently received reward for the selected option. Negative values indicate smaller rewards relative to the reward received the last time that option was chosen, while positive values indicate greater rewards relative to the last reward received for that option. On average, the placebo training group switched to the other option after a delta reward of -3.65 points, while the non-depressive control and active training groups switched after an average delta reward value of -4.48 and -5.89 points, respectively. Thus, although these differences were not statistically significant ($p > .1$), participants in the placebo-trained group had a numerically smaller threshold for switching following declines in reward relative to what they had previously received when they selected the same option.

The subjective measures of mood, collected before and after the training procedure, were also analyzed. Participants rated their current mood on two nine-point scales, from ‘not happy at all’ to ‘very happy’, and from ‘not sad at all’ to ‘very sad’. Pre and post mood scores were compared for both the placebo training and active attention training groups using paired t-tests. The pre and post mood scores were not significantly different for either question for either group ($p > .1$), indicating that changes in decision-making performance were not attributable to conscious changes in mood.

We also examined correlations between CES-D scores and the proportion of optimal selections in the decision-making task. Across participants who did not receive active attention training (non-depressive controls and placebo training with elevated depressive symptoms) CES-D scores were negatively correlated with the proportion of optimal selections $r(53) = -.296, p = .03$. Interestingly, this trend was present in the non-depressive control group alone, $r(31) = -.339, p = .05$, indicating that depressive symptoms may be negatively related to decision-making performance even at levels below the standard CES-D cutoff for depression. Correlations within both the depressive active training and placebo training groups were not significant ($p > .1$).

Computational Modeling

The behavioral results suggest that attention training attenuates the decision-making deficits observed in individuals with elevated self-reported depressive symptoms. To better understand the strategies that participants used to make decisions in the task, we fit a series of computational models, including a baseline model, a basic Reinforcement Learning (RL) model with a Softmax decision rule (Beevers et al., 2013; Frank & Kong, 2008; Lee, Zhang, Munro, & Steyvers, 2011; Steyvers, Lee, & Wagenmakers, 2009; Sutton & Barto, 1998; Worthy & Maddox, 2011) and a novel extension of the RL model that accounts separately for positive and negative prediction errors.

The basic Softmax RL model accounts for decision-making behavior by updating expected reward values (EVs) for each option, i , on each trial, t , in Equation 1. Expected values (EVs) are initialized at zero for each option and are updated based on prediction error δ :

$$\delta = r(t) - EV_{i,t} \quad (1)$$

Here $r(t)$ is the reward received for the chosen option. Prediction errors are used to update EVs each time an option is chosen based on the following update rule:

$$EV_{i,t+1} = EV_{i,t} + \alpha \cdot \delta \quad (2)$$

The recency parameter (α), $0 < \alpha < 1$, weighs the degree to which participants update the expected values for each option based on their most recently received rewards. As α approaches 1, recent rewards are given greater weight in updating EVs. An α of zero indicates that no learning took place and EVs were not updated from their initial starting point. The expected values for each option are used to determine the probabilities for selecting each option by the Softmax decision rule (Sutton & Barto, 1998):

$$P(a_i, t) = \frac{e^{\theta \cdot EV(a_i, t)}}{\sum_{j=1}^2 e^{\theta \cdot EV(a_j, t)}} \quad (3)$$

Where θ is an exploitation parameter of the degree to which the option with the highest EV is chosen. Higher values of θ indicate that the highest valued option is chosen more often, while a value of zero indicates that they are chosen equally often.

Depressives show enhanced punishment processing but deficient reward processing (i.e. Eshel & Roiser, 2010). Punishments are directly associated with negative prediction errors and rewards with positive prediction errors. To explore positive and negative reward prediction errors mechanistically we extended the model to account separately for positive and negative prediction errors. We let the model freely estimate two learning rate parameters that were used to update EVs in Equation 2 when prediction errors were positive (α_{pos}) or negative (α_{neg}). This rule allows the model to track the recency-weighted average rewards provided by each option.

A critical aspect of this standard RL model is that it assumes that on any given trial EVs for each option are represented by a single numerical value. An alternative assumption is that EVs are represented in the form of distributions around a mean value, rather than an exact single value. This assumption has been highlighted in recent RL and associative learning work (Doll, Jacobs, Sanfey, & Frank, 2009; Frank, Doll, Oas-Terpstra, & Moreno, 2009; Kruschke, 2008), and is likely more realistic in environments with variable rewards provided on each trial.

The variance, or noise around the mean of each distribution of expected values, could indicate the level of uncertainty people have as to what reward they will actually receive. Here we develop a model that assumes that the mean EV for the distribution of EVs is a recency-weighted average of past rewards for each option, as updated in Equation 1, and that Noise (N) around each mean is a recency-weighted average of the squared prediction errors on each trial:

$$N_{i,t+1} = N_{i,t} + \alpha_N \cdot [\delta^2 - N_{i,t}] \quad (4)$$

The degree to which the noise estimates are updated based on the most recent prediction errors is modulated by a recency parameter (α_N), $0 \leq \alpha_N \leq 1$, similar to how EVs are updated in Equation 2 above. Here, we also fit the model with separate α_N parameters for trials with positive $\alpha_{N(pos)}$ and negative $\alpha_{N(neg)}$ prediction errors. This approach is similar to recent approaches that have approximated noise, or the variance in the most likely outcome for each choice option by tracking the variance in recent outcomes (Doll et al., 2009; Nassar, Wilson, Heasly, & Gold, 2010).

The initial N estimate for each option (N_0) is a free parameter in the model that represents the initial uncertainty participants have regarding the average reward provided by each option. The Extended RL model includes this noise term in the Softmax rule to account for behavior where participants select options that they have greater or lesser uncertainty about:

$$P(a_t) = \frac{e^{[\theta \cdot EV_{a,t} + \theta_N \cdot N_{a,t}]}}{\sum_{j=1}^2 e^{[\theta \cdot EV_{j,t} + \theta_N \cdot N_{j,t}]}} \quad (5)$$

As in the Basic RL model we set a minimum value of 0 on θ , but we allowed θ_N to be positive or negative. Positive values indicate a preference for options that have greater uncertainty, while negative values indicate a preference for options with lesser uncertainty. Thus, the addition of the noise term allows us to account for behavior where participants attempt to reduce uncertainty regarding the noise around each expected reward value.

In sum the Extended RL model has seven free parameters: α_{pos} , α_{neg} , $\alpha_{N(pos)}$, $\alpha_{N(neg)}$, N_0 , θ , and θ_N , while the Basic RL model has two free parameters: α , and θ . We compared the relative fits of these models using the Akaike Information Criterion (AIC) for each model, which rewards goodness of fit but also includes a penalty for increasing the number of free parameters (Akaike, 1974). This was done to ensure that the increase in relative fit of the model was not outweighed by the flexibility of additional parameters.

The final model that we fit, baseline (or null) model, assumes fixed choice probabilities (Gureckis & Love, 2009; Worthy & Maddox, 2011; Yechiam & Busemeyer, 2005). The baseline model has one free parameter that represents the probability of selecting one of the two options on any given trial. This model does not assume that participants learn from rewards given on each trial, yet it provides a good fit for data when participants repeatedly choose the same option (Gureckis & Love, 2009).

Model-based Predictions

There are two clear ways in which models could capture the performance differences between groups. One possibility is that all groups will be best fit by the basic Softmax RL model, but with different parameter estimates indicating that learning is occurring at different rates. Specifically, we would expect to see higher variability in action choice in the placebo-trained group, represented by a lower exploitation parameter estimate, consistent with Beevers et al. (2013) and Kunisato et al. (2012).

A second possibility, which we believe to be the most likely, is that the observed behavioral differences can be attributed to increased attention to negative information in the placebo-trained group with elevated depressive symptoms, and that increased attention to negative information results in increased learning when prediction errors are negative, ultimately causing over-adjustment of the expected value for that option and sub-optimal switching. In the extended RL model negative and positive reward prediction errors update the expected values for each option with separate learning rates (recency parameters), allowing the model to account for differential learning from positive and negative prediction errors. Consistent with depressive's hypersensitivity to negative feedback and punishment (i.e. Eshel & Roiser, 2010), we would expect depressives' increased attention to negative information to result in increased learning rates for negative prediction errors. As attention training is thought to reduce this bias, we would expect to see reduced learning rates (lower recency parameter) for negative prediction errors in the active attention training group relative to the placebo training group.

Modeling Results—We fit each participant's data individually with the models detailed above. The models were fit on a trial-by-trial basis to the participant's response and the parameters were estimated by maximizing likelihood. We used Akaike Weights (Wagenmakers & Farrell, 2004) to compare the relative fits of the models. The Akaike Weights are derived from Akaike's Information Criterion (AIC: Akaike, 1974), which is defined for each model i as:

$$AIC_i = -2\text{Log}L_i + 2V_i \quad (6)$$

where L_i is the maximum likelihood for model i and V_i is the number of free parameters in the model. Notice that the AIC measure penalizes the model for each additional free parameter.

The AIC values were used to generate the Akaike Weight for each of the three models for each participant. The relative likelihood, L , of each model, i , is computed using the transform:

$$L(M_i|\text{data}) \propto \exp\left\{-\frac{1}{2} \cdot \Delta_i(\text{AIC})\right\} \quad (7)$$

where $\Delta_i(\text{AIC})$ represents the difference between the AIC for that model and the lowest AIC of all candidate models. The relative likelihoods of each candidate model are then normalized by dividing each of the likelihoods by the sum of all likelihoods for all k models:

$$w_i(\text{AIC}) = \frac{\exp\left\{-\frac{1}{2} \cdot \Delta_i(\text{AIC})\right\}}{\sum_{k=1}^K \exp\left\{-\frac{1}{2} \cdot \Delta_k(\text{AIC})\right\}} \quad (8)$$

These Akaike weights can be interpreted as the probability that the model is the best model for the data given the data set from the set of candidate models (Wagenmakers & Farrell, 2004).

We computed Akaike weights for each model for each participant. Figure 5 shows the average Akaike weights for each condition. The Extended RL model was clearly the best fitting model for all groups.

Extended RL Model Results

We used simulations to verify that the extended RL model provided a good account for the observed data. For each group, 1000 simulations were conducted by sampling (with replacement) a set of parameters from one subject in that group for each simulation.

The results of the simulation data were strikingly similar to the behavioral data (Figure 3b), indicating that the model provides a good account to the behavioral data. We also examined the frequency of switch trials from the simulation data, which were again similar to the switch trials observed in the behavioral data (Figure 4b). This suggests that model parameters are related to behavior observed in this study.

We examined the parameters of interest, the learning rates for negative reward prediction errors and positive reward prediction errors (Figure 6). We observed no main effect of group within the learning rate in positive prediction errors ($p > .1$); however, we observed a marginally significant main effect of group in learning from negative prediction errors $F(2,72)=2.910$, $p=.061$, and a significant linear contrast, $F(1,72)=5.755$, $p=.019$, partial- $\eta^2=.074$. In decomposing this effect we found that placebo trained group with elevated depressive symptoms had a significantly higher learning rate from negative prediction errors than non-depressives, $t(53)=2.441$, $p=.018$, Cohen's $d=.670$; the same learning rate parameter for the attention trained group was intermediate between these two groups.

We then compared the estimated parameter values for learning from negative prediction errors and positive prediction errors within each of the three groups. For the placebo training group, the learning rate (recency parameter α) for negative prediction errors was marginally higher than the learning rate for positive prediction errors, $t(42)=1.887$, $p=.066$, Cohen's $d=.582$. The negative and positive learning rates were not different between the other two groups, $p's > .1$, indicating that of the three groups, only the placebo training group may have learned more from negative prediction errors than from positive prediction errors.

In addition to providing a good fit for the behavioral data, simulations using the extended RL model accounted for the differences in the number of trials on which participants switched between options that we observed across the three groups. In summary, participants in the placebo training group learned more from negative prediction errors and switched to the sub-optimal choice following negative reward prediction errors more often than non-depressives. The group with elevated depressive symptoms that received active training and the non-depressive group did not show a difference between learning from negative and positive reward prediction errors, switched less between options, and performed better on the task.

Discussion

Individuals with major depression and elevated depressive symptoms face physical, social and cognitive challenges. The presence of depressive symptoms is associated with increased attention to negative information and decreased attention to positive information. This reward processing deficit has broad effects on cognition, including one's ability to make good decisions (Beevers et al., 2013; Kunisato et al., 2012; Maddox et al., 2012; Pizzagalli et al., 2008). Recent research suggests that these reward processing deficits may be malleable and can be attenuated with attention training (T. T. Wells & Beevers, 2010).

The overriding aim of the present study was to determine whether an attention training procedure that was successful in ameliorating depressive symptoms might be used to attenuate the prevalent deficit in decision-making. In line with previous research (Beevers et al., 2013; Kunisato et al., 2012; Maddox et al., 2012), we found that individuals with elevated depressive symptoms performed worse than non-depressives in maximizing rewards. Importantly, this performance deficit was significantly attenuated when given attention training toward positive stimuli prior to performing the decision-making task. In fact, performance of the active attention-trained group was no different from that of non-depressives.

Computational modeling was utilized to better understand the strategies being used by each group to solve the decision-making task. Analysis of the best-fitting model parameters suggest that of the three groups, the placebo-trained group had the highest learning rate for negative prediction errors and the greatest discrepancy in learning from negative prediction errors relative to positive prediction errors.

These findings indicate that individuals with elevated depressive symptoms place too much weight on immediate rewards, particularly when these rewards are less than expected, resulting in overcorrection of their expected value for the reward option. This is consistent with previous research demonstrating that people with elevated depressive symptoms were more likely to alter their expected values based on recently received rewards (Beevers et al., 2013). In our study, this overcorrection in response to negative prediction errors resulted in increased frequency of sub-optimal switching relative to non-depressives.

Training attention toward positive information decreased learning from negative prediction errors and decreased the frequency of switching to the sub-optimal choice, improving task performance. Importantly, performance of the active-trained group was not different from the non-depressive group. This is an exciting finding as it suggests that attention training can reduce depressives' bias toward negative information, which often leads to poor decision-making. Individuals with elevated depressive symptoms who experienced training had numerically lower frequency of switching and learning rates for negative prediction errors than those who did not receive training. Future work will explore the effect of longer-term training to see if repeated training further improves decision making and persists over time.

It should be noted that the performance difference between depressives and non-depressives in this task were much greater than the performance difference observed in Beevers et al. (2013). This may be attributable to the different reward structures underlying the previously

published paper and the current study, although a direct empirical comparison would be required to draw any strong conclusions. The task in the previous paper had a clear separation between the optimal and suboptimal choices, where the lowest reward for the “good” option was still higher than the highest reward from the “bad” option. Thus, because the rewards never overlap, the task was relatively easy and few performance or strategy differences were observed between the groups. These findings are consistent with the idea that depression has less of an effect on easier, automatic processing, but has a greater effect on more demanding tasks (Hartlage, Alloy, Vázquez, & Dykman, 1993; Hertel, 1994).

Limitations and Future Directions

One implication of this work is that depressive individuals do have the cognitive processing capabilities to perform well in a gain-maximizing task, and this ability can be enhanced with training. Attention training during a single session reduced the reward-learning deficit, likely by reducing learning from negative (punishment-based) prediction errors. Future work would benefit from utilization of punishment-based tasks to determine whether attention training increases reward learning or decreases punishment-based learning. Future work should also evaluate the effect of attention training on performance in history-dependent tasks for which reduced sensitivity to negative reward predictions errors may not be beneficial. While it is possible that attention training may improve performance on these tasks, it is also possible that attention training enhances the use implicit, reflexive strategies that may not benefit all types of reward-maximization tasks.

As similar training has also been shown to reduce reported depressive symptoms when completed four times over the course of two weeks (T. T. Wells & Beevers, 2010), future work should also evaluate the effect of longer-term training on protracted cognitive abilities and experience of depression symptoms. Decision-making deficits were attenuated with short-term training in this study but did not change mood. As the long-term training has been shown to reduce depressive symptoms, altering mood, there is reason to think that decision-making performance could be further improved. Long-term training efforts should seek to evaluate changes in reward sensitivity, depression, and cognitive abilities.

In the current work, the non-depressive group demonstrated a numerically greater but non-significant bias for learning due to negative over positive prediction errors. Thus, it is unclear how this training paradigm would affect those with normative emotional processing. It is worth noting that task performance varied within the non-depressive group. On average, performance was better for the non-depressive than the placebo-trained depressive group. However, some non-depressive individuals exhibit difficulty with reward maximization and could have benefited from attention training toward positive stimuli. It is therefore possible that attention training might also be helpful for decision-making in healthy individuals with relatively low reward processing. It is also possible that non-depressives would develop a strong bias for learning from positive prediction errors that could impair task performance. Future work should apply these training paradigms to non-depressive individuals to determine the impact of attention training on subsequent learning in healthy adults.

Several limitations of the current study should be noted. First, we did not complete diagnostic interviews and thus do not have extensive information regarding past

psychopathology, use of psychoactive substances or medications, or neurological diseases. Although all participants in the groups with elevated depressive symptoms exceeded a cut-point on the CESD-D commonly used to screen for major depressive disorder (Radloff, 1977), it is likely that some participants would not have met criteria for a major depressive episode. Thus, it is unknown whether clinically depressed individuals would perform similarly in the decision-making task used in this study, and whether they would be affected by attention training in the same way. We also do not know if the findings of this study will generalize to other decision-making tasks, including those for which actual monetary gains and losses are utilized. Future studies using clinically diagnosed patients and multiple decision-making tasks should be used to answer these questions.

Another limitation of the current study is that independent measures of attention were not collected before and after training. Obtaining an independent measure of attention before and after training (ideally with new stimuli) could provide additional evidence that the changes in decision-making performance observed in the current study were due to changes in attention rather than another underlying factor (for an example, see Amir et al., 2009). Finally, although our findings are consistent with previous work that identifies decision-making deficits in individuals with elevated levels of depressive symptoms (Beevers et al., 2013; Kunisato et al., 2012; Maddox et al., 2012; Pizzagalli et al., 2008), it is important to note that replication is a critical pillar of the scientific method. Given our small sample size, replication of these findings in future studies will be critical to increasing our confidence in the robustness of our experimental effects.

Conclusions

This work provides important new insights into the effects of attention training on decision-making performance in individuals with elevated depressive symptoms. Attention training can be used to attenuate depressives' deficits in gain maximization to a point where performance does not differ from non-depressives. Computational modeling suggests that training attention toward positive stimuli enhances performance by reducing a bias for learning from negative prediction errors, resulting in decreased sub-optimal switching and increased performance. Thus, training toward positive stimuli can be used to improve decision-making in individuals experiencing elevated levels of depressive symptoms.

References

- Akaike H. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*. 1974; 19(6):716–723.10.1109/TAC.1974.1100705
- Amir N, Beard C, Taylor CT, Klumpp H, Elias J, Burns M, Chen X. Attention training in individuals with generalized social phobia: A randomized controlled trial. *Journal of Consulting and Clinical Psychology*. 2009; 77(5):961–973.10.1037/a0016685 [PubMed: 19803575]
- Baert S, De Raedt R, Schacht R, Koster EHW. Attentional bias training in depression: therapeutic effects depend on depression severity. *Journal of behavior therapy and experimental psychiatry*. 2010; 41(3):265–274.10.1016/j.jbtep.2010.02.004 [PubMed: 20227062]
- Beck, AT. *Cognitive therapy and the emotional disorders*. Oxford, UK: International Universities Press; 1976.
- Beck, AT.; Steer, RA. *Beck Depression Inventory: Manual*. San Antonio, TX: The Psychological Corporation; 1993.

- Beekman ATF, Deeg DJH, Limbeek J, Braam AW, De Vries MZ, Tilburg W. Criterion validity of the Center for Epidemiologic Studies Depression scale (CES-D): results from a community-based sample of older subjects in the Netherlands. *Psychological Medicine*. 1997; 27(1):231–235. [PubMed: 9122304]
- Beevers CG, Worthy DA, Gorlick MA, Nix B, Chotibut T, Maddox WT. Influence of depression symptoms on history-independent reward and punishment processing. *Psychiatry Research*. 2013; 207(1):53–60.10.1016/j.psychres.2012.09.054 [PubMed: 23122555]
- Berenbaum H, Oltmanns TF. Emotional experience and expression in schizophrenia and depression. *Journal of abnormal psychology*. 1992; 101(1):37–44.10.1037/0021-843X.101.1.37 [PubMed: 1537971]
- Berridge KC, Robinson TE. Parsing reward. *Trends in neurosciences*. 2003; 26(9):507–513.10.1016/S0166-2236(03)00233-9 [PubMed: 12948663]
- Bradley BP, Mogg K, Lee SC. Attentional biases for negative information in induced and naturally occurring dysphoria. *Behaviour research and therapy*. 1997; 35(10):911–927.10.1016/S0005-7967(97)00053-3 [PubMed: 9401132]
- Bradley, MM.; Lang, PJ. Affective norms for English words (ANEW): Instruction manual and affective ratings. Gainesville, FL: 1999.
- Brainard DH. The psychophysics toolbox. *Spatial vision*. 1997; 10(4):433–436. [PubMed: 9176952]
- Cohen, J. *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale, New Jersey: 1988.
- Doll BB, Jacobs WJ, Sanfey AG, Frank MJ. Instructional control of reinforcement learning: A behavioral and neurocomputational investigation. *Brain research*. 2009; 1299:74–94.10.1016/j.brainres.2009.07.007 [PubMed: 19595993]
- Ellis AJ, Beevers CG, Wells TT. Attention Allocation and Incidental Recognition of Emotional Information in Dysphoria. *Cognitive therapy and research*. 2011; 35(5):425–433.10.1007/s10608-010-9305-3
- Eshel N, Roiser JP. Reward and Punishment Processing in Depression. *Biological psychiatry*. 2010; 68(2):118–124.10.1016/j.biopsych.2010.01.027 [PubMed: 20303067]
- Frank MJ, Kong L. Learning to avoid in older age. *Psychology and Aging*. 2008; 23(2):392–398.10.1037/0882-7974.23.2.392 [PubMed: 18573012]
- Frank MJ, Doll BB, Oas-Terpstra J, Moreno F. Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature neuroscience*. 2009; 12(8):1062–1068.10.1038/nn.2342
- Gotlib IH, Joormann J. Cognition and Depression: Current Status and Future Directions. *Annual review of clinical psychology*. 2010; 6(1):285–312.10.1146/annurev.clinpsy.121208.131305
- Gotlib IH, Krasnoperova E. Biased information processing as a vulnerability factor for depression. *Behavior Therapy*. 1998; 29(4):603–617.10.1016/S0005-7894(98)80020-8
- Gureckis TM, Love BC. Learning in noise: Dynamic decision-making in a variable environment. *Journal of Mathematical Psychology*. 2009; 53(3):180–193.10.1016/j.jmp.2009.02.004 [PubMed: 20161328]
- Hakamata Y, Lissek S, Bar-Haim Y, Britton JC, Fox NA, Leibenluft E, et al. Attention bias modification treatment: a meta-analysis toward the establishment of novel treatment for anxiety. *Biological Psychiatry*. 2010; 68(11):982–990.10.1016/j.biopsych.2010.07.021 [PubMed: 20887977]
- Hallion LS, Ruscio AM. A meta-analysis of the effect of cognitive bias modification on anxiety and depression. *Psychological Bulletin*. 2011; 137(6):940–58.10.1037/a0024355 [PubMed: 21728399]
- Hartlage S, Alloy LB, Vázquez C, Dykman B. Automatic and effortful processing in depression. *Psychological bulletin*. 1993; 113(2):247–278.10.1037/0033-2909.113.2.247 [PubMed: 8451334]
- Henriques JB, Davidson RJ. Decreased responsiveness to reward in depression. *Cognition & Emotion*. 2000; 14(5):711–724.
- Henriques JB, Glowacki JM, Davidson RJ. Reward fails to alter response bias in depression. *Journal of abnormal psychology*. 1994; 103(3):460–466.10.1037/0021-843X.103.3.460 [PubMed: 7930045]
- Hertel PT. Depression and memory: Are impairments remediable through attentional control? *Current Directions in Psychological Science*. 1994; 3(6):190–193.

- Kessler RC, Walters EE. Epidemiology of DSM-III-R major depression and minor depression among adolescents and young adults in the national comorbidity survey. *Depression and anxiety*. 1998; 7(1):3–14.10.1002/(SICI)1520-6394(1998)7:1<3::AID-DA2>3.0.CO;2-F [PubMed: 9592628]
- Kessler RC, Berglund P, Demler O, Jin R, Koretz D, Merikangas KR, et al. The Epidemiology of Major Depressive Disorder: Results From the National Comorbidity Survey Replication (NCS-R). *JAMA*. 2003; 289(23):3095–3105.10.1001/jama.289.23.3095 [PubMed: 12813115]
- Kruschke JK. Bayesian approaches to associative learning: From passive to active learning. *Learning & Behavior*. 2008; 36(3):210–226.10.3758/LB.36.3.210 [PubMed: 18683466]
- Kunisato Y, Okamoto Y, Ueda K, Onoda K, Okada G, Yoshimura S, et al. Effects of depression on reward-based decision making and variability of action in probabilistic learning. *Journal of Behavior Therapy and Experimental Psychiatry*. 2012; 43(4):1088–1094.10.1016/j.jbtep.2012.05.007 [PubMed: 22721601]
- Lee MD, Zhang S, Munro M, Steyvers M. Psychological models of human and optimal performance in bandit problems. *Cognitive Systems Research*. 2011; 12(2):164–174.10.1016/j.cogsys.2010.07.007
- MacLeod C, Mathews A, Tata P. Attentional bias in emotional disorders. *Journal of abnormal psychology*. 1986; 95(1):15–20.10.1037/0021-843X.95.1.15 [PubMed: 3700842]
- MacLeod C, Rutherford E, Campbell L, Ebsworthy G, Holker L. Selective attention and emotional vulnerability: Assessing the causal basis of their association through the experimental manipulation of attentional bias. *Journal of Abnormal Psychology*. 2002; 111(1):107–123.10.1037/0021-843X.111.1.107 [PubMed: 11866165]
- Maddox WT, Gorlick MA, Worthy DA, Beevers CG. Depressive symptoms enhance loss-minimization, but attenuate gain-maximization in history-dependent decision-making. *Cognition*. 2012; 125:118–124. [PubMed: 22801054]
- Mathews A, MacLeod C. Cognitive vulnerability to emotional disorders. *Annual review of clinical psychology*. 2005; 1:167–195.
- Mogg K, Bradley BP. Attentional Bias in Generalized Anxiety Disorder Versus Depressive Disorder. *Cognitive therapy and research*. 2005; 29(1):29–45.10.1007/s10608-005-1646-y
- Nassar MR, Wilson RC, Heasley B, Gold JL. An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment. *The Journal of Neuroscience*. 2010; 30(37):12366–12378.10.1523/JNEUROSCI.0822-10.2010 [PubMed: 20844132]
- Pelli DG. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial vision*. 1997; 10(4):437–422. [PubMed: 9176953]
- Pizzagalli DA, Holmes AJ, Dillon DG, Goetz EL, Birk JL, Bogdan R, et al. Reduced Caudate and Nucleus Accumbens Response to Rewards in Unmedicated Individuals With Major Depressive Disorder. *American Journal of Psychiatry*. 2009; 166(6):702–710.10.1176/appi.ajp.2008.08081201 [PubMed: 19411368]
- Pizzagalli DA, Iosifescu D, Hallett LA, Ratner KG, Fava M. Reduced hedonic capacity in major depressive disorder: Evidence from a probabilistic reward task. *Journal of Psychiatric Research*. 2008; 43(1):76–87.10.1016/j.jpsychires.2008.03.001 [PubMed: 18433774]
- Radloff LS. The CES-D scale a self-report depression scale for research in the general population. *Applied psychological measurement*. 1977; 1(3):385–401.
- Ridderinkhof KR, van den Wildenberg W. Neurocognitive mechanisms of cognitive control: the role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*. 2004; 56(2):129–140.10.1016/j.bandc.2004.09.016 [PubMed: 15518930]
- Sears CR, Thomas CL, LeHuquet JM, Johnson JCS. Attentional biases in dysphoria: An eye-tracking study of the allocation and disengagement of attention. *Cognition & Emotion*. 2010; 24(8):1349–1368.
- Steyvers M, Lee MD, Wagenmakers EJ. A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*. 2009; 53(3):168–179.10.1016/j.jmp.2008.11.002
- Sutton, RS.; Barto, AG. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press; 1998.
- Teasdale JD. Cognitive Vulnerability to Persistent Depression. *Cognition & Emotion*. 1988; 2(3):247–274.

- Wadlinger HA, Isaacowitz DM. Looking happy: the experimental manipulation of a positive visual attention bias. *Emotion*. 2008; 8(1):121–126.10.1037/1528-3542.8.1.121 [PubMed: 18266522]
- Wagenmakers EJ, Farrell S. AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*. 2004; 11(1):192–196.10.3758/BF03206482 [PubMed: 15117008]
- Weissman MM, Sholomskas D. Assessing depressive symptoms in five psychiatric populations: a validation study. *American Journal of Epidemiology*. 1977; 106(3):203–214. [PubMed: 900119]
- Wells KB, Trust PM. The functioning and well-being of depressed patients: results from the Medical Outcomes Study. 1989; 35(5):425–433.
- Wells TT, Beevers CG. Biased attention and dysphoria: Manipulating selective attention reduces subsequent depressive symptoms. *Cognition & Emotion*. 2010; 24(4):719–728.
- Worthy DA, Maddox WT. Age-Based Differences in Strategy Use in Choice Tasks. *Frontiers in neuroscience*. 2011; 5:1–10.10.3389/fnins.2011.00145 [PubMed: 21390287]
- Yechiam E, Busemeyer JR. Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*. 2005; 12(3):387–402.10.3758/BF03193783 [PubMed: 16235624]

1a)

1b)

Figure 1.

a) Sample screen shot from the training task. In the placebo training condition the positive and neutral words each preceded the dot-probe with equal probability. In the active training condition the positive word preceded the dot-probe on 85% of trials. *b)* Sample screen shot from the decision-making task. Participants were told that they were testing two oxygen extraction systems. The oxygen extracted on each trial was shown in the “Current” tank then transferred to the “Cumulative” tank before the next trial began.

Figure 2.

Decision-making task reward structure. Depending on their selection, participants received the reward corresponding to Option A or Option B. Option A gave a mean reward of 55 points, while option B gave a mean reward of 65 points.

Figure 3.

a) Proportion of optimal selections for the decision-making task. Error bars represent standard error of the mean. *b)* Proportion of optimal selections for each group averaged across 1000 simulations of the extended RL model. Error bars represent standard error of the mean.

Figure 4.

a) Average number of times that participants in each group switched between reward options across 150 trials. Error bars represent standard error of the mean. b) Average number switches between reward options based on 1000 simulations for each group using the extended RL model. Error bars represent standard error of the mean.

Figure 5.

Akaike weights compare goodness of fit for the baseline model, Softmax RL model, and extended RL model. Higher Akaike weights indicate better fit. Error bars represent standard error of the mean.

Figure 6.

Learning rate parameter values for positive and negative reward prediction errors. Error bars represent standard error of the mean.