Social anxiety and dynamic social reinforcement learning in a volatile environment

Miranda L. Beltzer

Charlottesville, Virginia

A predissertation research project presented to the graduate faculty of the University of Virginia in candidacy for the degree of Master of Arts

Department of Psychology

University of Virginia

September 2017

Bethany A. Teachman, Ph.D.

Steven M. Boker, Ph.D.

Abstract

Human behavior is driven by seeking rewards and avoiding punishments, so difficulties learning about rewards and punishments can lead to maladaptive behavior. In fact, aberrant monetary reinforcement learning has been implicated in depression, schizophrenia, and other disorders, and researchers are beginning to find evidence for altered social reinforcement learning in social anxiety disorder. But learning is not a static process. Because the social environment is uncertain and unstable, the present study examines how social anxiety affects probabilistic social reward and punishment learning and dynamic updating of learned probabilities in a volatile environment. N=222 online participants completed a novel version of Cyberball, a computerized ball-catching game, and questionnaires. Mixed-effects regression analysis was used to analyze throw patterns as a function of social anxiety symptoms. To assess the relative importance ascribed to new information in response to volatility, dynamic learning rates were calculated by applying Q-learning algorithms to overlapping sets of throws. The new Cyberball task appears to be a valid measure of social reward learning; throws to the computerized avatars varied according to their probabilities of social reward and punishment, and performance improved over time as participants became more certain of the probability of social reward. Moreover, higher social anxiety predicted fewer throws to the avatar that had been the punisher in the previous block, suggesting that social anxiety may be characterized by difficulty updating learned social probabilities such that socially anxious individuals miss the chance to learn that a once-punishing situation no longer poses a threat.

Key words: Social anxiety, reinforcement learning, volatility, Cyberball, reward, punishment

Social interaction is both uncertain and volatile; few people give exclusively positive or negative social feedback, and the probability of a person reacting positively or negatively towards you changes over time. Even your closest friend will sometimes be angry at you, and it is unlikely that you are still close with all of your childhood friends. Navigating this uncertain, volatile social environment requires learning the probabilities of social reward and punishment, and then updating these mental representations when those probabilities change. Despite research suggesting that socially anxious individuals process social information differently than non-anxious individuals, including information about social rewards and punishments, little research to date has addressed how socially anxious individuals update probabilities about *volatile* social environments. The present study aims to address this question, here defining volatility as probabilities of reward and punishment that change over time.

Reinforcement Learning and Social Anxiety

On a basic level, human behavior is driven by seeking rewards and avoiding punishments. *Reinforcement learning* is the fundamental process by which an agent learns to predict and optimize their behavior in an environment where taking actions leads to rewards and punishments (Dayan & Niv, 2008). The agent chooses actions based on what they anticipate will optimize their outcomes over time, and optimizing performance requires adjusting their predictions of what the outcomes of their actions will be as they notice errors in their predictions. Beyond its relevance to psychology, engineers and computer scientists have used reinforcement learning to solve other types of optimization problems, and in doing so, have developed algorithms describing the process. Interestingly, these reinforcement learning algorithms appear not just to describe affective decision-making, but also to describe the neural mechanisms underlying these decisions. For instance, the phasic activity of dopaminergic neurons in reward-related regions of the brain mirrors the prediction error term in temporal difference models (Dayan & Niv, 2008). These neural findings provide support and validation for these models as descriptive of human decision-making processes.

Problems with reinforcement learning, which may include slow learning of rewards and punishments and/or decision-making that misaligns with learned probabilities of rewards and punishments, can have enormous downstream effects in terms of maladaptive behavior patterns. Aberrant reinforcement learning has been implicated in many psychological disorders, including major depression, schizophrenia, social anxiety disorder, and addiction, and at many levels of analysis, including neural function, structure, and connectivity, and behavior (for a review, see Dayan, 2009). Studies of major depressive disorder suggest that anhedonia, or loss of pleasure, is related to a reduced ability to modulate behavior as a function of rewards, and this may be due to blunted phasic dopaminergic signaling and reduced reward anticipation (Whitton et al., 2016). Given the high comorbidity of depression and social anxiety disorder (the 12-month correlation is 0.43; Kessler, Chiu, Demler, & Walters, 2005), similar processes may underlie both, including blunted reward sensitivity.

Anxiety disorders are typically associated with hypervigilance for threat, but there is reason to also consider aberrant responses to reward, particularly in the case of social anxiety disorder. The integrated hierarchical model of anxiety and depression (Brown, Chorpita, & Barlow, 1998), which examines how higher order negative and positive affect load onto mood and anxiety disorders, suggests that, of the anxiety disorders, social anxiety disorder has particular ties to reduced positive affect. Although high negative affect is implicated in many mood, anxiety, and obsessive-compulsive disorders, low positive affect is mostly only associated with mood disorders and social anxiety disorder, and at similar magnitudes (Brown et al., 1998).

Further, a meta-analysis found that social anxiety is negatively related to positive affect beyond what can be explained by co-occurring depression (Kashdan, 2007). Ecological momentary assessment studies have also found that highly socially anxious people experience less intense positive emotions both when with other people and when alone (Kashdan, Weeks, & Savostyanova, 2011). These studies suggest that there may be processes unique to social anxiety that decrease positive affect, which may suggest blunted reward sensitivity, and these processes may also distinguish social anxiety disorder from other anxiety disorders.

In the case of social anxiety disorder, we need to look beyond the characteristic fear of negative evaluation to investigate how individuals respond to social reward. Individuals with social anxiety disorder tend to respond negatively not just to negative social evaluation, but also to *positive* social evaluation (Weeks, Heimberg, Rodebaugh, & Norton, 2008; Weeks & Howell, 2014), suggesting that social events that are rewarding for most people may be processed differently by socially anxious individuals.

Social vs. non-social reinforcement learning. Most studies of reinforcement learning, particularly those investigating its ties to depression, typically involve monetary or other non-social rewards, which leads to the question: How does learning from social rewards and punishments align with non-social reinforcement learning (i.e., can we generalize from the prior studies examining non-social rewards to understand social reward learning)? Neurally, there are two major possibilities: social reward learning may employ a wholly separate neural system, or it may employ the same system as non-social reward learning but connect with structures involved specifically in social cognition (Ruff & Fehr, 2014). Many of the same structures underlie expected value calculations of both social and non-social (e.g., monetary) rewards, though studies have also found activation in regions specific to social and non-social rewards (Ruff &

Fehr, 2014). This suggests that at least some processes are common to both social and non-social reinforcement learning, while there may be some elements that diverge.

Behaviorally, some studies have found specific deficits in social reward learning in highly socially anxious individuals. Compared to men high in empathy, men low in empathy, who are more likely to be more socially anxious, showed reduced activation in the right nucleus accumbens (part of the brain's reward system) during social reward anticipation, but increased activation during monetary reward anticipation (Gossen et al., 2014). Additionally, individuals with social anxiety disorder showed decreased bilateral nucleus accumbens activation during social reward anticipation relative to healthy control participants, but did not show a significantly different response during monetary reward anticipation (Richey et al., 2014). These studies suggest that non-social reward learning may be intact for highly socially anxious individuals, while social reward anticipation is blunted.

Sensitivity to reward and/or punishment. Notably, aberrant social reinforcement learning in social anxiety may be driven by *hypo*sensitivity to social reward, *hyper*sensitivity to social punishment, or some combination of both. As evidence of hypersensitivity to social punishment, highly socially anxious individuals have been found to avoid stimuli with high probabilities of social punishment (scowling faces) more than less socially anxious individuals (Abraham & Hermann, 2015). Individuals with social anxiety disorder also show facilitated fear conditioning to critical faces, evidenced by potentiation of the startle-blink reflex, versus conditioning to neutral or happy faces, an effect not seen in individuals without social anxiety disorder (Lissek et al., 2008). Beyond facilitated learning of this fear response for socially punishing stimuli, it may also be more difficult to extinguish this response. Compared to people low in rejection sensitivity, people high in rejection sensitivity, which is related to social anxiety, display a greater resistance to extinction of a learned fear response to angry faces, but not to neutral faces or non-social stimuli (Olsson, Carmona, Downey, Bolger, & Ochsner, 2013). These studies all suggest that social anxiety disorder may be characterized by an exaggerated response to social punishment.

Consistent with the idea that both hyposensitivity to social reward and hypersensitivity to social punishment are present in social anxiety disorder, one study found that during social evaluative threat, higher trait punishment sensitivity predicted better punishment learning and poorer reward learning, but social stress had an opposite pattern of effects for people lower in trait punishment sensitivity (Cavanagh, Frank, & Allen, 2011). Although this study did not select for social anxiety and used a non-social reinforcement learning task, more socially anxious individuals are often higher in behavioral inhibition and punishment sensitivity and particularly reactive to social evaluative threat, so the results may also apply to socially anxious individuals (Gray & McNaughton, 2003; Kimbrel, Mitchell, & Nelson-Gray, 2010; Kimbrel, Nelson-Gray, & Mitchell, 2012; Mogg & Bradley, 2002).

Adapting to Volatility

Beyond learning static probabilities of reward and punishment, our volatile environment requires individuals to update those mental representations when circumstances change. Consider a new coworker who initially seems standoffish and rude, but after a few weeks, starts asking you about your weekend and seems more interested in getting to know you. You might have been correct to associate this person with a low probability of social reward and a high probability of social punishment at first, but if you are too slow to update this mental representation of them when they become friendlier, you may miss out on an opportunity to build a collegial working relationship or even a friendship. The social environment is volatile; as people and circumstances change over time, so too do relationships. Tracking these changes dynamically and updating how you think of other people is essential for optimal social behavior (Ronay & von Hippel, 2014). To do this, people need to assess how volatile the environment is (i.e., how stable or unstable probabilities of reward and punishment are) and adjust their decision-making accordingly. Computationally, this process is reflected in the learning rate, or how heavily recent information is weighted relative to previously learned information in predicting the outcomes of future actions (Behrens, Woolrich, Walton, & Rushworth, 2007). To continue the previous example, if you think the social environment is more volatile, you would pay more attention to how your coworker treated you this morning than last month when considering whether he would likely accept or reject an invitation to join you for lunch today.

How might anxiety predict learning in volatile environments? To optimize decision-making, a person's learning rate should increase in more volatile environments because information learned further in the past is less predictive of current probabilities of reward and punishment. Research suggests, however, that individuals high in trait anxiety adjust their learning rate less in response to the volatility of an aversive environment than less anxious individuals do (Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015). One study found no significant difference in learning rate between participants high and low in trait anxiety in either stable (i.e., probability of punishment does not change) or volatile (i.e., probability of punishment changes several times) aversive environments, but did find that highly anxious participants showed a smaller increase in the learning rate between the stable and volatile environments, reflecting an insensitivity of the learning rate to volatility (Browning et al., 2015). Moreover, another study did find that more (vs. less) punishment-sensitive individuals showed a lower learning rate for reward when not under social evaluative threat (suggesting lower state anxiety; Cavanagh et al., 2011). Taken together, these studies suggest that anxiety may be associated with some differences in learning rate, but the relationship may be dependent on whether the individual is updating probabilities of reward or punishment, the volatility of the environment, and the individual's state anxiety. Given the paucity of research on the effects of anxiety on learning rate in response to environmental volatility (and there are no studies specific to social anxiety to our knowledge), further investigation is warranted to understand how these learning processes affect behavior in social situations, which are inherently volatile. Reinforcement learning has yet to incorporate measures of learning rate changing over time, and such a measure would be particularly valuable for understanding not only how individuals learn probabilities of reward and punishment, but also how quickly they update this information in response to a changing environment.

Overview and Hypotheses

To better understand how aberrant social reinforcement learning characterizes social anxiety disorder, the present study investigates how individual differences in social anxiety symptoms predict how people learn and update probabilities of social reward and punishment in a volatile social environment. The environment was created through a novel modification of Cyberball (Williams, Cheung, & Choi, 2000), a virtual game of catch that has been used in over 200 studies and induces strong feelings of ostracism by making computerized avatars exclude the participant (Hartgerink, van Beest, Wicherts, & Williams, 2015). Our modified version varies the probability of each avatar throwing the ball to the participant, such that one avatar is highly rewarding (includes the participant), one is neutral (throws the ball equally to all players), and one is highly punishing (tends to exclude the participant), and these probabilities change at several timepoints. For some analyses, we used the participants' throwing and catching behavior in a simple, raw form (i.e., better performance is indicated by more throws to the rewarder than the neutral and punishing avatar, and by receiving the ball more often), and for some analyses, we applied a reinforcement learning algorithm to model the learning rate in response to environmental volatility (i.e., a higher learning rate indicates faster updating of mental representation by more heavily weighting new information about the probabilities of reward and punishment for each avatar).

Given past research on hypersensitivity to social punishment and hyposensitivity to social reward, we hypothesized that participants higher (relative to lower) in social anxiety would be less likely to throw to the punishing avatar, but also less likely to throw to the rewarding avatar, resulting in relatively more throws to the neutral avatar. Because avoiding the punisher would improve performance on the task (receiving the ball more), but not choosing the rewarder would impair performance, it was unclear whether there would be a main effect of social anxiety on overall task performance based on number of catches. We also hypothesized that differences as a function of social anxiety might emerge in the learning rate analyses; that participants higher (vs. lower) in social anxiety symptoms would be slower to adjust their learning rate when the punishing avatar became more rewarding, and faster to adjust when the rewarder became more punishing. These hypotheses were based on the emerging literature about anxiety's effects on learning rate and on more established findings that socially anxious individuals show blunted social reward anticipation, but facilitated learning and impaired extinction of fear responses to socially punishing stimuli.

Method

Participants

N=292 participants aged 18 and above were recruited through Amazon's Mechanical Turk to complete an online study. Only high reputation (90% approval rate or higher) workers were recruited to improve data quality. Because of potential cultural differences in expression of social anxiety (Heinrichs et al., 2006; Hofmann, Asnaani, & Hinton, 2010), 51 participants from outside of Pacific, Mountain, Central, and Eastern Time Zones were excluded prior to conducting analyses. To remove participants who were not paying attention to the dynamic social reinforcement task or who were not adhering to instructions, participants who almost exclusively targeted only one avatar throughout the game (despite changing roles) were removed. This was done by summing each participant's throws to their two least-targeted avatars, performing a mixture distribution analysis on this sum (using the Mclust package in R), and removing the cluster of participants with low counts on this sum. This analysis yielded three clusters, and the 19 participants in the cluster with fewest throws to the two least-targeted avatars were removed from subsequent analyses, resulting in a final *N*=222 participants (M_{age} = 34.89 years, SD_{age} = 10.95; 107 female, 114 male, 1 prefer not to answer).[1]

Measures

Volatile Social Learning Task (VSLT). Social reinforcement learning was measured with a modified version of a popular psychological task, Cyberball (Williams et al., 2000), versions of which have been used in hundreds of research studies (Hartgerink et al., 2015). Previous research has found that this simple online game of catch can elicit surprisingly strong emotional and behavioral responses to social inclusion and exclusion (Eisenberger, Lieberman, & Williams,

---

[1] To confirm that the participants in the removed cluster were indeed outliers, participants whose number of throws to the rewarder were more than 1.5 standard deviations below the mean in at least two of the four blocks of the task were identified. There was a high degree of overlap between participants in the removed cluster and outliers identified in this alternative fashion.

2003; Hartgerink et al., 2015). In our novel modified version, the stated goal is to maximize the number of times you receive the ball. The game is played with three computerized players who differ in the probability of reward (throwing the ball to the participant) and punishment (excluding the participant). Specifically, one avatar throws the ball to the participant with 0.7 probability (the "rewarder"), one with 0.33 probability (the "neutral player"), and one with 0.1 probability (the "punisher"). These roles change at three time points in the game (after each block of 100 throws, with a total of 4 blocks, resulting in 400 throws), such that all three avatar roles are always represented, but which avatar is in which role switches. Order of these role changes and the starting locations of each role were counterbalanced (see Figure 1).

Similar modifications of Cyberball have been employed to study social cognition in social anxiety (Fang, Hoge, Heinrichs, & Hofmann, 2014), autism spectrum disorder (Andari et al., 2010), and after traumatic brain injury (Kelly, McDonald, & Kellett, 2014). The VSLT is novel because it incorporates a true neutral player (.33 versus .3 probability) and includes more role switches and more trials than other versions, providing more data on different types of role transitions than other versions, and the VSLT includes solely social rewards (other versions also include monetary incentives). For more information about the development of the VSLT and its instructions, see Appendix. VSLT files are available for use in other studies: https://osf.io/9g56x/

From this task, we examined the throwing behavior of each participant. The main variables of interest were each participant's number of throws to each avatar during each block and during 25-throw segments of each block (to capture learning processes over the course of a block), and the number of times each participant received the ball (task performance). We also used a Q-learning algorithm (Sutton & Barto, 1998) to computationally model the learning rate, or the weight given to new information relative to what a participant had learned from all previous trials. Each participant's trial-by-trial behavior was fit by this model, in which the expected value (Q) of any state-action pair (s,a, where state is the location of the ball, and action is where the ball is thrown) computed after each time the ball was thrown by the participant or any avatar. At each time point, the Q-value of the state-action pair was updated by adding their prediction error multiplied by a learning rate (α) to their previous Q-value, following the equation:

$$Q(s,a) = Q(s,a) + \alpha*(R + \gamma*\max_{a'} Q(s',a') - Q(s,a))$$

In this algorithm, the prediction error is the difference between the observed reward (R = 1 when the participant receives the ball, and R = 0 in all other states) plus the expected value of their next state (temporally discounted) and their previous Q-value. Here, the discount rate (γ) was set to one, indicating no temporal discounting. Higher learning rates indicate more heavily weighting prediction errors to update the expected value of a state-action pair, and lower learning rates indicate relying more on the previous expected value with less influence of recent prediction errors.

The best-fitting parameters were found by entering Q-values into a softmax logistic algorithm to determine the probability of each action at each time, using the following equation:

$$P_A(t) = e^{(QA(t)/\beta)} / (e^{(QA(t)/\beta)} + e^{(QB(t)/\beta)} + e^{(QC(t)/\beta)})$$

where throwing to each avatar was modeled as a separate action (A, B, or C) and the inverse gain parameter (β) was fixed to 1 so that Q-values directly corresponded to the probabilities of throwing to each avatar. From these probabilities, the estimated log likelihood of a participant taking the actions they did was computed. A brute-force search was then used to maximize this log likelihood, searching over the space from α=0.01 to 1 with a step size of 0.01. This procedure for calculating learning rate was performed over the 400 trials of the VSLT for sliding windows of 25 trials, with a 5-trial increment between overlapping windows. Learning rate was calculated in MATLAB.

Questionnaires[2]

The Social Interaction Anxiety Scale (SIAS; Mattick & Clarke, 1998) is a widely-used 20-item self-report scale measuring social anxiety in dyads and groups (e.g., “I find myself worrying that I won’t know what to say in social situations.”). Participants rate how much they endorse each item on a 5-point Likert-type scale ranging from “not at all” to “extremely.” Three of these items are reverse scored, and Rodebaugh, Woods, and Heimberg (2007) have demonstrated that these items do not load onto the same factor as the straightforwardly worded items, and in fact seem to reflect extraversion more than neuroticism. Further, removing these reverse scored items generally improves the psychometric properties of the scale. Following Rodebaugh et al.’s recommendations, only ratings on the straightforwardly worded items were included.

Procedures

After informed consent, participants were directed to Qualtrics to complete a brief demographic questionnaire, the VSLT, questionnaires, and debriefing, in that order. For the VSLT, participants were randomly assigned to one of 36 versions of the task (six different orders of role switches, crossed with six different starting locations/avatars for each role; see Figure 1). Participants were compensated $2.50 for their participation in the study, which took less than an hour to complete.

---

[2] Additional questionnaires were administered but not discussed here because they were not central to these hypotheses. These include the Rejection Sensitivity Questionnaire (Downey & Feldman, 1996), Behavioral Inhibition Scales/Behavioral Activation Scales (Carver & White, 1994), Positive and Negative Affect Scale – State (Watson, Clark, & Tellegen, 1988), and Depression, Anxiety and Stress Scale (Lovibond & Lovibond, 1995).

Statistical Analyses

Mixed-effects modeling was used to model the number of throws from the participant to each avatar, and the number of catches the participant made. Mixed-effects modeling was chosen because our outcome measures were taken at multiple time points for each participant, and mixed-effects modeling allows for modeling individual differences with a separate random intercept for each participant. Across all models, random effects included random intercepts for each participant and each location (sequences of role changes were counterbalanced across locations of each avatar; i.e., left, top, and right). For all models, block number was modeled as an ordered categorical variable with polynomial contrasts, and orthogonal contrasts were used for all other categorical variables.

First, we aimed to validate the task by determining whether participants were able to discriminate the avatars by their roles and throw the ball accordingly: most to the rewarder, least to the punisher, with the neutral avatar in between. To do so, we modeled the number of throws from the participant to each avatar in each block from fixed effects of block number and each avatar's role, the interaction of block number and avatar role, and the aforementioned random effects. The interaction between block number and avatar role could not be included due to rank deficiency.

Second, we sought to determine whether social anxiety symptoms affected task performance, defined as succeeding at the task of receiving the ball as much as possible. For this analysis, we modeled the number of times the participant received the ball in each 25-trial segment of the game from fixed effects of the participant's social anxiety score (sum on the SIAS), the block number containing that segment, and the position of that segment within its block; and the aforementioned random effects. All two- and three-way interaction terms were included for social anxiety score, block number, and position of segment within block.

Third, we tested how social anxiety, avatar role, and previous avatar role affected participants' frequency of throws to each avatar over the course of each block for the whole game. For this analysis, we modeled the number of throws from the participant to an avatar during each 25-trial segment of the game from fixed effects of the participant's social anxiety score (calculated as above), the avatar's role during that segment, the role of that avatar in the previous block, the position of that segment within its block, and the block number containing that segment; and a random intercept for location but not for each participant, as the model with a random intercept for each participant was not structurally sound (rank deficient). All two-, three-, and four-way interactions were included among fixed effects, excluding block number, due to rank deficiency. This analysis excluded the first block of 100 trials, as avatars did not have previous roles during this block.

Finally, we used a novel approach to investigate how individuals updated their mental representations of the social environment after a role transition as a function of social anxiety symptoms and the type of role shifts occurring. As noted above, learning rates were estimated for each participant for sliding windows of 25-trials, with a five-trial increment between overlapping windows. As an individual perceives more volatility in the environment, they are expected to more heavily weight new information in their mental representations of the environment because old information is becoming less relevant to predicting current probabilities of reward and punishment. Accordingly, we would expect learning rate to increase as the participant perceives greater volatility and updates their mental representations of the other avatars more quickly. Once their learning rate reaches a peak and starts decreasing, the participant is presumably updating their mental representations less quickly. Thus, the first peak in their learning rate after a role transition represents the point at which the participant perceives the greatest environmental volatility and is relying most heavily on new information to update their mental representations of other avatars. For this reason, we calculated when this first peak in learning rate occurred after each role transition for each participant by finding the point at which the slope of learning rate became negative after the role transition. Because learning rate was estimated for overlapping windows, we defined role transitions as occurring between two adjacent windows: when the majority of trials comprising the window (15 out of 25) occurred before the transition, and when the majority (again, 15 out of 25) occurred after the transition. We refer to this as the lag after a transition, with shorter lags indicating quicker updating after a transition. Given the novel analytic approach and the conceptual novelty of examining dynamic changes in learning rate, this analysis was more exploratory.

For this analysis, we modeled each participant's lag after each role transition from fixed effects of their social anxiety symptom score, the previous role of the rewarder and the punisher, and interactions between social anxiety and previous role of rewarder, and social anxiety and previous role of punisher; and the aforementioned random effects (i.e., random intercepts for participant and location). Because the previous role of the neutral avatar is dependent on the previous roles of the rewarder and punisher, we opted to exclude the previous role of the neutral avatar from this model to avoid rank deficiency. We also excluded block number to avoid rank deficiency, as certain role switches only occurred between certain blocks. This analysis excluded the first block of 100 trials, as avatars did not have previous roles during this block.

All analyses were also performed with social anxiety split categorically into low and high groups. Participants with SIAS scores less than or equal to three quarters of a standard deviation

(10 or under) below the mean of a previous community sample (*M*=18.8, *SD*=11.8; Mattick & Clarke, 1998) were included in the low social anxiety group, and those scoring greater than or equal to three quarters of a standard deviation (28 or greater) above the mean were included in the high social anxiety group. Results were largely similar to those described below and are not included to reduce redundancy but are available from the first author.

Results

Validating the Task

To validate the task, a model was run predicting the number of throws from the participant to each avatar in each block from each avatar's role and the block number; see Table 1. As expected, avatar role significantly predicted how frequently participants threw to that avatar: participants threw to the rewarder significantly more frequently than to the neutral avatar ($\beta = 7.97$, $p < 0.001$), and participants threw to the punisher significantly less frequently than to the neutral avatar ($\beta = -5.08$, $p < 0.001$); see Figure 2.

There were also a number of significant interactions with block number. As block number increased linearly, there were relatively fewer throws to the rewarder than to the neutral avatar ($\beta = -1.79$, $p < 0.001$). As block number increased quadratically, there were relatively fewer throws to the punisher versus the neutral avatar ($\beta = -1.05$, $p = 0.01$). As block number increased cubically, there were relatively more throws to the punisher ($\beta = 0.89$, $p = 0.02$) and fewer to the rewarder ($\beta = -2.38$, $p < 0.001$) versus the neutral avatar. Visual inspection of the mean number of throws to each avatar in each block (Figure 3) revealed that 95% confidence intervals for throws to the neutral avatar and the punisher overlapped in blocks 2 and 3. Throws to the rewarder decreased between blocks 1 and 2, increased between blocks 2 and 3, and decreased between blocks 3 and 4. Throws to the punisher increased between blocks 1 and 2, and then slightly decreased from blocks 2 through 4. Throws to the neutral avatar slightly increased between blocks 1 and 2, slightly decreased between blocks 2 and 3, and increased between blocks 3 and 4. In summary, participants threw to avatars in accordance with their roles (most to the rewarder and least to the punisher), though there were slight differences in this effect in different blocks.

Task Performance

To assess the effect of social anxiety on task performance, a model was run predicting the number of times the participant received the ball (catches) in each 25-trial segment of the game from the participant's SIAS score, the block number containing that segment, and the position of that segment within its block; see Table 2. SIAS did not significantly predict the number of catches ($\beta < 0.001$, $p = 0.99$). There was a significant main effect of position of segment within block ($\beta = 0.13$, $p < 0.001$), such that as a block progressed, participants tended to receive the ball more. There was also a significant cubic effect of block number ($\beta = -0.35$, $p = 0.03$). Visual inspection revealed that catches were highest in block 1, lowest in block 2, increased in block 3, but not to the level of block 1, and then slightly decreased again in block 4. No two- or three-way interactions of SIAS, position of segment within block, and block number were significant. Task performance did not vary as a function of social anxiety symptoms, but it did improve over the course of each block, with some differences in task performance across blocks.

Throws to Each Avatar as a Function of Social Anxiety

To assess whether social anxiety affected the frequency of throws to each avatar, a model was run predicting number of throws from the participant to each avatar during each 25-trial segment of the game from the participant's SIAS score, the avatar's role during that segment, the role of that avatar in the previous block, the position of that segment within its block, and the block number containing that segment; see Table 3. The avatar's current role significantly predicted how frequently the participant threw to that avatar: participants threw to the rewarder significantly more frequently than to the neutral avatar ($\beta = 1.27$, $p < 0.001$), and participants threw to the punisher significantly less frequently than to the neutral avatar ($\beta = -0.81$, $p < 0.001$). The avatar's previous role was also a significant predictor, but only for one contrast. Participants threw significantly more to avatars who, in the previous block, had been the rewarder compared to neutral ($\beta = 0.73$, $p < 0.001$), but there was no significant difference based on whether the avatar had just been the punisher or neutral ($\beta = -0.28$, $p = 0.14$). We found no main effects of SIAS ($\beta < 0.001$, $p = 0.38$), segment position within its block ($\beta = 0.01$, $p = 0.16$), or block number.

Several interactions were also significant predictors of a participant's throws. The interaction between an avatar's current role and the position of the segment within the block was significant: as the block progressed, participants threw more to the rewarder compared to the neutral avatar ($\beta = 0.32$, $p < 0.001$) and less to the punisher compared to the neutral avatar ($\beta = -0.18$, $p = 0.01$). There was also a significant interaction between a participant's SIAS score and an avatar's previous role, as anticipated: participants with higher social anxiety threw less to the avatar who had just been the punisher versus neutral ($\beta = -0.01$, $p = 0.02$), in line with the expectation of greater sensitivity to social punishment (see Figure 4). The interaction between SIAS and previous avatar role was not significant for the contrast of previous rewarder and previous neutral avatar ($\beta < 0.001$, $p = 0.28$). The interaction of an avatar's current and previous roles was not significant, nor was the interaction of SIAS and segment position within a block. None of the three- or four-way interactions were significant. In summary, the differentiation of throws by avatar role became more extreme over the course of a block, and higher social anxiety symptoms predicted fewer throws to the avatar who had been the punisher in the previous block, suggesting less updating of expectations of social punishment in a more neutral or rewarding direction.

Learning Rate

To assess whether social anxiety affected the speed at which participants updated their mental representations of the avatars after role transitions, a model was run predicting each participant's lag to peak learning rate after each role transition from their SIAS score, the previous roles of the rewarder and the punisher, and the interactions between SIAS and previous roles of the rewarder and punisher; see Table 4. None of the main effects were statistically significant, nor was either interaction.

Discussion

This study used a novel, modified version of Cyberball to assess the effect of social anxiety symptoms on updating mental representations of volatile probabilities of social reward and punishment. The results partially supported hypotheses. Participants were able to distinguish the avatars by their probabilities of social reward and punishment, throwing more frequently to the rewarder than the neutral avatar and more frequently to the neutral avatar than the punisher, suggesting the task was a valid way to demonstrate reinforcement learning. Moreover, throws to the rewarder increased and throws to the punisher decreased relative to the neutral avatar over the course of a block, as participants had more time to learn these probabilities, providing further evidence that participants learned the probabilities of reward and punishment. Results also supported the hypothesis that social anxiety is characterized by hypersensitivity to social punishment, as participants higher in social anxiety symptoms were less likely to throw to the avatar who had been the punisher (versus neutral) in the prior block, suggesting that they were slower to update their expectations of reward after experiencing social punishment. We did not find support for the hypothesis of hyposensitivity to social reward in social anxiety, nor did we find significant differences in learning rate as a function of social anxiety or the type of role change occurring. Thus, while *past* probabilities of social punishment interacted with social anxiety to predict throw patterns, *current* probabilities did not. This suggests that social anxiety does not affect real-time learning of probabilities of social reward and punishment, but it does affect updating of learned probabilities.

Our results suggest that socially anxious individuals may have difficulty adapting after someone has rejected them (or they perceive that they have been rejected) and then later becomes more rewarding. This extends previous research on negative biases in social anxiety disorder. Our finding that participants higher in social anxiety were less likely to throw to a previously punishing avatar is in line with theorized and observed attentional and memory biases for threat in social anxiety disorder (Hirsch & Clark, 2004; Mogg & Bradley, 2002). Hirsch and Clark's (2004) influential information-processing account of social anxiety disorder posits that one factor maintaining the disorder is the tendency to overestimate the probability of future social threat, which our findings suggest might occur due to lack of updating these probabilities after a rejection experience. Attentional bias for (Mogg & Bradley, 2002) and difficulty disengaging from socially threatening stimuli (Buckner, Maner, & Schmidt, 2010) may contribute to this update bias, as preferential allocation of attentional resources may render socially punishing experiences more salient in memory. Similarly, the tendency for socially anxious (compared to non-anxious) individuals to show greater fear conditioning to social stimuli (Lissek et al., 2008; Pejic, Hermann, Vaitl, & Stark, 2013) and impaired extinction of conditioned responses to socially threatening (i.e., angry) faces (Olsson et al., 2013) may also contribute to the difficulties “letting go” of a prior punishment experience. It will be interesting in future work to determine the specific mechanisms that interfere with effective updating when a punishment has ended. In addition to testing the roles of attention and memory biases, using behavioral economic games to probe the role of trust in this updating difference may be useful.

The findings of hypersensitivity to social punishment may also be relevant to the growing literature on post-event processing in social anxiety, which generally finds evidence of negatively biased recall for social feedback over time (Cody & Teachman, 2010; Koban et al., 2017). The post-event processing literature typically refers to the persistence of negative *explicit* self-relevant thoughts, and our results suggest that negative *implicitly* learned rules (e.g., that person does not like me because they barely throw me the ball) may also persist over time. Probing participants’ conscious awareness of the changes in throw patterns would be interesting in future research to determine in what ways the rule learning is implicit versus explicit.

While the hypothesis tied to hypersensitivity to social punishment was supported, we did not find support for the expectations of hyposensitivity to reward or learning rate differences. An intriguing possibility that may partially explain our null findings is that social anxiety might affect how individuals experience the *magnitudes* of reward and punishment within the paradigm. Reward and punishment were modeled consistently across all participants (e.g., it was assumed that a probability of 33% reward was experientially equivalent to a probability of 33% punishment), but it may be that more socially anxious individuals experience receiving the ball as less rewarding and not receiving the ball as more punishing than do less socially anxious participants. This might reflect individual differences in values and goals; for socially anxious participants, the cost of being excluded might be much higher than the reward of receiving the ball, so receiving the ball might be better modeled as a lack of punishment (really, a relief) than receiving a reward. Future studies using inverse reinforcement learning techniques may investigate potential social anxiety-related differences in reward function. Also, qualitative interviews with anxious participants after they have completed the task would be informative.

Clinical Implications

Punishment in the modified Cyberball task (VSLT) is a relatively minimal rejection experience (a generic cartoon representing a person unknown to the participant virtually throws a cartoon ball to the participant less frequently than to other players), yet it was sufficient to elicit decreased updating in participants higher in social anxiety. Real social interactions are much richer, with information being communicated through many channels: facial expressions, body language, vocal tone, words, and timing of responses, to name a few. The current results are consistent with the possibility that perceiving rejection in any of these numerous streams of information might lead a socially anxious person to be loath to update their expectations of this person in the future. Unwilling to let go of some past perceived slight, they may not be able to think of this person as a potential friend or even as a neutral acquaintance and consequently avoid social interactions – potential opportunities to learn that this person is actually more rewarding than they thought.

As such, updating expectations may be a potentially fruitful therapeutic target. In a one-on-one therapeutic setting, this might involve exploring negative social expectations about peers or other people in the client's life and encouraging the client to draw on more recent versus historical interactions when anticipating how these people will treat the client in the future. This more recent focus may mitigate their bias against updating negative social expectations. Computerized cognitive trainings, such as cognitive bias modification, may also consider targeting this update bias by rewarding socially anxious users for more quickly trusting previously threatening others whose behavior has become more rewarding. Determining when to encourage this trust is not simple, however. It is an interesting challenge for both socially anxious and non-anxious individuals to determine when it is healthy to move beyond a prior actual or perceived slight and give someone another chance. The line between adaptive self-protection and maladaptive avoidance is not always obvious, and the literature on forgiveness may have some interesting lessons for how to overcome the updating bias observed in the present study.

Limitations and Future Directions

There are several limitations to this study that should be addressed in future research. First, the null results for learning rate may reflect limitations due to the design of the modified Cyberball task (VSLT) and the analytic approach. The VSLT was designed so that all roles would be represented at all times and all possible role shifts would occur, but this design could be improved. As can be seen in Figure 1, two sets of role shifts occurred between blocks one and two, and these same two sets occurred between blocks three and four, but never occurred between blocks two and three. Similarly, the three sets of role shifts that occurred between blocks two and three never occurred between blocks one and two or three and four. Further, the three sets of role shifts occurring between blocks two and three all held one avatar role constant and switched the roles of the other two avatars, which would arguably create less volatility between blocks two and three. It is possible that this design necessitated less learning and behavior change between these blocks, as can be seen in Figure 2, which demonstrates throws to the rewarder increasing and throws to the other avatars decreasing relatively linearly over the course of blocks two and three. This design might also be responsible for the block number effects found in some models. This was unintentional and could be changed in future versions of the task.

There are two other limitations in our calculation of learning rate that may have contributed to our null results. First, we selected a 25-trial window for calculating each learning rate, but 25 trials may be insufficient to estimate a valid learning rate, especially considering that all throws (those done by participants, but also by avatars) were included in this 25-trial window. Second, we calculated the lag to peak learning rate after a role switch for each individual participant, but the individual data was very noisy. Some studies (e.g., Reiter, Heinze, Schlagenhauf, & Deserno, 2016; Schiller, Levy, Niv, LeDoux, & Phelps, 2008) pool participants to analyze learning rate on a group, rather than individual, level, which may reduce the individual-level noise and reveal a signal. However, there are benefits to analyzing individual differences in reinforcement learning, as models with individual learning parameters may explain more variance in the data (Cohen, 2007).

Another limitation is that while we measured trait social anxiety with the SIAS, we did not collect any markers of state social anxiety, perceived exclusion, or other constructs during the task. Given alterations in learning rate as a function of state positive affect (Bakic, Jepma, De Raedt, & Pourtois, 2014) and other model-based decision-making parameters as a function of acute psychosocial stress (Radenbach et al., 2015), these state markers might differ between groups and be important in modeling reinforcement learning. Inclusion of psychophysiological measures throughout the VSLT and/or subjective units of distress ratings before, during, and after the VSLT would improve our ability to detect state effects of social anxiety on learning. Further, we recruited a non-clinical sample, so it is possible that using a sample seeking treatment for social anxiety disorder might find different results. However, our analysis of social anxiety symptoms as a continuous variable allows us to investigate a fuller range of functioning, from healthy to disordered, in line with the National Institute of Mental Health's Research Domain Criteria initiative (Insel et al., 2010).

Collecting data online is another possible limitation of our research, as the setting in which participants complete the study cannot be controlled and participants may become distracted during the study. In order to mitigate these risks, comprehension questions were added following VSLT instructions and only highly rated mTurk workers were recruited. Further, participants whose throwing patterns on the VSLT indicated that they were likely not paying attention or were not following instructions were identified and excluded from analyses. Benefits of online data collection for clinical research include fast, inexpensive data collection from large samples, and higher prevalence of social anxiety symptoms than the general population, which was particularly helpful for recruitment for this study (Shapiro, Chandler, & Mueller, 2013).

The VSLT is a novel paradigm that, with small adaptations, could be well-suited to investigating other questions related to reinforcement learning. Measuring concurrent physiological activity with electroencephalography (EEG), eye tracking, or other monitoring devices might provide insight into the biological processes underlying learning and updating probabilities of social reward and punishment. Numerous studies have demonstrated the sensitivity of EEG for measuring electrophysiological correlates of updating social value expectations (e.g., Gutz, Renneberg, Roepke, & Niedeggen, 2015) and reinforcement learning more generally (see Frank et al., 2015; Lee, D., Seo H., 2012). Eye tracking might also provide important information about neural processes, given spontaneous blink rate has been connected to striatal dopamine levels and reinforcement learning (e.g., Slagter, Georgopoulou, & Frank, 2015), pupil dilation can reflect expectancy violations (Browning et al., 2015), and gaze fixation patterns can be used to measure attentional engagement and disengagement (Buckner et al., 2010). Further, the VSLT could be easily altered to include more specific social cues, such as photos and names of the participant and their peers, or could be adapted to measure updating of volatile probabilities of non-social reward by changing the other avatars to generic shapes, which could probe the domain-specificity of effects.

Conclusion

In summary, a novel paradigm, the Volatile Social Learning Task, was used to assess how updating volatile probabilistic information about social reward and punishment varies as a function of trait social anxiety. Evidence of biased updating in social anxiety was found; more socially anxious participants were less likely to update their negative expectations of players who had previously been punishing. This update bias may contribute to avoidance behavior and maintenance of social anxiety disorder (though that cannot be established from this first study's correlational design). Computational methods examining individualized learning parameters may improve our understanding of learning processes contributing to psychopathology, and incorporating volatility into learning environments allows us to examine dynamic processes relevant to real-world social interactions.

References

Abraham, A., & Hermann, C. (2015). Biases in probabilistic category learning in relation to social anxiety. *Frontiers in Psychology*, *6*, 1218. http://doi.org/10.3389/fpsyg.2015.01218

Andari, E., Duhamel, J.-R., Zalla, T., Herbrecht, E., Leboyer, M., & Sirigua, A. (2010). Promoting social behavior with oxytocin in high-functioning autism spectrum disorders. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *107*(9), 4389–4394. http://doi.org/10.1073/pnas.0910249107

Bakic, J., Jepma, M., De Raedt, R., & Pourtois, G. (2014). Effects of positive mood on probabilistic learning: Behavioral and electrophysiological correlates. *Biological Psychology*, *103*, 223–232. http://doi.org/10.1016/j.biopsycho.2014.09.012

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221. http://doi.org/10.1038/nn1954

Brown, T. A., Chorpita, B. F., & Barlow, D. H. (1998). Structural relationships among dimensions of the DSM-IV anxiety and mood disorders and dimensions of negative affect, positive affect, and autonomic arousal. *Journal of Abnormal Psychology*, *107*(2), 179–192. http://doi.org/http://dx.doi.org/10.1037/0021-843X.107.2.179

Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*, *18*(4), 590–596. http://doi.org/10.1038/nn.3961

Buckner, J. D., Maner, J. K., & Schmidt, N. B. (2010). Difficulty disengaging attention from social threat in social anxiety. *Cognitive Therapy and Research*, *34*(1), 99–105. http://doi.org/10.1007/s10608-008-9205-y

Carver, C. S., & White, T. L. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS / BAS Scales. *Journal of Personality and Social Psychology*, *67*(2), 319–333. http://doi.org/http://dx.doi.org/10.1037/0022-3514.67.2.319

Cavanagh, J. F., Frank, M. J., & Allen, J. J. B. (2011). Social stress reactivity alters reward and punishment learning. *Social Cognitive and Affective Neuroscience*, *6*(3), 311–320. http://doi.org/10.1093/scan/nsq041

Cody, M. W., & Teachman, B. A. (2010). Post-event processing and memory bias for performance feedback in social anxiety. *Journal of Anxiety Disorders*, *24*(5), 468–479. http://doi.org/10.1016/j.janxdis.2010.03.003

Cohen, M. X. (2007). Individual differences and the neural representations of reward expectation and reward prediction error. *Social Cognitive and Affective Neuroscience*, *2*(1), 20–30. http://doi.org/10.1093/scan/nsl021

Dayan, P. (2009). Dopamine, reinforcement learning, and addiction. *Pharmacopsychiatry*, *42*(Suppl1), S56–S65. http://doi.org/10.1055/s-0028-1124107

Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, *18*(2), 185–196. http://doi.org/10.1016/j.conb.2008.08.003

Downey, G., & Feldman, S. I. (1996). Implications of rejection sensitivity for intimate relationships. *Journal of Personality and Social Psychology*, *70*(6), 1327–1343. http://doi.org/http://dx.doi.org/10.1037/0022-3514.70.6.1327

Eisenberger, N. I., Lieberman, M. D., & Williams, K. D. (2003). Does rejection hurt? An FMRI study of social exclusion. *Science*, *302*(5643), 290–292. http://doi.org/http://dx.doi.org/10.1126/science.1089134

Fang, A., Hoge, E. A., Heinrichs, M., & Hofmann, S. G. (2014). Attachment style moderates the effects of oxytocin on social behaviors and cognitions during social rejection: Applying a research domain criteria framework to social anxiety. *Clinical Psychological Science*, *2*(6), 740–747. http://doi.org/http://dx.doi.org/10.1177/2167702614527948

Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. (2015). fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *The Journal of Neuroscience*, *35*(2), 485–494. http://doi.org/10.1523/JNEUROSCI.2036-14.2015

Gossen, A., Groppe, S. E., Winkler, L., Kohls, G., Herrington, J., Schultz, R. T., … Spreckelmeyer, K. N. (2014). Neural evidence for an association between social proficiency and sensitivity to social reward. *Social Cognitive and Affective Neuroscience*, *9*(5), 661–670. http://doi.org/10.1093/scan/nst033

Gray, J. A., & McNaughton, N. (2003). *The neuropsychology of anxiety: An enquiry into the functions of the septo-hippocampal system*. Oxford University Press.

Gutz, L., Renneberg, B., Roepke, S., & Niedeggen, M. (2015). Neural Processing of Social Participation in Borderline Personality Disorder and Social Anxiety Disorder, *124*(2), 421–431. http://doi.org/10.1037/a0038614

Hartgerink, C. H. J., van Beest, I., Wicherts, J. M., & Williams, K. D. (2015). The ordinal effects of ostracism: a meta-analysis of 120 Cyberball studies. *PloS One*, *10*(5), e0127002. http://doi.org/10.1371/journal.pone.0127002

Heinrichs, N., Rapee, R. M., Alden, L. A., Bögels, S., Hofmann, S. G., Oh, K. J., & Sakano, Y. (2006). Cultural differences in perceived social norms and social anxiety. *Behaviour Research and Therapy*, *44*(8), 1187–1197. http://doi.org/10.1016/j.brat.2005.09.006

Hirsch, C. R., & Clark, D. M. (2004). Information-processing bias in social phobia. *Clinical Psychology Review*, *24*(7), 799–825. http://doi.org/10.1016/j.cpr.2004.07.005

Hofmann, S. G., Asnaani, A., & Hinton, D. E. (2010). Cultural aspects in social anxiety and social anxiety disorder. *Depression and Anxiety*, *27*(12), 1117–1127. http://doi.org/10.1002/da.20759

Insel, T., Cuthbert, B., Garvey, M., Heinssen, R., Pine, D. S., Quinn, K., … Wang, P. (2010). Research domain criteria (RDoC): Toward a new classification framework for research on mental disorders. *American Journal of Psychiatry*, *167*(7), 748–751. http://doi.org/10.1176/appi.ajp.2010.09091379

Kashdan, T. B. (2007). Social anxiety spectrum and diminished positive experiences: Theoretical synthesis and meta-analysis. *Clinical Psychology Review*, *27*(3), 348–365. http://doi.org/10.1016/j.cpr.2006.12.003

Kashdan, T. B., Weeks, J. W., & Savostyanova, A. A. (2011). Whether, how, and when social anxiety shapes positive experiences and events: A self-regulatory framework and treatment implications. *Clinical Psychology Review*, *31*(5), 786–799. http://doi.org/10.1016/j.cpr.2011.03.012

Kelly, M., McDonald, S., & Kellett, D. (2014). Development of a novel task for investigating decision making in a social context following traumatic brain injury. *Journal of Clinical and Experimental Neuropsychology*, *36*(9), 897–913. http://doi.org/10.1080/13803395.2014.955784

Kessler, R. C., Chiu, W. T., Demler, O., & Walters, E. E. (2005). Prevalence, Severity, and Comorbidity of Twelve-Month DSM-IV Disorders in the National Comorbidity Survey Replication (NCS-R). *Archives of General Psychiatry*, *62*(6), 617–627.

http://doi.org/10.1001/archpsyc.62.6.617

Kimbrel, N. A., Mitchell, J. T., & Nelson-Gray, R. O. (2010). An examination of the relationship between behavioral approach system (BAS) sensitivity and social interaction anxiety. *Journal of Anxiety Disorders*, *24*(3), 372–378. http://doi.org/10.1016/j.janxdis.2010.02.002

Kimbrel, N. A., Nelson-Gray, R. O., & Mitchell, J. T. (2012). BIS, BAS, and Bias: The role of personality and cognitive bias in social anxiety. *Personality and Individual Differences*, *52*(3), 395–400. http://doi.org/10.1016/j.paid.2011.10.041

Koban, L., Schneider, R., Ashar, Y. K., Andrews-Hanna, J. R., Landy, L., Moscovitch, D. A., … Arch, J. J. (2017). Social anxiety is characterized by biased learning about performance and the self. *Emotion*. http://doi.org/10.1037/emo0000296

Lee, D., Seo, H., & Jung, M. W. (2012). Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience*, *35*, 287–308. http://doi.org/10.1146/annurev-neuro-062111-150512

Lissek, S., Levenson, J., Biggs, A. L., Johnson, L. L., Ameli, R., Pine, D. S., & Grillon, C. (2008). Elevated fear conditioning to socially relevant unconditioned stimuli in social anxiety disorder. *The American Journal of Psychiatry*, *165*(1), 124–132. http://doi.org/10.1176/appi.ajp.2007.06091513

Lovibond, P. F., & Lovibond, S. H. (1995). The structure of negative emotional states: comparison of the Depression Anxiety Stress Scales (DASS) with the Beck Depression and Anxiety Inventories. *Behaviour Research and Therapy*, *33*(3), 335–343. http://doi.org/http://dx.doi.org/10.1016/0005-7967(94)00075-U

Mattick, R. P., & Clarke, J. C. (1998). Development and validation of measures of social phobia scrutiny fear and social interaction anxiety. *Behavior Research and Therapy*, *36*(4), 455–

470. http://doi.org/10.1016/S0005-7967(97)10031-6

Mogg, K., & Bradley, B. P. (2002). Selective orienting of attention to masked threat faces in social anxiety. *Behaviour Research and Therapy*, *40*(12), 1403–1414. http://doi.org/10.1016/S0005-7967(02)00017-7

Olsson, A., Carmona, S., Downey, G., Bolger, N., & Ochsner, K. N. (2013). Learning biases underlying individual differences in sensitivity to social rejection. *Emotion*, *13*(4), 616–621. http://doi.org/10.1037/a0033150

Pejic, T., Hermann, A., Vaitl, D., & Stark, R. (2013). Social anxiety modulates amygdala activation during social conditioning. *Social Cognitive and Affective Neuroscience*, *8*(3), 267–276. http://doi.org/10.1093/scan/nsr095

Radenbach, C., Reiter, A. M. F., Engert, V., Sjoerds, Z., Villringer, A., Heinze, H. J., … Schlagenhauf, F. (2015). The interaction of acute and chronic stress impairs model-based behavioral control. *Psychoneuroendocrinology*, *53*, 268–280. http://doi.org/10.1016/j.psyneuen.2014.12.017

Reiter, A. M., Heinze, H.-J., Schlagenhauf, F., & Deserno, L. (2017). Impaired flexible reward-based decision-making in Binge Eating Disorder: Evidence from computational modeling and functional neuroimaging. *Neuropsychopharmacology*, *42*(3), 628–637. http://doi.org/10.1038/npp.2016.95

Richey, J. A., Rittenberg, A., Hughes, L., Damiano, C. R., Sabatino, A., Miller, S., … Dichter, G. S. (2014). Common and distinct neural features of social and non-social reward processing in autism and social anxiety disorder. *Social Cognitive and Affective Neuroscience*, *9*(3), 367–377. http://doi.org/10.1093/scan/nss146

Rodebaugh, T. L., Woods, C. M., & Heimberg, R. G. (2007). The reverse of social anxiety is not always the opposite: The reverse-scored items of the social interaction anxiety scale do not belong. *Behavior Therapy*, *38*(2), 192–206. http://doi.org/http://dx.doi.org/10.1016/j.beth.2006.08.001

Ronay, R., & von Hippel, W. (2014). Sensitivity to changing contingencies predicts social success. *Social Psychological and Personality Science*, *6*(1), 23–30. http://doi.org/10.1177/1948550614542348

Ruff, C. C., & Fehr, E. (2014). The neurobiology of rewards and values in social decision making. *Nature Reviews Neuroscience*, *15*(8), 549–562. http://doi.org/10.1038/nrn3776

Schiller, D., Levy, I., Niv, Y., LeDoux, J. E., & Phelps, E. A. (2008). From fear to safety and back: Reversal of fear in the human brain. *The Journal of Neuroscience*, *28*(45), 11517–11525. http://doi.org/10.1523/JNEUROSCI.2265-08.2008

Shapiro, D. N., Chandler, J., & Mueller, P. A. (2013). Using mechanical turk to study clinical populations. *Clinical Psychological Science*, *1*(2), 213–220. http://doi.org/10.1177/2167702612469015

Slagter, H. A., Georgopoulou, K., & Frank, M. J. (2015). Spontaneous eye blink rate predicts learning from negative, but not positive, outcomes. *Neuropsychologia*, *71*, 126–132. http://doi.org/10.1016/j.neuropsychologia.2015.03.028

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning : An Introduction*. Cambridge, Mass: A Bradford Book.

Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of Personality and Social Psychology*, *54*(6), 1063–1070. http://doi.org/http://dx.doi.org/10.1037/0022-3514.54.6.1063

Weeks, J. W., Heimberg, R. G., Rodebaugh, T. L., & Norton, P. J. (2008). Exploring the relationship between fear of positive evaluation and social anxiety. *Journal of Anxiety Disorders*, *22*(3), 386–400. http://doi.org/10.1016/j.janxdis.2007.04.009

Weeks, J. W., & Howell, A. N. (2014). Fear of positive evaluation: The neglected fear domain in social anxiety. In J. W. Weeks (Ed.), *The Wiley Blackwell handbook of social anxiety disorder* (pp. 433–453). http://doi.org/10.1002/9781118653920.ch20

Whitton, A. E., Treadway, M. T., & Pizzagalli, D. A. (2016). Reward processing dysfunction in major depression, bipolar disorder and schizophrenia. *Current Opinion in Psychiatry*, *28*(1), 7–12. http://doi.org/10.1097/YCO.0000000000000122

Williams, K. D., Cheung, C. K. T., & Choi, W. (2000). Cyberostracism: effects of being ignored over the Internet. *Journal of Personality and Social Psychology*, *79*(5), 748–762. http://doi.org/http://dx.doi.org/10.1037/0022-3514.79.5.748

Table 1. *Model effects for validation analysis, predicting number of participant's throws to an avatar*.

Predictor	Fixed Effect	t	p
Intercept	10.30	76.01	<0.001*
Role (P vs N)	-5.08	-26.49	<0.001*
Role (R vs N)	7.97	41.60	<0.001*
Block number (linear)	-0.08	-0.31	0.75
Block number (quadratic)	0.03	0.10	0.92
Block number (cubic)	-0.18	-0.66	0.51
Role (P vs N) x block number (linear)	0.66	1.72	0.09
Role (R vs N) x block number (linear)	-1.79	-4.66	<0.001*
Role (P vs N) x block number (quadratic)	-1.05	-2.75	0.01*
Role (R vs N) x block number (quadratic)	0.58	1.53	0.13
Role (P vs N) x block number (cubic)	0.89	2.31	0.02*
Role (R vs N) x block number (cubic)	-2.38	-6.20	<0.001*
	Random Effect Variance		
Participant intercept	0.00		
Location intercept	2.62 e^{-15}		

Note: P = punisher, N = neutral, R = rewarder.

Table 2. *Model effects for performance analysis, predicting participant's number of catches during a 25-trial bin.*

Predictor	Fixed Effect	t	p
Intercept	7.38	67.30	<0.001*
SIAS	-5.56 e^{-5}	-0.02	0.99
Position (of bin in block)	0.13	4.53	<0.001*
Block number (linear)	-0.14	-0.83	0.41
Block number (quadratic)	0.07	0.45	0.65
Block number (cubic)	-0.35	-2.15	0.03*
SIAS x position	1.29 e^{-4}	0.13	0.89
SIAS x block number (linear)	2.47 e^{-4}	0.05	0.96
SIAS x block number (quadratic)	2.44 e^{-3}	0.47	0.64
SIAS x block number (cubic)	1.12 e^{-3}	0.21	0.83
Position x block number (linear)	7.17 e^{-3}	0.12	0.90
Position x block number (quadratic)	0.01	0.20	0.84
Position x block number (cubic)	0.07	1.18	0.24
SIAS x position x block number (linear)	-3.59 e^{-5}	-0.02	0.99
SIAS x position x block number (quadratic)	-6.07 e^{-4}	-0.32	0.75
SIAS x position x block number (cubic)	-6.83 e^{-4}	-0.36	0.72
	Random Effect Variance		
Participant intercept	0.42		
Location intercept	2.40 e^{-4}		

Note: SIAS = Social Interaction Anxiety Scale.

Table 3. *Model effects for role by social anxiety analysis, predicting participant's number of throws to an avatar during a 25-trial bin.*

Predictor	Fixed Effect	t	p
Intercept	2.42	18.11	<0.001*
SIAS	$3.82\ e^{-3}$	0.88	0.38
Role (P vs N)	-0.81	-4.28	<0.001*
Role (R vs N)	1.27	6.72	<0.001*
Prev role (P vs N)	-0.28	-1.47	0.14
Prev role (R vs N)	0.73	3.86	<0.001*
Position (of bin in block)	0.01	1.41	0.16
Block number (linear)	0.02	0.45	0.66
Block number (quadratic)	.02	0.33	0.74
Role (P vs N) x prev role (P vs N)	0.12	0.40	0.69
Role (R vs N) x prev role (P vs N)	-0.03	-0.12	0.91
Role (P vs N) x prev role (R vs N)	-0.01	-0.27	0.78
Role (R vs N) x prev role (R vs N)	-0.30	-1.01	0.31
Role (P vs N) x position	-0.18	-2.54	0.01*
Role (R vs N) x position	-0.32	4.58	<0.001*
Prev role (P vs N) x position	0.03	0.38	0.70
Prev role (R vs N) x position	-0.06	-0.94	0.35
Role (P vs N) x SIAS	$-3.99\ e^{-3}$	-0.65	0.52
Role (R vs N) x SIAS	$-4.60\ e^{-3}$	-0.76	0.45
Prev role (P vs N) x SIAS	-0.01	-2.30	0.02*
Prev role (R vs N) x SIAS	0.01	1.09	0.28
Position x SIAS	$-7.52\ e^{-4}$	-0.48	0.63
Role (P vs N) x prev role (P vs N) x position	0.03	0.31	0.76
Role (R vs N) x prev role (P vs N) x position	-0.08	-0.88	0.38
Role (P vs N) x prev role (R vs N) x position	-0.08	-0.85	0.39
Role (R vs N) x prev role (R vs N) x position	0.18	1.62	0.10
Role (P vs N) x prev role (P vs N) x SIAS	0.01	1.07	0.29
Role (R vs N) x prev role (P vs N) x SIAS	-0.01	-0.75	0.45
Role (P vs N) x prev role (R vs N) x SIAS	-0.01	-0.89	0.37
Role (R vs N) x prev role (R vs N) x SIAS	0.01	1.57	0.12
Role (P vs N) x position x SIAS	$4.43\ e^{-4}$	0.20	0.84
Role (R vs N) x position x SIAS	$1.04\ e^{-3}$	0.47	0.64
Role (P vs N) x position x SIAS	$2.10\ e^{-3}$	0.93	0.35
Role (R vs N) x position x SIAS	$-3.09\ e^{-3}$	-1.41	0.16
Role (P vs N) x prev role (P vs N) x position x SIAS	$-1.77\ e^{-3}$	-0.49	0.63
Role (R vs N) x prev role (P vs N) x position x SIAS	$2.32\ e^{-3}$	0.79	0.43
Role (P vs N) x prev role (R vs N) x position x	$3.31\ e^{-3}$	1.12	0.26

SIAS			
Role (R vs N) x prev role (R vs N) x position x SIAS	-0.01	-1.51	0.13
	Random Effect Variance		
Location intercept	5.58 e^{-15}		

Note: P = punisher, N = neutral, R = rewarder, Prev role = role of that avatar in the previous block, SIAS = Social Interaction Anxiety Scale.

Table 4. *Model effects for learning rate analysis, predicting the lag from a role switch to a peak in learning rate.*

Predictor	Fixed Effect	t	p
Intercept	3.45	8.56	<0.001*
SIAS	-0.01	-0.55	0.59
Prev role of R (P vs N)	0.29	0.90	0.37
Prev role of R(R vs N)	-0.43	-0.85	0.39
Prev role of P (P vs N)	-0.16	-0.31	0.76
Prev role of P (R vs N)	0.16	0.48	0.63
SIAS x prev role of R (P vs N)	-0.02	-1.76	0.08
SIAS x prev role of R(R vs N)	0.02	0.87	0.39
SIAS x prev role of P (P vs N)	-0.01	-0.76	0.45
SIAS x prev role of P (R vs N)	$2.99\ e^{-3}$	-0.28	0.78
	Random Effect Variance		
Participant intercept	2.67		
Location intercept	0.11		

Note: P = punisher, N = neutral, R = rewarder, Prev role = role of that avatar in the previous block, SIAS = Social Interaction Anxiety Scale.

Figure 1. Schematic of role switches in the Volatile Social Learning Task. Each of the six orders of role switches is depicted separately. Additionally, the starting location (right, left, and top) of each avatar was counterbalanced, so each of these six orders started from six possible locations, leading to thirty-six versions of the task to which participants were randomly assigned.

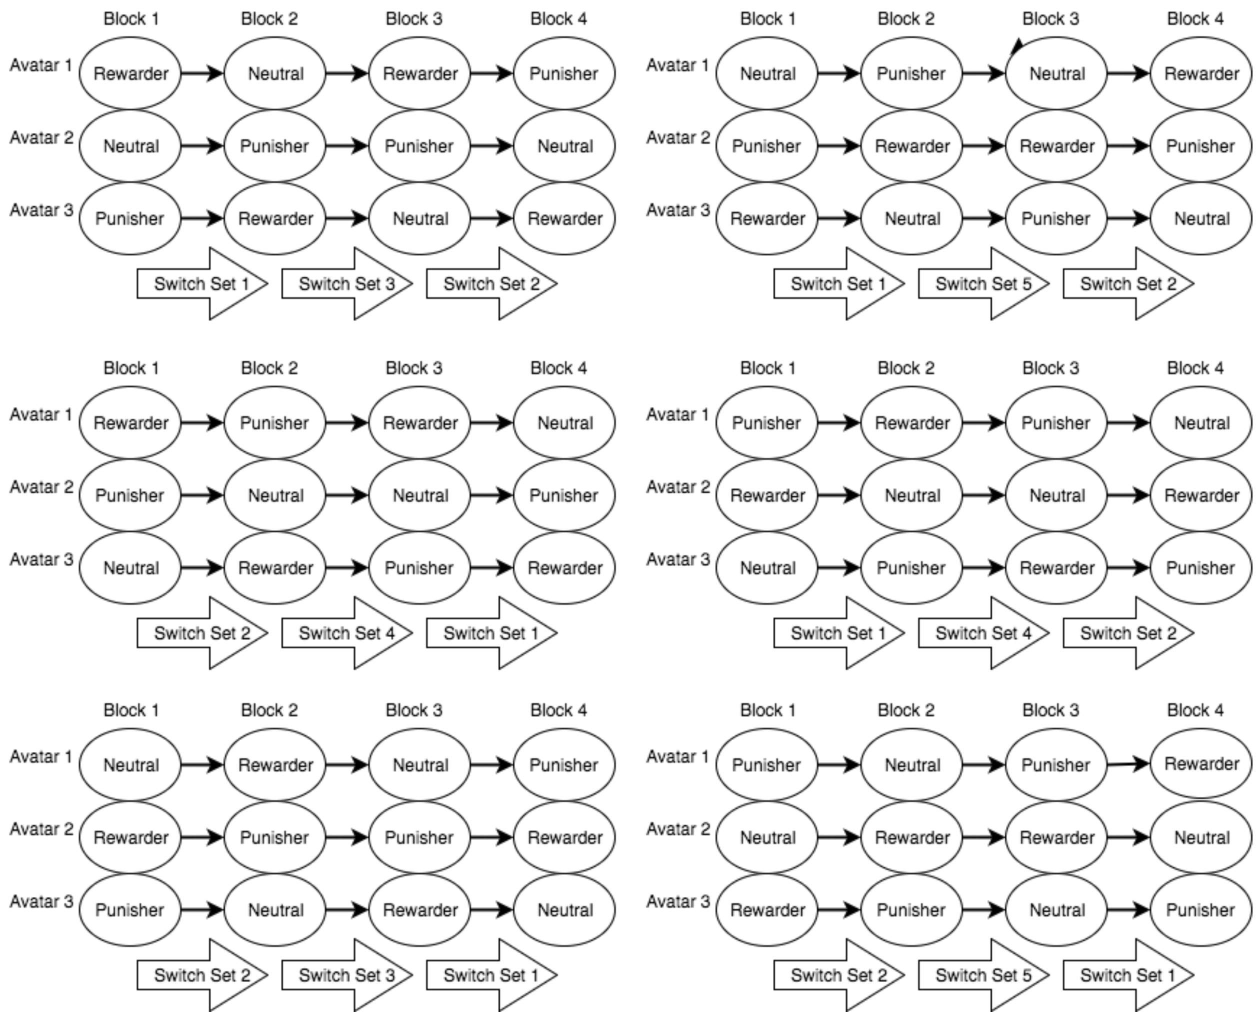

Figure 2. Participant throws to each avatar role over the course of the Volatile Social Learning Task, grouped into 25-trial bins.

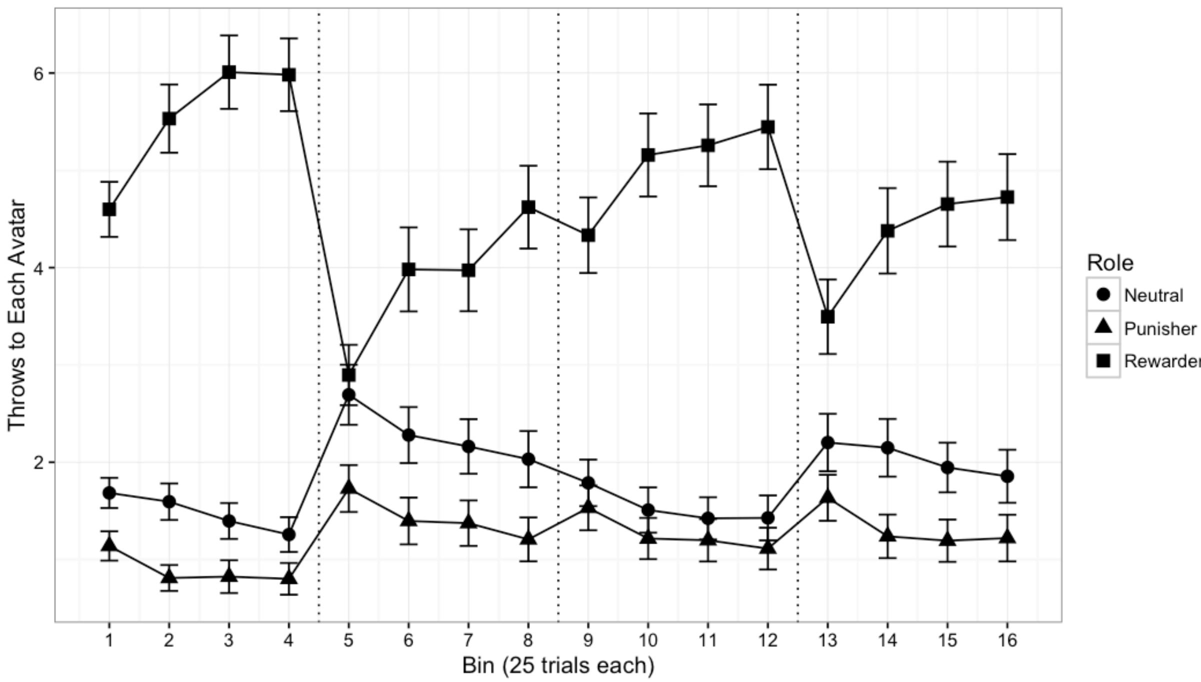

Note: error bars represent the 95% confidence interval. Dotted vertical lines represent role shifts between blocks.

Figure 3. Average participant throws to each avatar role over the Volatile Social Learning Task, grouped into 100-trial blocks.

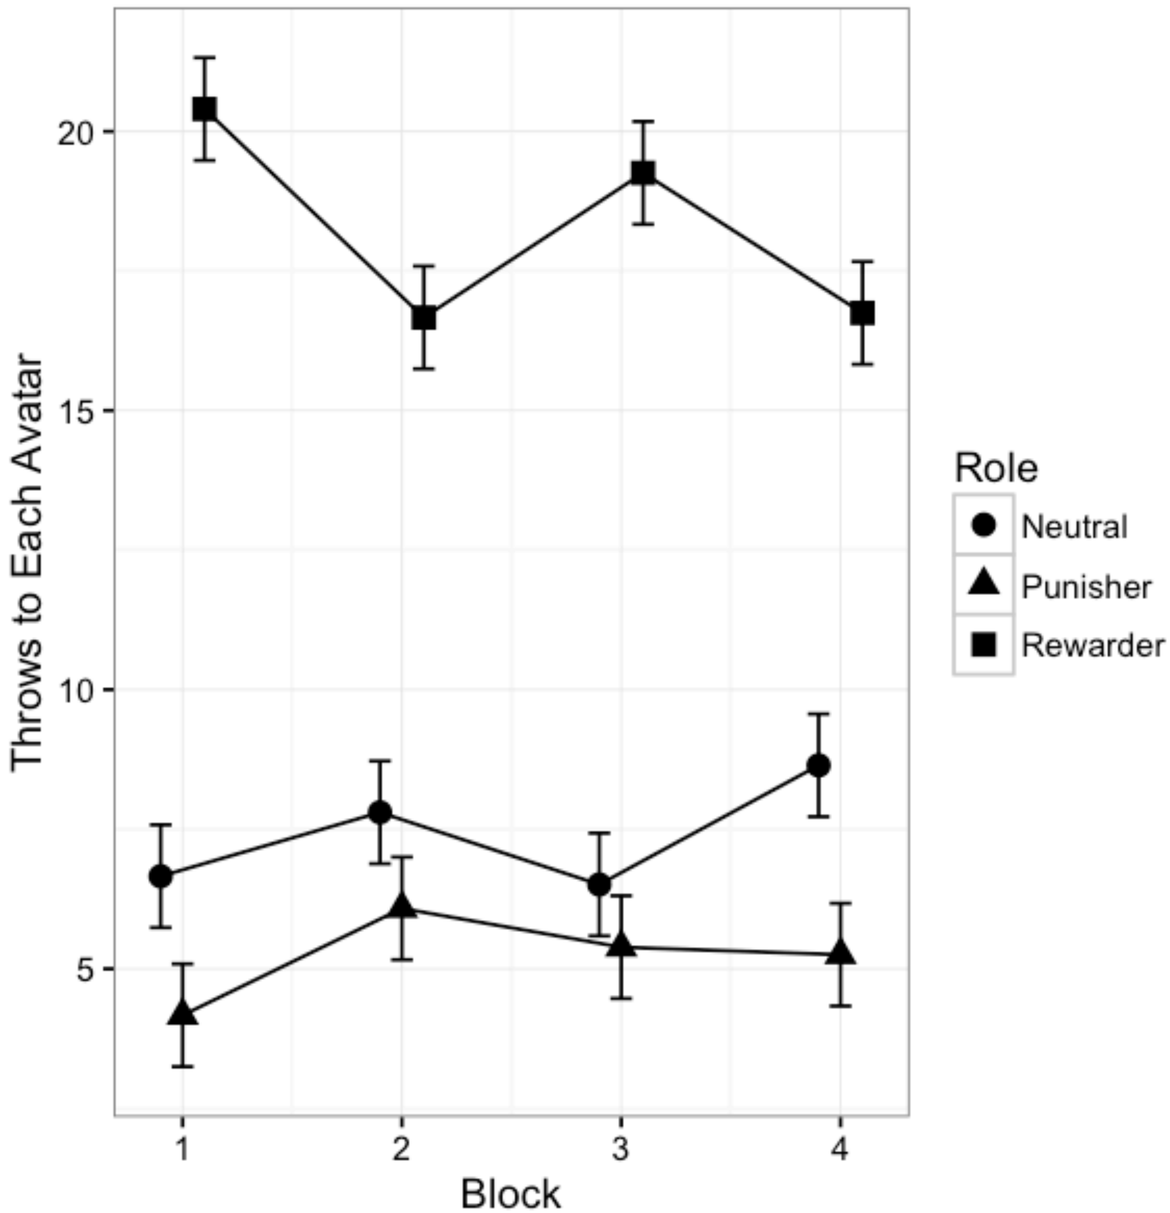

Note: error bars represent the 95% confidence interval.

Figure 4. Average participant throws to an avatar in each 25-trial bin, grouped by social anxiety symptom level and previous role of the avatar.

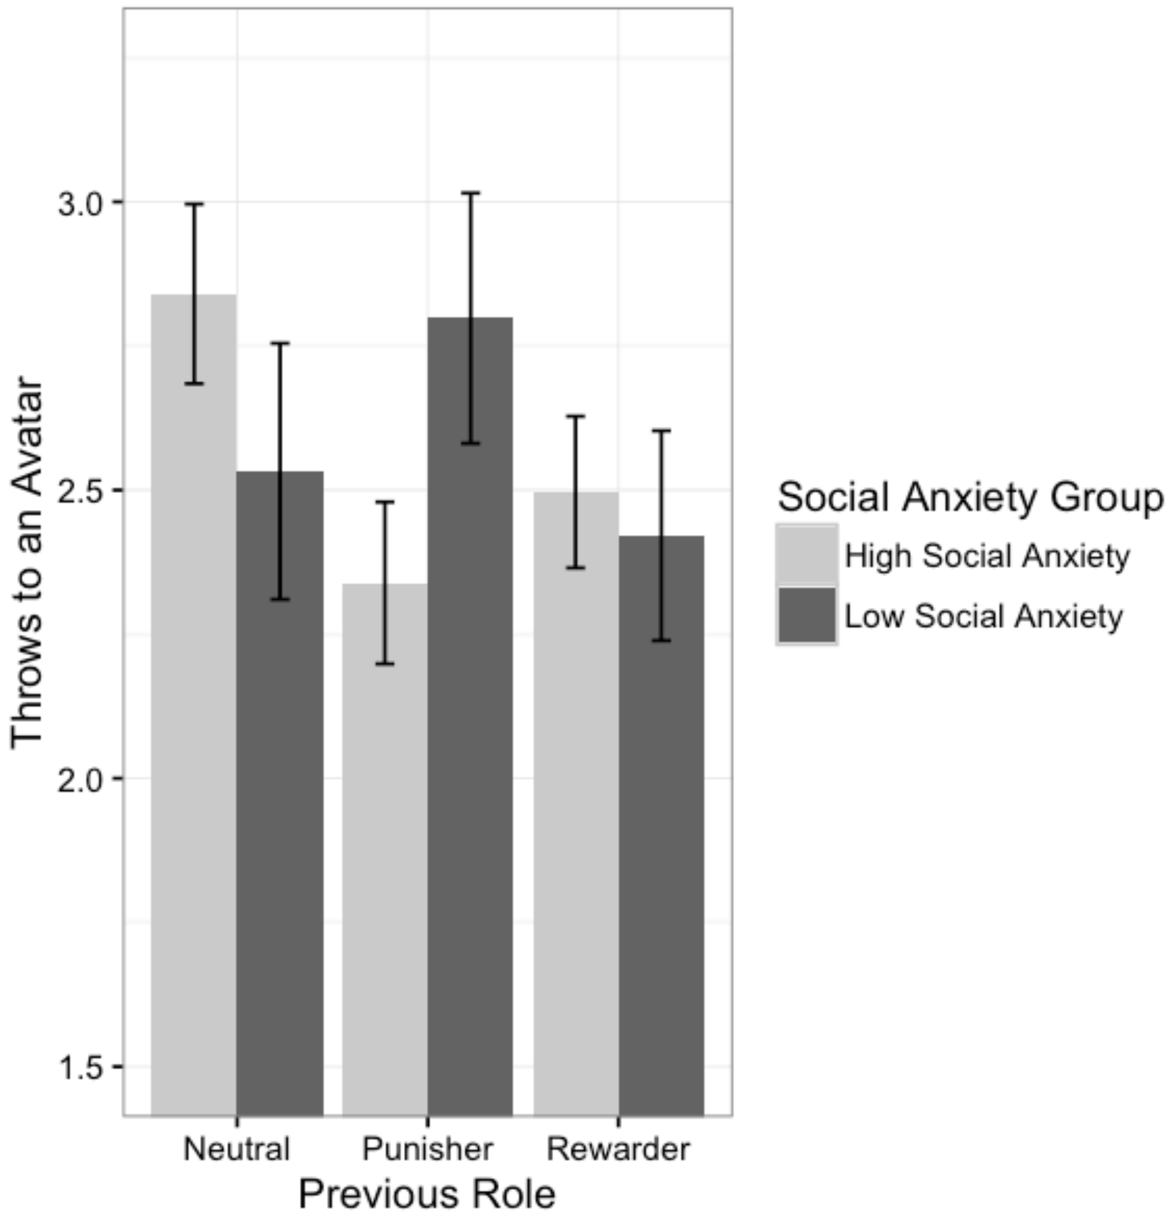

Note: error bars represent 95% confidence intervals. Although social anxiety is split into extreme groups for this figure, analyses reported were performed on social anxiety as a continuous variable. Here, participants with SIAS scores less than or equal to three quarters of a standard deviation (10 or under) below the mean of a previous community sample (*M*=18.8, *SD*=11.8;

Mattick & Clarke, 1998) were included in the low social anxiety group, and those scoring greater than or equal to three quarters of a standard deviation (28 or greater) above the mean were included in the high social anxiety group. Participants higher in social anxiety were less likely to throw to the avatar that had been the punisher in the prior block, independent of that avatar's current role.

Appendix

An initial version of the Volatile Social Learning Task (VSLT) was piloted with shorter blocks (70 trials each) and slightly different instructions. Examination of pilot results with N=44 online mTurk participants suggested that many participants aimed for fairness, throwing to less-targeted avatars so that all would receive the ball equally, or did not change their behavior in response to role switches. For the version of the task used in analyses reported here, VSLT instructions were modified to underscore the goal of receiving the ball as much as possible by throwing to the avatar most likely to return the ball, and comprehension questions were added (see Figure A1 for the final instructions and Figure A2 for the comprehension questions). Participants received corrective feedback if they chose the wrong answer on the comprehension questions and could not proceed to the game until correcting their answers (see Figure A3 for corrective feedback). Block length was also increased to 100 trials to increase the likelihood that participants would learn the correct avatar roles by the end of the block. In both the piloted version and the final version of the VSLT, the participant played the game with three computerized avatars (see Figure A4).

Figure A1. Instructions on VSLT.

Cyberball

Welcome to Cyberball, the Interactive Ball-Tossing Game!

You will play an online ball-tossing game with other players.

The game is very simple. When you receive the ball, click on the name of another player to throw the ball to them.

Your goal is to catch the ball as much as possible.

Throw the ball to whomever you think is most likely to throw the ball back to you.

Keep in mind that some players may throw the ball to you more often than others, and this may change over time.

The game will take about 20 minutes to complete. When the game is over, you will be directed to the next part of the study, which includes questionnaires.

Please click on the "Next" button below to answer a few questions, then begin the game:

Figure A2. Comprehension questions before starting VSLT.

What is my goal in the game?

- ○ For everyone to get the ball equally
- ○ To finish as quickly as possible
- ○ To maximize how many times I receive the ball

How do I maximize the number of times I receive the ball?

- ○ Throw the ball to everyone equally
- ○ Figure out who throws the ball to me the most, and then throw it to them

Once I figure out who throws the ball to me the most, will they keep throwing it to me the most for the whole game?

- ○ Yes
- ○ No

Figure A3. Corrective feedback for comprehension questions before starting VSLT.

Your goal is to maximize the number of times you receive the ball.

What is my goal in the game?

- ○ For everyone to get the ball equally
- ◉ To finish as quickly as possible
- ○ To maximize how many times I receive the ball

Figure out who throws the ball to you the most, then throw the ball to them.

How do I maximize the number of times I receive the ball?

- ◉ Throw the ball to everyone equally
- ○ Figure out who throws the ball to me the most, and then throw it to them

Players can change their throwing patterns over the course of the game.

Once I figure out who throws the ball to me the most, will they keep throwing it to me the most for the whole game?

- ◉ Yes
- ○ No

Figure A4. Screenshot of the VSLT.