Published in final edited form as:

Stress. 2013 May ; 16(3): 289–299. doi:10.3109/10253890.2012.744391.

Enhanced Avoidance Learning in Behaviorally-Inhibited Young Men and Women

Jony Sheynin^{1,2,6}, Saima Shikari³, Mark A. Gluck⁴, Ahmed A. Moustafa^{5,6}, Richard J. Servatius^{1,2,6}, and Catherine E. Myers^{2,6,7}

¹Joint Biomedical Engineering Program, New Jersey Institute of Technology and Graduate School of Biomedical Sciences, University of Medicine and Dentistry of New Jersey, Newark, NJ

²Stress & Motivated Behavior Institute, New Jersey Medical School, University of Medicine and Dentistry of New Jersey, Newark, NJ

³Honors College, Rutgers University, Newark, NJ

⁴Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ

⁵Marcs Institute for Brain and Behaviour & School of Social Sciences and Psychology, University of Western Sydney, Sydney, NSW, Australia

⁶Department of Veterans Affairs, New Jersey Health Care System, East Orange, NJ

⁷Department of Psychology, Rutgers University, Newark, NJ

Abstract

Behavioral inhibition (BI) is a temperamental tendency to avoid or withdraw from novel social and nonsocial situations and has been shown to predispose individuals to anxiety disorders. However, adequate means to assess individual differences in avoidance learning in humans are presently limited. Here, we tested whether individuals with high self-reported BI show faster associative learning on a purely cognitive task, and whether such inhibited individuals are more prone to avoid aversive outcomes. In Experiment 1, we tested 74 healthy undergraduate students (mean age 19.5 years; 55.4% female) on a computer-based probabilistic classification task, where participants were asked to classify four distinct visual stimuli into two categories. Two stimuli were associated with reward (point gain) and two were associated with punishment (point loss). In Experiment 2, 79 participants from the same population (mean age 19.8 years; 62% female) were tested on a novel modification of the same task, where they also had the option to opt out of responding on each trial and thus, avoid any chance of being punished (or rewarded) on that trial. Results show that inhibited participants demonstrated better associative learning in Experiment 1, while exhibiting a greater tendency to opt out in Experiment 2 (repeated-measures ANOVAs, main effects of BI, both $p < 0.05$). These results suggest that the facilitated classically-conditioned learning previously observed in inhibited individuals can be extended to a cognitive task, and also highlight a specific preference in inhibited individuals for withdrawal (“opting-out”) as a response strategy, when multiple strategies are available to avoid punishment.

Corresponding Author: Jony Sheynin, VA-NJHSC, 385 Tremont Avenue, Mailstop 129, East Orange, NJ 17018, Phone: (973) 676-1000 x 2220, sheynijo@umdnj.edu.

Declaration of Interest

This work was supported by the NSF/NIH Collaborative Research in Computational Neuroscience (CRCNS) Program, by NIAAA (R01 AA018737), by Award Number I01CX000771 from the Clinical Science Research and Development Service of the VA Office of Research and Development and by additional support from the Stress & Motivated Behavior Institute (SMBI).

Keywords

Decision making; avoidance; PTSD (posttraumatic stress disorder); behavioral inhibition; associative learning; anxiety disorders

Introduction

Decision-making is probably one of the most important challenges that one faces daily, and it provides crucial feedback that directly affects future choices. There are two distinct behavior strategies that ideally, should be employed together, in order to maximize benefit and minimize loss; individuals can learn and repeat actions that were rewarded in the past, or alternatively, they can avoid actions that are associated with past loss (Wachter et al., 2009). Yet, these strategies are not equivalent with regards to adverse mental health outcomes; the tendency to avoid may have implications for anxiety and the development of anxiety disorders.

Avoidance is a common symptom of all anxiety disorders (DSM-IV; American Psychiatric Association, 1994) and its severity parallels the overall growth and persistence of some of the disorders. For example, after exposure to a traumatic event, the degree of avoidance behavior differentiates those who develop post-traumatic stress disorder (PTSD) and those who recover (Karamustafalioglu et al., 2006; Marshall et al., 2006; North et al., 1999; O'Donnell et al., 2007). Although there is a growing appreciation that personality can shape one's tendency to avoid, and despite the direct relevance of avoidance to anxiety disorders, assessment of human avoidance behavior has been very limited (Dymond and Roche, 2009). In an experimental setting, avoidance can be assessed by measuring an individual's ability to learn to avoid or escape physically aversive stimuli, such as electric shock (Delgado et al., 2009), or aversive visual and auditory stimuli such as loud noises or unpleasant pictures (Dymond et al., 2011). Research has also been extended to computer-based tasks where the aversive outcome is point loss or monetary loss (Declercq and De Houwer, 2008; Molet et al., 2006; Schlund and Cataldo, 2010). However, since these tasks are often designed so that avoidance is the only possible way to prevent the aversive outcome, participants who perform well might simply be demonstrating facilitated associative learning, rather than a tendency to acquire avoidance behaviors in preference to other kinds of responses; tools to dissociate these two behavior aspects are lacking.

Behavioral inhibition (BI), a vulnerability factor for anxiety disorders (Rosenbaum et al., 1993), is a relatively stable personality trait conferring a tendency to avoid or withdraw from novel social and non-social situations (Kagan et al., 1989). A recent study showed a correlation between self-reported BI score and self-reported PTSD avoidance symptoms in veterans (Myers et al., 2012c). Moreover, inhibited individuals learn faster than their uninhibited peers on classical eyeblink conditioning paradigm, where a conditioned response (eyeblink) is produced by a conditioned stimulus (tone) that predicts an aversive unconditioned stimulus (airpuff) (Myers et al., 2012b). Research with animal models of BI has also showed enhanced learning by inhibited subjects of both an operantly-conditioned avoidance response (Servatius et al., 2008) and classically-conditioned eyeblink response (Ricart et al., 2011). However, much less is known about how BI might modulate associative learning where the outcomes are not physically aversive stimuli such as airpuffs, but rather on-screen feedback such as point loss in a computer task. In particular, rather than promoting a specific type of learning or response (e.g., acquiring somatomotor responses that promote avoidance of an aversive stimulus), individuals with inhibited temperament may show non-specific enhancement in associative learning. It is also unknown whether individuals with inhibited temperament exhibit different behavior patterns than uninhibited

peers on tasks which involve probabilistic and/or ambiguous contingencies. Ambiguous contingencies might lead to uncertainty, which can lead to significant anxiousness, especially among those who are intolerant for such feelings (Buhr and Dugas, 2002). It is possible that inhibited individuals will show higher tendency to avoid such uncertain situations.

Therefore, the central questions addressed in this work are (1) Do inhibited individuals learn faster than uninhibited individuals on a cognitive task, where the aversive outcome to be avoided is merely point loss and the response-outcome contingencies are probabilistic and ambiguous? (2) If so, is this reflected in a general facilitation of learning in inhibited individuals, or is the facilitation selective to learning to avoid punishment, rather than learning to obtain rewards? To answer those questions, we used a computer-based probabilistic classification task that has previously been used to demonstrate individual differences (Bódi et al., 2009). This task interleaves trials in which participants learn to make a correct classification in order to obtain a reward (point gain) and trials in which participants learn to make a correct classification in order to avoid a punishment (point loss). In Experiment 1 (“Forced-Choice”), we utilized the task to test if there are differences between behaviorally inhibited vs. uninhibited healthy young adults on reward and punishment learning. We expected that inhibited participants would outperform their uninhibited counterparts on the classification task, similar to what was shown in the eyeblink experiments. In Experiment 2 (“Opt-Out”), we considered a modified version of the task in which participants had the option to opt out of responding on any trial (rather than being forced to make a classification decision), and thereby avoid any chance of being punished (or rewarded) on that trial. We expected that inhibited participants would show more opt-out responding than uninhibited participants, particularly on trials where a punishment outcome was possible. If so, this would suggest that the inhibited individuals not only show generally faster associative learning, but also a specific preference to respond to the threat of aversive outcomes by withdrawal.

Methods

Participants

In Experiment 1, participants were 74 healthy young adults (Rutgers University undergraduates; mean age 19.49 years, SD 1.55; 55.4% female), recruited via a departmental subject pool, in which available research studies are posted and students sign up to participate in exchange for research credits in a psychology class. Experiment 2 included 79 participants from the same institution (mean age 19.82 years, SD 1.62; 62% female). None had participated in Experiment 1. Most participants were recruited through the departmental subject pool, as for Experiment 1. However, an additional nine participants were recruited via flyers posted around the campus and received cash payments in the amount of \$20. No difference on any behavioral, demographic, or questionnaire measure was observed between participants that received cash vs. participants that received credits (all $p > 0.200$); analyses were therefore pooled over all 79 participants. In both experiments, procedures followed guidelines established by Rutgers University and the Declaration of Helsinki for the protection of human subjects. All participants provided written informed consent before initiation of any experimental procedures.

Procedure

Participants were tested individually; the participant and experimenter sat in a quiet testing area during the experiment. Participants completed the Adult Measure of Behavioural Inhibition (AMBI), a self-report questionnaire, that assesses current temperamental tendency to respond to new stimuli with inhibition and/or avoidance, has high internal consistency

(Cronbach's $\alpha=0.87$) and high test-retest consistency (intraclass correlation coefficient 0.86, $p<0.001$), and has been shown to be a measure of anxiety proneness (Gladstone and Parker, 2005). Individuals scoring less than 16 are classified as "uninhibited", whereas individuals scoring 16 or higher are classified as "inhibited."

In addition, participants in Experiment 1 were administered a computer-based probabilistic classification task; full description of the task is given in Bódi et al. (2009). In brief, on each of 160 trials, participants viewed one of four images (stimuli S1, S2, S3 and S4; Figure 1A), and learned whether it belonged to category A or category B. On any given trial, stimuli S1 and S3 belonged to category A with 80% probability and to category B with 20% probability; stimuli S2 and S4 belonged to category B with 80% probability and to category A with 20% probability. Stimuli S1 and S2 were used on reward-learning trials; if the participant made a correct classification response, a reward of +25 points was received (Figure 1B) but if the participant made an incorrect classification response, no feedback message appeared. Stimuli S3 and S4 were used in punishment-learning trials; if the participant made an incorrect classification response on a trial with either of these stimuli, a punishment of -25 points was received, but correct classification received no feedback message. Thus, the no-feedback outcome, when it arrived, was ambiguous, as it could signal lack of reward for an incorrect response (if received during a trial with S1 or S2) or lack of punishment for a correct response (if received during a trial with S3 or S4). Participants did not receive any monetary reward related to point accumulation during the experiment.

Trials were divided into four blocks of 40 intermixed trials, with each stimulus appearing 10 times per block (for a total of 20 reward and 20 punishment trials intermixed per block).

Participants in Experiment 2 received the same task, except that on each trial, in addition to responding "A" or "B", the participant could choose a third option, "skip," which allowed the participant to opt out of responding on that trial. If the subject chose this option, the following acknowledgment appeared on the screen "OK, skipping this one...", and the program advanced to the next trial without further feedback and without point gain or loss.

Data analysis

Statistical analyses were conducted using SPSS version 17.0 (SPSS Inc., Chicago, IL). Alpha was set to 0.05 (two-tailed), with Bonferroni correction used as appropriate to protect against inflated risk of family-wise error. Effects and interactions that did not approach significance were not reported (all $p>0.100$). Questionnaire scores were analyzed using *t*-test for continuous values and chi-square for categorical values, with Yates continuity correction for 2x2 tables. Internal consistency of the AMBI questionnaire was analyzed using Cronbach's α .

In Experiment 1, for each participant and each trial, the computer recorded whether the participant's classification response was correct and whether it was optimal. For a reward trial, the response was correct if it resulted in point gain; for a punishment trial, the response was correct if it resulted in no point loss. For each stimulus, the response was optimal if the participant chose the category that was most often associated with that stimulus. For example, stimulus S1 belonged to category A on 80% of trials; thus, response "A" was always optimal for that stimulus, but it would only be correct (i.e., result in point gain) on 80% of trials, because on the remaining 20% of trials S1 would belong to category B. Percent optimal and total correct classification responses were calculated separately for reward and punishment-learning trials. Repeated-measures analysis of variance (rmANOVA) was used to assess responding, with between-subject factor of BI (inhibited vs. uninhibited) and within-subject factors of feedback type (reward vs. punishment trials) and, in some cases, trial block (four blocks of 20 trials with each feedback type). Post-hoc

ANOVAs and *t*-tests were performed, as appropriate. When a correlation between continuous behavior measures was performed, a Pearson correlation coefficient was used.

We performed several further analyses, to test whether other cognitive factors could account for any observed behavioral differences between BI groups. First, we calculated a learning bias score for each participant, as the difference between percent optimal classification responses on reward vs. punishment trials; calculation was done separately for the first and the last blocks. Following Simon et al. (2010), we defined as “balanced learners” those participants whose absolute value of the bias score on the first block was less than 25% (i.e., no strong difference in performance on reward vs. punishment trials). Next, we compared the stability of bias across first vs. last block. We defined “stable learners” as those participants who had the same bias (balanced or unbalanced) on the first and the last block. Pearson chi-square was used to compare the distribution of bias balance and stability between BI groups, with Yates continuity correction for 2×2 tables.

To further investigate possible learning biases, we considered trial-to-trial behaviors. Following previous work (Frank and Kong, 2008; Simon et al., 2010), “win-stay” behavior was defined as the percentage of trials in the first block on which the participant repeated the response which was correct (i.e., either rewarded or not punished) on the last trial where the same stimulus appeared. “Lose-shift” behavior was defined as the percentage of trials on which the participant switched his or her response after making an incorrect response (i.e., response that was punished or not rewarded) on the last time the same stimulus appeared. Pearson correlation coefficient was used to analyze correlation between win-stay and lose-shift behaviors. We used rmANOVA to examine percent of trials where participants displayed win-stay or lose-shift behaviors, with BI as the between-subject factor and behavior type (win-stay vs. lose-shift) as the within-subject factor.

Lastly, response latencies in milliseconds were assessed for each feedback type (reward vs. punishment) and response (optimal vs. non-optimal). For each participant, outlying values of latency (>3 SD from mean) were excluded from the analyses. We used rmANOVA on average latency, with BI as the between-subject factor and response and feedback type as within-subject factors.

In Experiment 2, for each trial, the computer recorded whether the participant chose to opt out of responding or, if not, whether the classification response was correct and/or optimal. Dependent variables were total number of opt-out responses to reward and punishment stimuli, as well as percent optimal and total correct classifications on those trials where a classification response was made. However, considering only percent optimal (or total correct) classifications potentially obscures learning in participants who may have learned to make skipping responses to punishment trials in order to minimize risk of punishment. Therefore, we also considered a measure of “adaptive responses” which included all punishment trials on which either an opt-out response or an optimal classification was made; for reward trials, only optimal classifications were “adaptive responses” since on these trials opting-out would cause omission of a possible reward. Participants were also divided into “skippers” vs. “non-skippers”, where participants who made at least one opt-out response during the experimental session were categorized as “skippers.” Such “skipping” classification was often included as a between-subject factor, enabling comparison of performance between skippers and non-skippers. Otherwise, data analysis was similar to Experiment 1.

Results

Experiment 1: “Forced-Choice”

Mean score on AMBI questionnaire was 13.450 (SD 4.561, Cronbach's $\alpha=0.727$). Twenty-two participants (29.7%) were “inhibited.” There were no differences in gender or age between BI groups (both $p>0.300$).

Participants' mean optimal classification performance on reward and punishment trials is shown in Figure 2; inhibited participants outperformed uninhibited participants [rmANOVA, $F(1,72)=4.620$, $p=0.035$]. A main effect of trial block [$F(3,216)=12.844$, $p<0.001$] was found, and the interaction between trial block and BI was not significant [$F(3,216)=2.209$, $p=0.088$]. Total correct responding for reward and punishment trials revealed a similar main effect of BI [rmANOVA, $F(1,72)=5.574$, $p=0.021$], with inhibited participants making more correct responses than uninhibited participants; no other effects or interactions were found.

Next, we tested whether the differences between inhibited and uninhibited participants could be attributed to differences in learning bias (Figure 3). In both BI groups, approximately 70% of participants were balanced and a small majority of around 55% were classified as stable learners (both $p>0.800$).

There was also a negative correlation between win-stay and lose-shift behaviors on the first block ($r=-0.544$, $p<0.001$). Overall, participants showed a higher percentage of win-stay than lose-shift responses [$F(1,72)=20.750$, $p<0.001$], with no main effect of BI and no interactions (Figure 4).

Lastly, when latency was analyzed, one inhibited and one uninhibited participant were excluded since they did not have any non-optimal reward trials. Overall, participants had greater latency on punishment than reward trials [$F(1,70)=29.108$, $p<0.001$] (Figure 5); the main effect of response was not significant [$F(1,70)=3.472$, $p=0.067$]. There was no effect of BI and no interactions.

Experiment 2: “Opt-Out”

Mean score on AMBI questionnaire was 14.770 (SD 5.359, Cronbach's $\alpha=0.790$). Twenty-nine of the participants (36.7%) were classified as “inhibited” and 44 participants (55.7%) as “skippers.” In both BI groups, almost half of the participants did not show any skipping responses (Figure 6A). There were no significant differences in gender or age between BI groups, and no differences in gender, age or AMBI scores between skippers vs. non-skippers (all $p>0.300$).

Opt-out performance of participants classified as skippers is shown in Figure 6B–C. Inhibited participants skipped more trials overall than uninhibited participants [rmANOVA, $F(1,42)=5.887$, $p=0.020$]. There were also main effects of trial block [$F(3,126)=17.744$, $p<0.001$], with skipping behavior increasing over trial blocks, and feedback type [$F(1,42)=34.383$, $p<0.001$], with participants skipping more punishment than reward trials. There was an interaction between trial block and feedback type [$F(3,126)=21.799$, $p<0.001$]. Post-hoc tests revealed a significant effect of block only in the punishment trials [$F(3,129)=20.064$, $p<0.001$], possibly due to a floor effect (minimal skipping) in the reward trials. Although six participants skipped 20% of punishment trials on the very first block of training, none of these participants skipped $>5\%$ of the reward trials on this block.

Figure 7 shows percent optimal classification responses on all reward and punishment trials that were not skipped, in participants categorized as skippers vs. non-skippers. Eleven

participants (13.9%) were not included in the following analysis, since these participants skipped all the trials on at least one block and thus, did not have classification data that could be analyzed on that block. There were no significant differences in gender or age between those 11 participants and the remaining 33 skippers (both $p > 0.200$). However, this group had significantly higher BI scores [$t(42) = 2.753$, $p = 0.009$], which is consistent with the finding that inhibited participants showed more skipping responses. Among the remaining 68 participants, there were main effects of BI [rmANOVA, $F(1,64) = 6.839$, $p = 0.011$] and trial block [$F(3,192) = 7.023$, $p < 0.001$] on optimal classification performance, showing that although both groups improved performance across blocks, classification performance was better in uninhibited participants. There was also an interaction between trial block and feedback type [$F(3,192) = 8.527$, $p < 0.001$], and an interaction between trial block, feedback type and skipping [$F(3,192) = 5.241$, $p = 0.002$]. Post-hoc tests revealed that among skippers (Figures 7C–D), classification performance on reward trials was better than on punishment trials on the last block [$t(35) = 3.336$, $p = 0.002$]; this difference did not occur in non-skippers (Figures 7A–B). Total correct responding was also higher in uninhibited than inhibited participants [$F(1,77) = 15.134$, $p < 0.001$]. Specifically, uninhibited participants received less punishments [$t(47.5) = 2.801$, $p = 0.007$], and more rewards [$t(77) = 2.139$, $p = 0.036$]. There were no other main effects or interactions.

The finding of better classification performance in uninhibited participants in Experiment 2 contrasted the results of Experiment 1, where inhibited participants showed better classification performance. Further, on punishment trials in Experiment 2, frequency of opt-out responses was negatively correlated with optimal classification performance ($r = -0.352$, $p = 0.019$). Since skipping of punishment trials can be considered an adaptive response, because it eliminates the risk of punishment on that trial, we repeated the previous ANOVA analysis on “adaptive responses,” counting either skipping responses or optimal classification responses for punishment trials (Figures 7B, E), but only optimal classification responses for reward trials (Figures 7A, C). Now, data from all 79 participants could be included. Most importantly, in contrast to the previous analysis, this time no main effect of BI was found ($p = 0.196$). Other main effects of trial block [$F(3,225) = 21.178$, $p < 0.001$], feedback type [$F(1,75) = 5.645$, $p = 0.020$] and skipping [$F(1,75) = 6.898$, $p = 0.010$] were observed, as was the interaction between trial block and skipping [$F(3,225) = 4.656$, $p = 0.004$]. The interaction between trial block and feedback type approached significance [$F(3,225) = 2.579$, $p = 0.054$]. Post-hoc t -tests revealed that the difference between skippers and non-skippers was significant only on the last block [$t(77) = 3.105$, $p = 0.003$].

Figure 8 depicts the mean latency on trials where a classification response was made. As in Experiment 1, one inhibited and one uninhibited participant were excluded since they did not have any non-optimal reward trials that could be analyzed. Overall, there were longer latencies in skippers than non-skippers [rmANOVA, $F(1,73) = 5.111$, $p = 0.027$], on punishment than reward trials [$F(1,73) = 21.023$, $p < 0.001$], and in participants choosing the non-optimal vs. optimal response [$F(1,73) = 8.597$, $p = 0.004$]. Although no main effect of BI was found ($p = 0.910$), there were interactions between BI and skipping [$F(1,73) = 7.915$, $p = 0.006$], as well as between feedback type and skipping [$F(1,73) = 5.903$, $p = 0.018$] and between performance and feedback type [$F(1,73) = 7.803$, $p = 0.007$]. To further investigate the BI $\times$ skipping interaction, post-hoc tests revealed that inhibited skippers responded significantly more slowly than inhibited non-skippers [$t(26) = 2.565$, $p = 0.016$], while the uninhibited participants did not show any difference between skippers and non-skippers ($p = 0.587$).

Lastly, as in Experiment 1, additional analyses were performed to verify that other potential differences between participants did not account for the different performance observed between BI groups. Learning bias and bias stability did not differ between inhibited and

uninhibited participants, or between skippers and non-skippers; all four groups had approximately 77% participants who were balanced learners (both $p > 0.700$) and 44.4–63.6% who were stable learners (both $p > 0.200$). Moreover, win-stay and lose-shift behavior strategies were not different between inhibited and uninhibited groups or between skippers and non-skippers (both $p > 0.800$).

Discussion

The aim of the present study was to better understand the relation between BI, a temperamental trait conferring vulnerability to anxiety disorders, and the acquisition and expression of avoidance behavior, a predominant symptom in anxiety disorders. We considered two types of avoidance: learning optimal classification responses that minimized risk of punishment, and “opt-out” responding that allowed the participant eliminate any risk of punishment. Previously, BI was linked to faster acquisition of a classically-conditioned eyeblink response (Myers et al., 2012b); in Experiment 1, inhibited participants showed facilitated classification learning to on-screen feedback in the form of verbal statements (“Correct”/“Incorrect”) and point gain or loss, which did not translate into money or any primary reinforcer or punisher. In contrast, a recent study which used the same task (Bódi et al., 2009) showed that harm avoidance is specifically associated with enhanced performance on punishment trials; this highlights a potential difference between BI and the related trait harm avoidance, which was proposed to be indicative of BI (Cloninger, 1987).

In Experiment 2, we tested participants’ tendency to withdraw or “opt out” of responding, when this option is available. Although many prior studies have examined forced choice behavior on classification and related tasks (e.g., Bódi et al., 2009; Frank et al., 2004), to our knowledge, this experiment is the first in which participants are given the option to opt out of trials. Surprisingly, almost half of the participants in Experiment 2 never took advantage of the opt-out response; none of the variables we investigated appeared to differentiate those who made the response at least once (“skippers”) vs. those who did not (“non-skippers”). Nevertheless, as hypothesized, inhibited participants opted out significantly more than their uninhibited peers. Figures 6B–C showed that skipping was generally selective to punishment trials, and tended to increase across trials, rather than represent as a pre-existing tendency. In fact, skipping was a negatively reinforced response, in which the removal of risk of punishment resulted in an increase in frequency of the response (a.k.a., omission training).

However, the main conclusion from Experiment 2 is that inhibited and uninhibited individuals may both achieve comparable performance on the task, indexed as adaptive responding, but they may preferentially use different strategies. Specifically, while uninhibited individuals learn to make classification responses to minimize probability of punishment, inhibited individuals are more likely to skip punishment trials altogether. While this latter strategy results in fewer overall punishments obtained, it may be that in real-world situations, a tendency to rely on such “opting-out” behavior may confer vulnerability for development of pathological avoidance.

In both experiments we ruled out several other aspects of behavior that could account for the differences between BI groups, including learning bias, learning strategy and latency. Our finding that the majority of participants showed balanced learning is consistent with previous work on a related probabilistic selection task (Simon et al., 2010), although that experiment did not consider BI. The finding that participants showed greater win-stay than lose-shift behavior is also consistent with previous literature (Frank and Kong, 2008; Simon et al., 2010). Lastly, in Experiment 2, an unexpected and intriguing latency difference emerged: inhibited skippers responded more slowly than inhibited non-skippers. Longer

latencies might reflect conflict that some inhibited individuals experience when determining how to respond to a potential threat.

Besides these core issues, the current study is one of a small but growing literature to use the AMBI questionnaire as a self-report measure for current (adult) BI and to show relationships between AMBI scores and behavior (Caulfield et al., 2011; Holloway and Servatius, 2010; Myers et al., 2012a; Myers et al., 2012b). In both experiments, AMBI scores were similar, with about one-third of participants classified as inhibited. Future research might consider whether other instruments assessing BI (such as the widely adopted BIS/BAS scale; Carver and White, 1994) would show similar associations.

The current study has several limitations that could be addressed in future work. First, it would be beneficial to consider a richer set of demographic variables such as ethnicity, marital status, history of psychiatric illness, brain injury and substance use/abuse. Second, larger samples might allow investigating skipping as a continuous variable, rather than classifying participants as skippers vs. non-skippers. Furthermore, it is unclear why only about half of the participants used the skipping option in Experiment 2. Motivation or demand characteristics might have been involved, such that participants saw little value in skipping or felt this would be considered “cheating.” It is possible that by adding a monetary reinforcement, more participants would be motivated to opt out of punishment trials. Finally, the negative correlation between optimal performance and skipping on punishment trials in Experiment 2 should be further investigated. Currently, it is unclear whether high rates of skipping impair learning, by reducing opportunities for feedback, or poor learning produces high rates of skipping, since participants who do not learn the optimal classification on specific stimuli may begin skipping trials with those stimuli. One way to explore this issue in future might be by providing feedback regarding the optimal classification response on skipped trials, so that skipping does not entail reduced feedback.

Moreover, uncertainty is an important attribute of a probabilistic classification task, because it is impossible to perfectly predict the correct response on every trial. Individuals who have high intolerance to uncertainty are also at higher risk for anxiety disorders (Buhr and Dugas, 2002), and it is possible that BI and intolerance of uncertainty are related constructs. It would be interesting to test individuals on version of the task that reduces ambiguity: either a deterministic version of the task where each stimulus is 100% predictive of category membership, as opposed to 80% in the current task, or one that eliminates the no-feedback outcome (after skipping response, or after classification that results in failure to obtain reward or successful avoidance of punishment). In such future studies, it may also be of interest to directly assess both BI and uncertainty of tolerance as personality constructs, to better understand their presumed relation and mutual involvement in human behavior patterns. Finally, a prior study has also associated performance on this probabilistic classification task (“Forced-Choice” version) and the midbrain dopaminergic system (Bódi et al., 2009). Given the prior literature linking BI with brain substrates including cerebellum, amygdala and hippocampus (e.g., Blackford et al., 2012; Blackford et al., 2009), an interesting direction for future work would be to explore brain activation in inhibited vs. uninhibited individuals on this task.

In conclusion, this work explored relationships between performance on a computer-based task and personality characteristics, specifically, trait BI. A main reason for studying behavior patterns in inhibited individuals is that BI confers increased risk for anxiety disorders, which entail pathologically exaggerated avoidance as a predominant symptom. The current work suggests that BI is associated with enhanced associative learning, which leads us to hypothesize that the exaggerated avoidance in anxious individuals may be a byproduct of facilitated associative learning rather than a stand-alone symptom. Our current

findings suggest that higher tendency to avoid is a pre-existing behavior pattern that can be observed in healthy young adults. As such, it might precede development of an anxiety disorder (and potentially contribute to its etiology and pathogenesis). Longitudinal studies would be required to determine whether non-anxious individuals who do show accelerated learning and/or high tendency to opt out on this task do indeed show higher long-term risk for development of anxiety disorders. Such results have implications for understanding the transition from temperamental vulnerabilities to symptomatic behaviors seen in anxiety disorders, and potentially, might assist with designing better therapeutic strategies to help anxious individuals.

References

- American Psychiatric Association. Diagnostic and Statistical Manual of Mental Disorders. 4. Washington, DC: APA; 1994.
- Blackford JU, Allen AH, Cowan RL, Avery SN. Amygdala and hippocampus fail to habituate to faces in individuals with an inhibited temperament. *Soc Cogn Affect Neurosci*. 2012 in press.
- Blackford JU, Avery SN, Shelton RC, Zald DH. Amygdala temporal dynamics: temperamental differences in the timing of amygdala response to familiar and novel faces. *BMC Neurosci*. 2009; 10:145. [PubMed: 20003287]
- Bódi N, Kéri S, Nagy H, Moustafa A, Myers CE, Daw N, Dibo G, Takats A, Bereczki D, Gluck MA. Reward-learning and the novelty-seeking personality: a between- and within-subjects study of the effects of dopamine agonists on young Parkinson's patients. *Brain*. 2009; 132:2385–2395. [PubMed: 19416950]
- Buhr K, Dugas MJ. The Intolerance of Uncertainty Scale: psychometric properties of the English version. *Behav Res Ther*. 2002; 40:931–945. [PubMed: 12186356]
- Carver CS, White TL. Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS Scales. *Journal of Personality and Social Psychology*. 1994; 67:319–333.
- Caulfield MD, McAuley JD, Zhu DC, Servatius RJ. Facilitated acquisition of eyeblink conditioning in those at risk for anxiety disorders: Concordance among scales of inhibited temperament. Society for Neuroscience Annual Meeting. 2011:684.617/GG629. [abstract].
- Cloninger CR. A systematic method for clinical description and classification of personality variants. A proposal. *Arch Gen Psychiatry*. 1987; 44:573–588. [PubMed: 3579504]
- Declercq M, De Houwer J. On the role of US expectancies in avoidance behavior. *Psychon Bull Rev*. 2008; 15:99–102. [PubMed: 18605487]
- Delgado MR, Jou RL, Ledoux JE, Phelps EA. Avoiding negative outcomes: tracking the mechanisms of avoidance learning in humans during fear conditioning. *Front Behav Neurosci*. 2009; 3:33. [PubMed: 19847311]
- Dymond S, Roche B. A contemporary behavior analysis of anxiety and avoidance. *Behav Anal*. 2009; 32:7–27. [PubMed: 22478511]
- Dymond S, Schlund MW, Roche B, Whelan R, Richards J, Davies C. Inferred threat and safety: symbolic generalization of human avoidance learning. *Behav Res Ther*. 2011; 49:614–621. [PubMed: 21767825]
- Frank MJ, Kong L. Learning to avoid in older age. *Psychol Aging*. 2008; 23:392–398. [PubMed: 18573012]
- Frank MJ, Seeberger LC, O'Reilly RC. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*. 2004; 306:1940–1943. [PubMed: 15528409]
- Gladstone G, Parker G. Measuring a behaviorally inhibited temperament style: development and initial validation of new self-report measures. *Psychiatry Res*. 2005; 135:133–143. [PubMed: 15922458]
- Holloway JL, Servatius R. Facilitated associative learning as a function of anxiety vulnerability in college aged participants. Society for Neuroscience Annual Meeting. 2010:599.593/JJJ598. [abstract].
- Kagan J, Reznick JS, Gibbons J. Inhibited and uninhibited types of children. *Child Dev*. 1989; 60:838–845. [PubMed: 2758880]

- Karamustafalioglu OK, Zohar J, Guveli M, Gal G, Bakim B, Fostick L, Karamustafalioglu N, Sasson Y. Natural course of posttraumatic stress disorder: a 20-month prospective study of Turkish earthquake survivors. *J Clin Psychiatry*. 2006; 67:882–889. [PubMed: 16848647]
- Marshall RD, Turner JB, Lewis-Fernandez R, Koenan K, Neria Y, Dohrenwend BP. Symptom patterns associated with chronic PTSD in male veterans: new findings from the National Vietnam Veterans Readjustment Study. *J Nerv Ment Dis*. 2006; 194:275–278. [PubMed: 16614549]
- Molet M, Leconte C, Rosas JM. Acquisition, extinction and temporal discrimination in human conditioned avoidance. *Behav Processes*. 2006; 73:199–208. [PubMed: 16806735]
- Myers, CE.; Moustafa, AA.; Sheynin, J.; VanMeenen, KM.; Gilbertson, MW.; Orr, SP.; Beck, KD.; Pang, KCH.; Servatius, RJ. Learning to obtain reward and avoid punishment is correlated with different PTSD symptom clusters in veterans: Empirical data and computational modeling. 2012a. (under review)
- Myers CE, Vanmeenen KM, McAuley JD, Beck KD, Pang KC, Servatius RJ. Behaviorally inhibited temperament is associated with severity of post-traumatic stress disorder symptoms and faster eyeblink conditioning in veterans. *Stress*. 2012b; 15:31–44. [PubMed: 21790343]
- Myers CE, Vanmeenen KM, Servatius RJ. Behavioral inhibition and PTSD symptoms in veterans. *Psychiatry Res*. 2012c; 196:271–276. [PubMed: 22397911]
- North CS, Nixon SJ, Shariat S, Mallonee S, McMillen JC, Spitznagel EL, Smith EM. Psychiatric disorders among survivors of the Oklahoma City bombing. *JAMA*. 1999; 282:755–762. [PubMed: 10463711]
- O'Donnell ML, Elliott P, Lau W, Creamer M. PTSD symptom trajectories: from early to chronic response. *Behav Res Ther*. 2007; 45:601–606. [PubMed: 16712783]
- Ricart TM, Jiao X, Pang KC, Beck KD, Servatius RJ. Classical and instrumental conditioning of eyeblink responses in Wistar-Kyoto and Sprague-Dawley rats. *Behav Brain Res*. 2011; 216:414–418. [PubMed: 20801161]
- Rosenbaum JF, Biederman J, Bolduc-Murphy EA, Faraone SV, Chaloff J, Hirshfeld DR, Kagan J. Behavioral inhibition in childhood: a risk factor for anxiety disorders. *Harv Rev Psychiatry*. 1993; 1:2–16. [PubMed: 9384823]
- Schlund MW, Cataldo MF. Amygdala involvement in human avoidance, escape and approach behavior. *Neuroimage*. 2010; 53:769–776. [PubMed: 20600966]
- Servatius RJ, Jiao X, Beck KD, Pang KC, Minor TR. Rapid avoidance acquisition in Wistar-Kyoto rats. *Behav Brain Res*. 2008; 192:191–197. [PubMed: 18501974]
- Simon JR, Howard JH, Howard DV. Adult age differences in learning from positive and negative probabilistic feedback. *Neuropsychology*. 2010; 24:534–541. [PubMed: 20604627]
- Wachter T, Lungu OV, Liu T, Willingham DT, Ashe J. Differential effect of reward and punishment on procedural learning. *J Neurosci*. 2009; 29:436–443. [PubMed: 19144843]

Figure 1.

The feedback-based probabilistic classification task (Bódi et al., 2009). (A) On each trial, the participant saw one of four stimuli and was asked whether this stimulus belonged to category A or B. Each stimulus belonged to one of the categories on 80% of trials and to the other category on 20% of trials. (B) For some stimuli, correct responses were rewarded with visual feedback and 25 points, whereas for other stimuli, incorrect responses were punished with visual feedback and the loss of 25 points.

Figure 2.

Mean optimal classification performance on reward (A) and punishment (B) trials in Experiment 1, in inhibited ($n=22$) vs. uninhibited ($n=52$) participants. Both BI groups showed learning over blocks (rmANOVA, $p<0.001$), but the inhibited participants outperformed their uninhibited counterparts ($p=0.035$). Dotted horizontal line represents chance performance (50%). Error bars indicate SEM.

Figure 3.

Individual bias scores (percent optimal on reward trials minus percent optimal on punishment trials), reflecting a participant's bias to learn from either positive or negative feedback on the first and the last block of Experiment 1. Each pair of adjacent bars represents an individual participant. Within each BI group, participants are ordered by the magnitude of their bias on the first block. Blank spaces represent values of zero. The dashed horizontal lines represent the thresholds for defining learning biases; participants whose absolute value of the bias score on the first block was less than 25% were defined as "balanced learners" (grey areas). Stable learners are those who had similar bias on the first and last blocks; three stable learners are labeled as examples: (1) is reward biased on both the first and last blocks; (2) is punishment biased on both blocks, and (3) is a balanced learner on both blocks. Approximately 70% and 55% of participants in both BI groups were balanced and stable, respectively (Pearson chi-square, both $p > 0.8$).

Figure 4.

Percent of trials during the first block of Experiment 1 in which participants showed “win-stay” and “lose-shift” behaviors, in inhibited ($n=22$) vs. uninhibited ($n=52$) participants. In general, participants showed a higher percentage of win-stay than lose-shift trials (rmANOVA, $p<0.001$). However, inhibited participants did not differ from their uninhibited counterparts. Asterisks indicate significant differences ($p<0.050$). Error bars indicate SEM.

Figure 5.

Mean response latency (msec) in Experiment 1, as a function of feedback type (reward vs. punishment) and participant's response (optimal vs. non-optimal) in inhibited ($n=21$) vs. uninhibited ($n=51$) participants. While responses on punishment trials were generally slower than responses on reward trials (rmANOVA, $p<0.001$), there was no significant difference in latency between optimal or non-optimal responses ($p=0.067$) or between inhibited participants and their uninhibited counterparts ($p=0.733$). Asterisks indicate significant differences ($p<0.050$). Error bars indicate SEM.

Figure 6.

Skipping behavior in Experiment 2. (A) Distribution of skipping frequency, in inhibited ($n=29$) and uninhibited ($n=50$) participants. Nearly half of the participants made no skipping responses. (B–C) Mean percent reward and punishment trials skipped per block, among participants who made at least one opt-out response. Participants showed increased numbers of opt-out responses across blocks (rmANOVA, $p<0.001$), with more skipping on punishment than reward trials ($p<0.001$) and more skipping performed by inhibited than uninhibited participants ($p=0.020$). Asterisks indicate significant differences ($p<0.050$). Error bars indicate SEM.

Figure 7.

Percent optimal classification responses in Experiment 2 (on trials where a classification response was made), among “non-skippers” ($n=35$) on reward (A) and punishment (B) trials, and among “skippers” ($n=44$) on reward (C) and punishment (D) trials. (E) Adaptive responding on punishment trials (including both optimal classification and skipping response) by skippers. The significant difference observed between BI groups in figures A–D (rmANOVA, $p=0.011$) is eliminated when skipping is scored as an adaptive response (E; $p=0.196$). Dotted horizontal line represents chance performance (50%). Asterisks indicate significant differences ($p < 0.050$). Error bars indicate SEM.

Figure 8.

Mean response latency (msec) in Experiment 2 for trials on which a classification response was made, as a function of feedback type (reward vs. punishment), participant's response (optimal vs. non-optimal) and skipping (skippers vs. non-skippers) in inhibited vs. uninhibited participants. Across response and feedback types, inhibited skippers ($n=16$) responded significantly more slowly than inhibited non-skippers ($n=12$; rmANOVA, $p=0.016$), while the uninhibited skippers ($n=27$) did not differ from uninhibited non-skippers ($n=22$; $p=0.587$). Error bars indicate SEM.