The Role of Outcome Expectancies in Reversal Learning Performance

People often have to use feedback to adjust to changing circumstances in their environment. This kind of learning takes places in social settings, school, gambling situations, etc. In the laboratory, cognitive flexibility based on feedback is often studied using reversal learning tasks. In reversal learning tasks, participants make choices over time and alter them based off of received feedback. One choice option is rewarded over another and, after this is learned, the contingencies are reversed. The previously "good" option is now "bad", and vice-versa. The variable of interest in this task is the ability to quickly begin making good choices post-reversal.

Past research in the area of reversal learning has been done predominantly with individuals with frontal lobe damage (e.g., Fellows, 2007). These individuals perform similar to healthy individuals on most measures of cognitive ability, but they exhibit deficits in social interactions as well as reversal learning tasks. Studies of reversal learning failure by people with frontal lobe damage indicate that frontal lobe damage interferes with the ability to create expectations for outcomes (Schoenbaum, Saddoris, & Stalnaker, 2007). This study tested the hypothesis that individual differences in reversal learning performance in people without brain injury are due to the differences in expectations people develop about choice outcomes.

Methods

Participants

Ninety-six participants were recruited from the Washington State University Department of Psychology online subject pool. Participants received compensation in the form of course credit.

Materials

This task was designed as a computerized gambling task programmed using E-Prime®. Reversal learning tasks have been used extensively to study cognitive flexibility in individuals with brain damage (Fellows, 2007), but these tasks were likely to produce ceiling effects in our neurologically intact college student population. Hence we designed a probabilistic reversal learning task in which participants used feedback in the forms of gains and losses to learn which of the two choices would result in the best outcomes. The task consisted of 96 choice trials, broken into four trial blocks (1,2,3a,3b,4). An example of one choice trial is shown in figure 1:

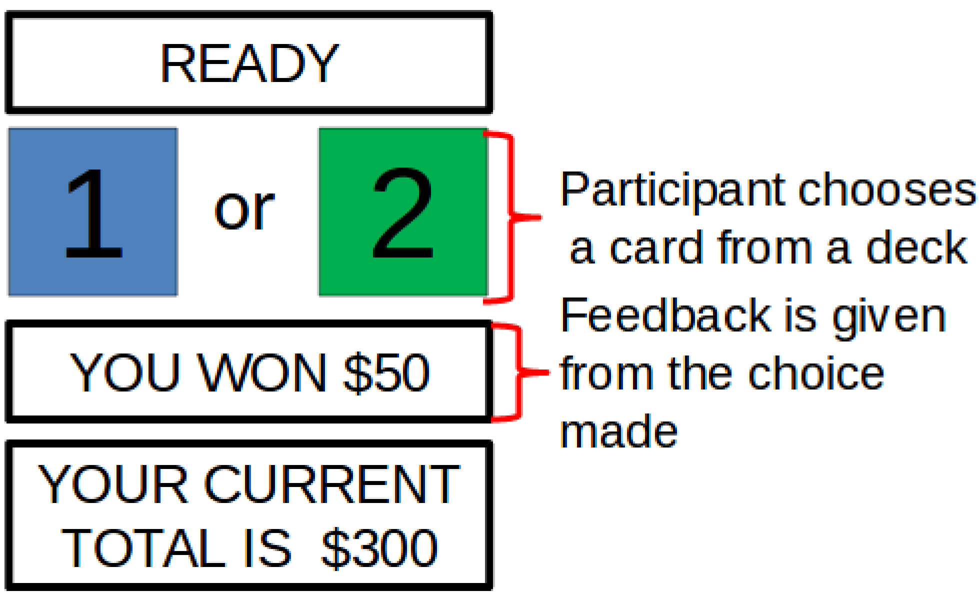

Figure 1

The participant saw a “Ready” screen, and then both decks appeared on screen. At that point, the participant made a choice between deck 1 and deck 2. After making this choice by clicking either the right or left mouse button, they saw feedback from their choice as well as their running total from all trials. Gains and losses were all hypothetical. Both decks produced gains and losses, but the “good deck” averaged a gain of $75 (SD =100) and the “bad deck” averaged a loss of $75 (SD= 100).

Design/Procedure

Participants were first given a consent form detailing what they would experience in the study. They were then attached to skin conductance response (SCR) units. This measure of affective arousal was recorded throughout the entire task. After that, a research assistant gave basic instructions regarding the task and asked them to give their full effort. The task was run on computers at individual stations, with no more than three participants running at one time.

Participants first saw instructions on the computer screen detailing task instructions. They ran through a set of practice trials to ensure that they understood how to perform the task. After the second block of 24 trials, participants were asked about their knowledge of deck outcomes. This was done in the form of the question, "Thinking about the results you received from DECK 1, what do you think would be a likely result if you chose another card from DECK 1?" The question was then repeated regarding deck 2. After 60 trials (halfway through the third block of 24 trials) the deck contingencies reversed (the formerly "good" deck became "bad" and vice versa). After completing the 96 choice trials, participants were given a debriefing form explaining the study.

Results

Choice Data

The goal of this experiment was to determine if there are individual differences in the ability to learn, and recover from, a reversal of contingencies. We split our group of 96 participants into three groups of 32; the low, medium, and high performers. Individuals were put in these groups based on the amount of times they chose the good deck over the bad deck in block 4 (recovery phase), which shows how well they were able to understand and adapt to the reversal in contingencies. Figure 2 shows the three groups' performance by block. A 3x2 analysis of variance (ANOVA) showed that all three groups showed a significant difference in performance before and after the reversal ($F(1,93)=113.08$, $p<.001$), indicating that the reversal did impact their choices. A 3x3 ANOVA showed that there was a significant main effect for group ($F(2,93)=9.82$, $p<.001$), with the high performing group performing significantly better pre-reversal than the low or medium groups, who did not differ from each other.

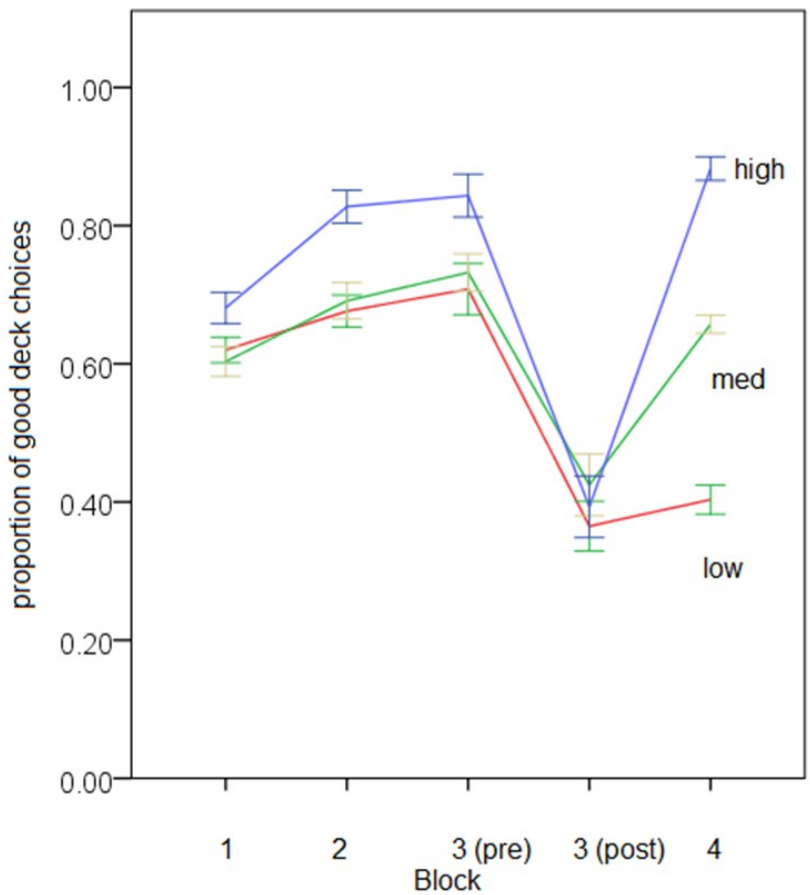

Figure 2

Knowledge Probes

Figure 3 shows the results obtained from analysis of the knowledge probe questions administered to participants after block 2. These knowledge probes allowed us to assess the expectations for each deck that participants developed over the course of the experiment. We hypothesized that individuals who had the most accurate expectations of outcomes would be the best performers, but in fact we find that those who performed the best in the reversal task were actually the individuals who significantly overestimated the difference between the good and bad decks, and most notably exaggerating the expected value of the good deck. This may lead to an "adaptive bias" which actually makes them more sensitive to the reversal since their exaggerated expectations are easily violated by a shift in card contingency. A 3x2 ANOVA revealed a significant main effect for deck ($F(1,91)=23.37$, $p<.001$), which indicates that knowledge probe accuracy differed depending on whether participants were rating the good or bad deck. It also revealed a marginally significant main effect for group ($F(1,91)=2.68$, $p=.07$), indicating that knowledge probe accuracy was related to task performance.

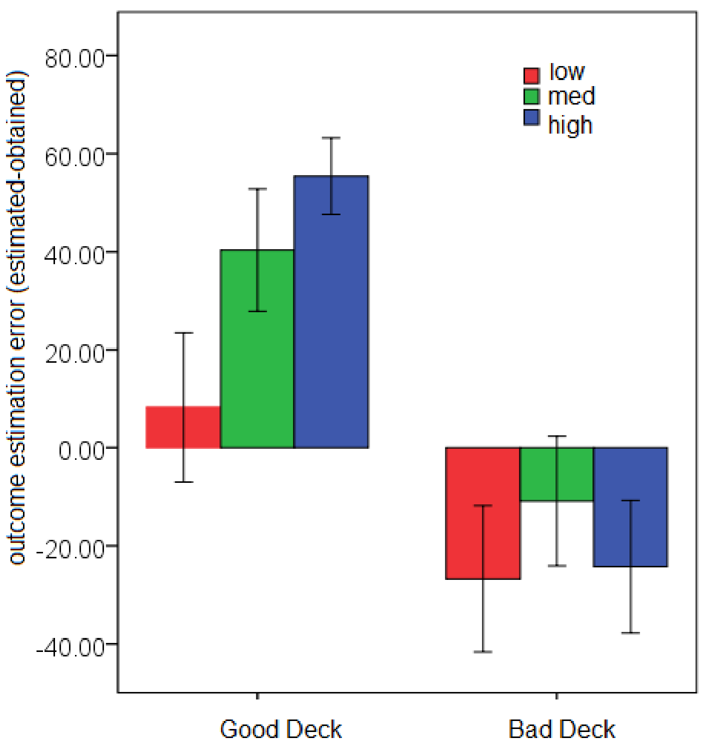

Figure 3

Skin Conductance Response

Our procedure used SCR in order to assess physiological arousal at different points of interest throughout the task. Figure 4 shows the SCR for the deck that starts out good but becomes bad post reversal. Figure 5 shows the SCR for the deck that starts out bad and becomes good post reversal. The two graphs are relatively similar pre-reversal, but differ once the reversal occurs. SCR drops in Figure 4 as all three groups learn that the deck is good and will average gains. SCR raises post-reversal as all three groups are either uncertain of their choice or displeased with their situation. An 8x3 ANOVA found a significant effect of block ($F(7,651)=6.82$, $p<.001$) which indicates that SCR levels differed across all groups depending on what block they were in. Since SCR is unidirectional, it cannot be determined what emotion causes the participants' spike in affect. Figure 5 shows all three groups experiencing a similar level of affect towards the bad deck pre-reversal. However, post-reversal SCR differs between the high performers and the other two groups in that the good performers' SCR levels decreased after learning that the deck had changed and now averages gains, while the medium and low performers' SCR levels remain high, reflecting their uncertainty. Another 8x3 ANOVA revealed a main effect of block ($F(7,651)=6.25$, $p<.001$) and a marginally significant main effect of group ($F(1,93)=2.79$, $p=.06$). The ANOVA also revealed a significant block by group interaction ($F(14,651)=2.42$, $p=.003$). Further analysis indicated that the groups differed in SCR levels in the last block ($F(1,93)=41.27$, $p<.001$).

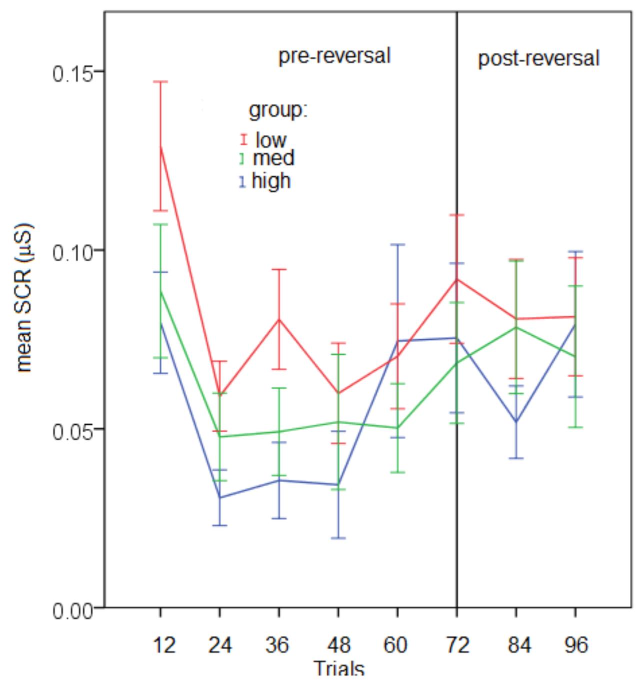

Figure 4

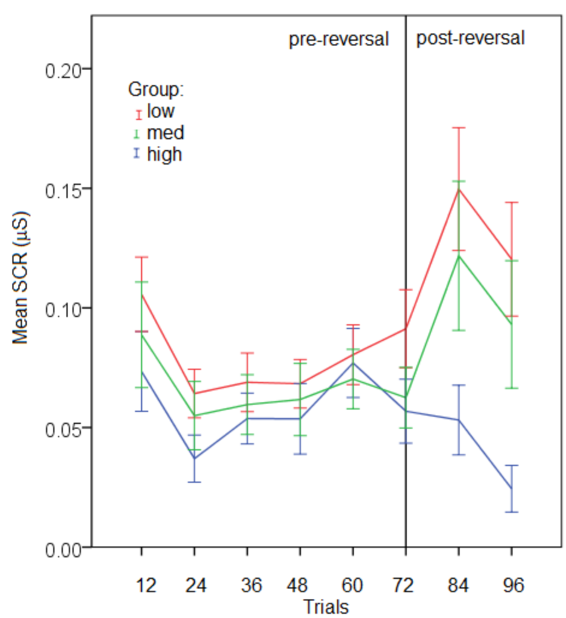

Figure 5

Discussion

This experiment set out to examine individual differences of expected outcomes in a neurologically intact group of participants in a probabilistic reversal learning task. Our hypothesis that difficulty in adapting to the contingency reversal would be related to accuracy in estimating deck outcomes was supported, but not in the way that we expected. The best performers did not have more accurate knowledge of deck values, but rather their biased impressions of the deck outcomes exaggerated the real differences, perhaps making it easier for them to detect changes in contingency. Individuals with frontal lobe lesions are unable to form accurate expectations of outcomes, which causes them to fail at reversal learning tasks. In contrast to this, the participants who performed the best at this reversal learning task had the most inaccurate knowledge of deck outcomes and the poor performers had the most accurate knowledge.

We sought to investigate individual differences in our reversal learning task as an analog of the reversal learning that occurs in everyday life, most of the time without our knowing it. When the ability to learn from feedback is impaired, many deficits occur in varying severity. However, our results show that one does not need to be incredibly accurate to successfully reversal learn, but rather understand the situational contingencies (to an exaggerated degree, in this experiment). The adaptive bias seen in the high performing group in this experiment may also be pervasive in other decision making tasks that require learning over trial periods. Future research should investigate the presence and utilization of this adaptive bias in other related tasks to examine whether or not exaggerated knowledge in the correct vein facilitates high performance.

References

Fellows, L. K. (2007). The Role of Orbitofrontal Cortex in Decision Making. *Annals of the New York Academy of Sciences*, *1121*, 421-430. doi:10.1196/annals.1401.023

Schoenbaum, G., Saddoris, M. P., & Stalnaker, T. A. (2007). Reconciling the Roles of Orbitofrontal Cortex in Reversal Learning and the Encoding of Outcome Expectancies. *Annals of the New York Academy of Sciences*, *1121*, 320-335. doi:10.1196/annals.1401.001