Anxiety Impedes Adaptive Social Learning Under Uncertainty

Amrita Lamba (amrita_lamba@brown.edu)
Cognitive, Linguistic, and Psychological Sciences, 190 Thayer Street
Brown University, Providence, RI 02912, United States

Michael J. Frank (michael_frank@brown.edu)
Cognitive, Linguistic, and Psychological Sciences, 190 Thayer Street
Brown University, Providence, RI 02912, United States

Oriel FeldmanHall (oriel_feldmanhall@brown.edu)
Cognitive, Linguistic, and Psychological Sciences, 190 Thayer Street
Brown University, Providence, RI 02912, United States

Abstract:

Very little is known about how individuals learn under uncertainty in social contexts. Given that social information is especially noisy and ambiguous, we propose that humans are particularly tuned to social uncertainty, which may be exacerbated in those who are uncertainty-sensitive. For example, anxious individuals generally report lower tolerance of uncertainty, which may be further heightened in social contexts. We employed a Bayesian-RL model in a dynamic Trust Game and matched slot machine task to probe reward learning dynamics across social and nonsocial domains. We find healthy subjects are particularly good at learning under negative social uncertainty (e.g., potential monetary losses imposed by others), as this buffers an individual from being exploited, which results in swiftly learning when to stop investing in an exploitative social partner. In contrast, anxious subjects showed equivalent sensitivity for monetary gains across both social and nonsocial contexts and thus sub-optimally overinvested in others. In addition, those with anxiety had difficulty in adjusting their learning rates as the task dynamics shifted. Our results suggest that humans are particularly tuned to negative social uncertainty, which likely facilitates adaptive social learning.

Keywords: anxiety; social learning; uncertainty sensitivity; Bayesian reinforcement learning; Trust Game

Introduction

The constantly evolving nature of real-life environmental contingencies results in action-outcome pairings that are often ambiguous and dynamic (Niv, Duff, & Dayan, 2005; Payzan-LeNestour & Bossaerts, 2011). In social situations, this is likely to occur with high frequency, as social information is often sparser and noisier compared to information acquired in nonsocial contexts (FeldmanHall & Shenhav, 2019). Yet, humans appear to be remarkably adept at learning these hidden states (e.g., another's intentions and motivations) in the social domain. If uncertainty is accentuated in the social domain, it would suggest a greater need to rely on generative models to make inferences about the environment. However, very little is known to date about the computational dynamics of social learning. Given the importance of social relationships to an individual's prospects and opportunities for wellbeing (e.g., more social partners and resources; Barkow, Cosmides, & Tooby, 1992), it is possible that humans are uniquely tuned to the subtle fluctuations in uncertainty encountered in social settings and are therefore quicker to adjust their learning from monetary losses imposed by others (e.g., negative prediction errors). If this is the case, then those that are particularly averse to uncertainty (i.e. those suffering from anxiety; Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015) may be disproportionately perturbed—and therefore slower to learn—from negative prediction errors encountered during social exchanges.

Methods

Experimental Design

To test our learning predictions across social and nonsocial contexts, we used a modified version of the well-vetted trust game (TG) paradigm with a matched nonsocial analog task (within-subjects, counterbalanced). On each trial of TG, subjects were endowed with $1.00 and paired with one of three players. They were asked how much of the $1.00 they wanted to invest using a slider bar set in $0.10 increments (Figure 1a). Endowments were quadrupled when sent to the other player, who could then decide how much (if any) money they wanted to return to the subject (actually a preprogrammed algorithm; Figure 1b). All aspects of the task were exactly matched between the social task (TG) and a nonsocial slot machine (SM) task, except for the social component. The task involved change points to elicit positive prediction errors (monetary gains) and negative prediction errors (monetary losses).

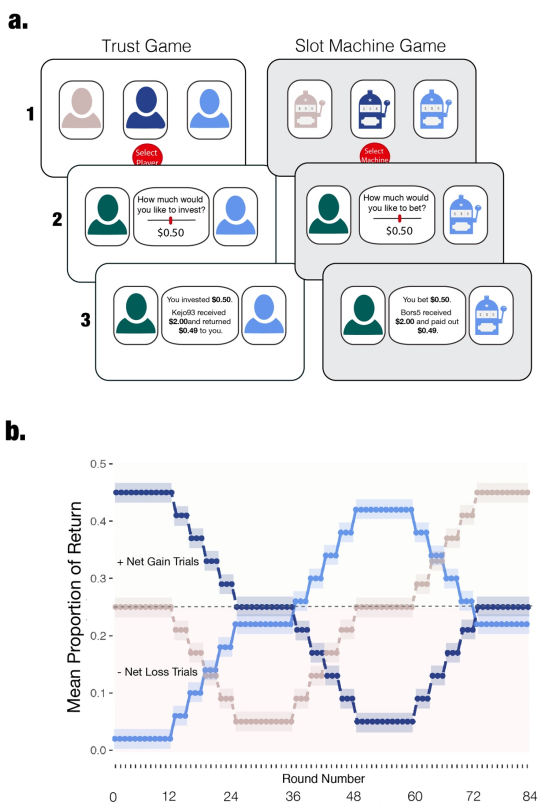

Figure 1. Depiction of task structure. **Fig. 1a** illustrates the trial structure of the Trust Game and matched slot machine game. The graph in **Fig. 1b** displays the preprogrammed algorithm underlying monetary returns for each online player in the Trust Game and each machine in the slot machine game. The y-axis of the graph corresponds to the proportion of the investment that was returned to the subject on a given trial. Subjects rotated through engaging with each online player once every three rounds and played a total of 28 rounds with each online player. The grey dotted line partitions trial types into net gain and net loss trials. The shaded areas around player returns correspond to a 4% uniform boundary, in which actual returns were randomly drawn from the corresponding return interval. Critically, all players were exactly monetarily equivalent in summed returns over the course of the game.

All data was collected on Amazon Mechanical Turk (N = 354, 53.1% female, age range: 18-45, *Mean* = 34.53 yrs). Based on the GAD-7 Scale, 97 subjects reported clinically-significant symptoms of Generalized-Anxiety Disorder.

Computational Model

In our computational analyses we compared three distinct models that differed in their degree of psychological relevance. Our primary model of interest was a 6-parameter Dynamic Bayesian RL (DB-RL) model that captures flexibility in learning (i.e. ability to adjust one's behavior in a non-stationary environment) and sensitivity to trial-to-trial changes in uncertainty (Daw, Niv, & Dayan, 2005; Franklin & Frank, 2015). The DB-RL model represents subjects' current beliefs about the payoff-maximizing strategy implemented on each trial. Subject beliefs are summarized by a beta distribution, which includes both the mean and the uncertainty about beliefs which are separately updated with each positive or negative outcome.

Because our task involves change-points (e.g., the trustworthy-start player depicted as the dark blue line in Figure 1b, who slowly shifts from always reciprocating to always defecting as the task progresses), subjects should down-weight past outcomes over the course of the task, which in turn increases the overall uncertainty in the posterior distribution and allows more recent trial-level feedback to be more informative than past outcomes. A decay parameter γ was fit for each subject to estimate their degree of learning flexibility (low $\gamma =$ more decay), separately for positive and negative outcomes (γ_{pos} and γ_{neg}, respectively). We further allowed the decay to increase as uncertainty increases about the other player's strategy (e.g., when the strategy changes: quantified as entropy *H* in the posterior distribution; Franklin & Frank, 2015).

Formally, we modeled γ_0 and γ_1 for positive and negative outcomes as separate free parameters to account for valence-dependent asymmetries, using a logit transform to maintain a 0-1 range.

$$logit(\gamma_{pos}) = \gamma_{0_{pos}} + \gamma_{1_{pos}} \cdot \Delta H$$
$$logit(\gamma_{neg}) = \gamma_{0_{neg}} + \gamma_{1_{neg}} \cdot \Delta H$$

We predicted that our DB-RL model would best capture learning in healthy subjects but that subjects with anxiety would not efficiently use the change points (where uncertainty is greatest) to appropriately adjust their learning. To account for these potential differences, we compared the fit of DB-RL with simplified Bayesian RL that was equivalent in all respects, except that it was modeled without γ_1 and ΔH parameters. Additionally, we compared the fit of both models to standard RL.

Results

Behavioral Results

Anxiety subjects overinvested in TG players compared to healthy controls, particularly when reward dynamics were downward-trending. Figure 2a shows mean investments broken out across net gain (positive valence) and net loss (negative valence) trials by task block (4 trial bins). Effectively, anxious subjects gave significantly more money during negative valence blocks for both the neutral start player (RMANOVA for neutral start player investments, healthy vs. anxious × valence: *F(2,704)* = 4.30, *p* = 0.024) and the trustworthy start player (RMANOVA for trustworthy start player investments, healthy vs. anxious × valence: *F(2,704)* = 6.35, *p* < 0.004), and therefore were slower to learn when to stop investing in an exploitative social partner. This suggests that anxious individuals are slower to learn the statistics of negative outcomes relative to healthy controls.

Computational Results

Using Bayesian model selection, we compared the relative fit of DB-RL, B-RL, and RL, finding DB-RL to be the winning model for healthy subjects (pxp > 0.99). However, when models were individually compared using pairwise comparisons for the anxiety group, there was no clear model-fit difference between DB-RL and simplified B-RL (SM pxp = 0.54, TG pxp = 0.56), indicating that anxious subjects were fit equally well by both DB-RL and B-RL models. Thus, the DB-RL model did not do a better job of capturing behavior in anxious subjects when compared to simplified Bayesian-RL. This suggests that healthy subjects were in fact, using fluctuations in environmental uncertainty to adaptively guide learning, whereas anxious subjects exhibited these effects to a lesser degree.

Accordingly, to further examine whether individuals with anxiety were able to learn about the drifting reward statistics of the environment in the same way as healthy controls, we directly compared the decay rate across groups. We assessed whether these groups would exhibit differences in the decay of past positive versus negative experiences ($\gamma_{0_{pos}} - \gamma_{0_{neg}}$). Because decay allows for flexibility in updating one's learning rule, overweighting past rewards relative to losses ($\gamma_{0_{pos}} > \gamma_{0_{neg}}$) biases one towards consistently overinvesting. Conversely, perseverating on past losses relative to rewards produces a bias towards under-investing ($\gamma_{0_{neg}} > \gamma_{0_{pos}}$).

While both healthy and anxious subjects showed a general bias towards weighting rewards more heavily than losses in SM game compared to TG ($\gamma_{0_{pos}} > \gamma_{0_{neg}}$; *F(1,352)* = 12.94, *p* < .001; Figure 2b) resulting in overinvesting, only healthy subjects selectively adjusted their learning in the TG by demonstrating a greater likelihood of weighting losses more heavily than rewards ($\gamma_{0_{neg}} > \gamma_{0_{pos}}$).

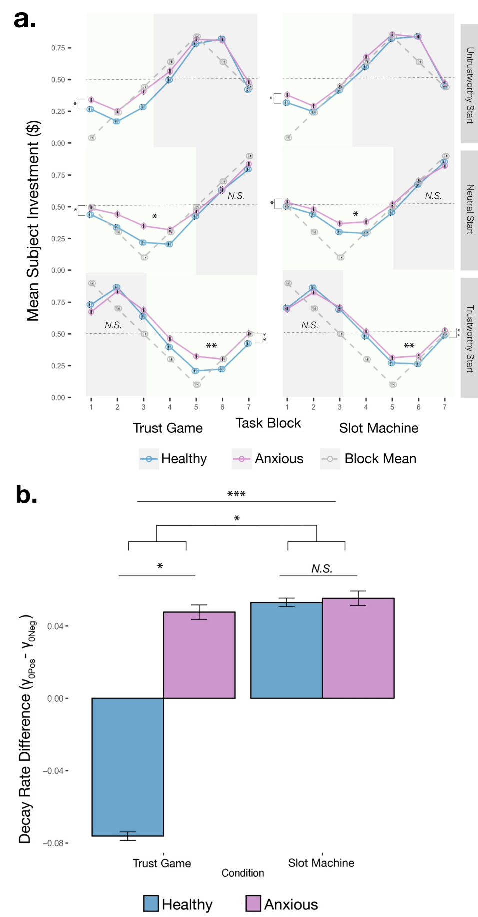

Figure 2. Behavioral and computational modeling results. **Fig. 2a** shows learning curves for healthy and anxious subjects. The y-axis of the graph indicates mean investments in TG and SM per block (4 trial bins, indicated on the x-axis). The angled grey line in each panel corresponds to the rescaled proportion of return that the agent was set to per block (rescaled and transposed from Figure 1b). The pale green and light grey panels correspond to negative and positive valence blocks, respectively. **Fig. 2b** shows mean and standard error estimates of decay rate difference ($\gamma_{0_{pos}} - \gamma_{0_{neg}}$) across healthy vs. anxious groups.

The results reveal that overall, context selectively moderates the impact of valence on reward-learning in healthy subjects (Mixed design ANOVA, healthy vs. anxious × condition: $F(1, 352) = 4.14$, $p = 0.042$). Post-hoc pairwise comparisons further revealed that there was a significant difference in decay rate between healthy and anxious subjects in the TG ($t(352) = -2.57$, $p = 0.011$), whereas no differences emerged in the slot machine game ($t(352) = -0.045$, $p = 0.96$).

Discussion

Although learning and uncertainty are tightly coupled, little is known about how this relationship unfolds in the social domain. Understanding the dynamics of learning under social uncertainty is especially pertinent given that social interactions are particularly noisy and ambiguous. According to past work, social and nonsocial reward learning are governed by largely overlapping neural circuitry, suggesting that social and nonsocial cognition arise from parallel computations (Behrens, Hunt, Woolrich, & Rushworth, 2008).

In the work presented here, we find that this may not strictly be the case. By comparing social and nonsocial learning under uncertainty we found that healthy individuals exhibit asymmetrical learning profiles across contexts: in the slot machine task, healthy subjects showed a distinct pattern of overweighting rewards relative to losses which resulted in consistently overinvesting in trials that had negative prediction errors. Conversely, in the social domain, healthy subjects switched their learning strategy, such that they were more likely to weight losses more heavily than rewards, suggesting that healthy subjects were highly sensitive to exploitative behavior and were quicker to learn to stop investing.

In addition, these data provide the first evidence that we are aware of showing that people can selectively adjust learning across contexts to avoid social exploitation. Although previous work in the nonsocial domain illustrates that individuals with trait anxiety have difficulty learning the statistics of volatile reward environments (i.e., learning flexibility is impaired in anxious individuals; Browning et al., 2015; Delgado, Frank, & Phelps, 2005), we show that these effects are exacerbated in the social domain. Together, the work presented here offers novel evidence that learning under uncertainty uniquely unfolds across social and nonsocial domains, providing a candidate mechanism for how learning under uncertainty varies as a function of the environment and uncertainty sensitivity.

Acknowledgments

The current study was funded by the NIH's Centers of Biomedical Research Excellence (COBRE) program.

References

Barkow, J. H., Cosmides, L., & Tooby, J. (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. New York, NY, US: Oxford University Press.

Behrens, T. E., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. (2008). Associative learning of social value. *Nature, 456*(7219), 245-249. doi:10.1038/nature07538

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. *Nat Neurosci, 18*(4), 590-596. doi:10.1038/nn.3961

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci, 8*(12), 1704-1711. doi:10.1038/nn1560

Delgado, M. R., Frank, R. H., & Phelps, E. A. (2005). Perceptions of moral character modulate the neural systems of reward during the trust game. *Nat Neurosci, 8*(11), 1611-1618. doi:10.1038/nn1575

FeldmanHall, O., & Shenhav, A. (2019). Resolving uncertainty in a social world. *Nature Human Behaviour*. doi:10.1038/s41562-019-0590-x

Franklin, N. T., & Frank, M. J. (2015). A cholinergic feedback circuit to regulate striatal population uncertainty and optimize reinforcement learning. *Elife, 4*. doi:10.7554/eLife.12029

Niv, Y., Duff, M. O., & Dayan, P. (2005). Dopamine, uncertainty and TD learning. *Behav Brain Funct, 1*, 6. doi:10.1186/1744-9081-1-6

Payzan-LeNestour, E., & Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Comput Biol, 7*(1), e1001048. doi:10.1371/journal.pcbi.1001048