

Anxiety Impedes Adaptive Social Learning Under Uncertainty



Amrita Lamba¹, Michael J. Frank^{1,2}, and Oriel FeldmanHall^{1,2} 

¹Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, and

²Carney Institute for Brain Science, Brown University

Psychological Science

1–12

© The Author(s) 2020

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/0956797620910993

www.psychologicalscience.org/PS



Abstract

Very little is known about how individuals learn under uncertainty when other people are involved. We propose that humans are particularly tuned to social uncertainty, which is especially noisy and ambiguous. Individuals exhibiting less tolerance for uncertainty, such as those with anxiety, may have greater difficulty learning in uncertain social contexts and therefore provide an ideal test population to probe learning dynamics under uncertainty. Using a dynamic trust game and a matched nonsocial task, we found that healthy subjects ($n = 257$) were particularly good at learning under negative social uncertainty, swiftly figuring out when to stop investing in an exploitative social partner. In contrast, subjects with anxiety ($n = 97$) overinvested in exploitative partners. Computational modeling attributed this pattern to a selective reduction in learning from negative social events and a failure to enhance learning as uncertainty rises—two mechanisms that likely facilitate adaptive social choice.

Keywords

anxiety, social learning, uncertainty sensitivity, Bayesian reinforcement learning, trust game, open data

Received 6/22/19; Revision accepted 1/4/20

Real-world environmental dynamics are noisy and evolving. Thus, deciding on the best action to take in a given moment requires us to appropriately weight potential rewards and losses in fundamentally uncertain settings (Dayan & Niv, 2008; Mathys, Daunizeau, Friston, & Stephan, 2011; Payzan-LeNestour & Bossaerts, 2011; Platt & Huettel, 2008; Rangel, Camerer, & Montague, 2008). Because humans have hidden and dynamic intentions that are not always observable to others, social information in particular is exceedingly uncertain (FeldmanHall & Shenhav, 2019). Despite this, humans appear to be remarkably proficient at adapting their behavior in response to the reward statistics of the current social landscape. Continually making adaptive social choices (e.g., those that facilitate functional and appropriate behavior in everyday life; Wehmeyer, 2013) should therefore require individuals to incrementally adjust their behavior to maximize social rewards (e.g., mutual cooperation toward a common goal) while avoiding exploitation from others. The mechanisms that support adaptive social learning under uncertainty—such as

figuring out whether people can be trusted and, if so, whether they should continue to be trusted if contexts change—are still poorly understood. To comprehend what facilitates adaptive social choice, we investigated how humans learn the reward statistics of social exchanges that have multidimensional and evolving hidden states.

Insights acquired from traditional theories in the non-social domain suggest that flexible, on-line learning is largely facilitated through the coupling of learning and uncertainty-perception systems (Franklin & Frank, 2015; Mathys et al., 2011; McGuire, Nassar, Gold, & Kable, 2014; Nassar, Wilson, Heasly, & Gold, 2010; Niv, Duff, & Dayan, 2005; O'Reilly, 2013; Rushworth & Behrens, 2008; Yu & Dayan, 2005). When humans perceive an increase in environmental uncertainty (e.g., through the addition of volatility in reward contingencies), new information

Corresponding Author:

Oriel FeldmanHall, Brown University, 190 Thayer St., Box 1821, Providence, RI 02912

E-mail: oriel.feldmanhall@brown.edu

should be weighted more heavily than past information, and therefore learning should be adjusted upward (Aylward et al., 2019; Behrens, Woolrich, Walton, & Rushworth, 2007; Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015; Franklin & Frank, 2015). People also asymmetrically learn from positive and negative outcomes, depending on the context (Collins & Frank, 2014; Gershman, 2016; Niv, Edlund, Dayan, & Doherty, 2012). The lion's share of these formal learning models have been focused on learning in nonsocial situations (e.g., gambling and foraging; Auer, Cesa-Bianchi, & Fischer, 2002; Montague, Dayan, Person, & Sejnowski, 1995; Sutton & Barto, 1998; Weinstein & Littman, 2012); thus, how learning unfolds during social interactions (Behrens, Hunt, Woolrich, & Rushworth, 2008; Diaconescu et al., 2014; Diaconescu et al., 2017), especially under uncertainty, remains an open question.

Given the importance of social relationships to people's prospects and opportunities for success and well-being, it is possible that humans are uniquely attuned to the subtle fluctuations in uncertainty encountered in social settings, compared with nonsocial settings (FeldmanHall & Shenhav, 2019). The implications of this are threefold. First, if uncertainty is accentuated in the social domain, this would suggest a greater need to rely on internal generative models (e.g., probabilistic learning) to make approximate meta-inferences about the structure of the environment (e.g., the variance of social rewards), which should augment learning in complex social exchanges. Second, if humans are more perceptive of uncertainty in social settings, then they should be able to reduce aversive experiences of uncertainty through learning the statistics of the environment, thus exhibiting faster learning. Third, evidence of asymmetric reward learning in the nonsocial domain (Collins & Frank, 2014) should be even greater in social situations, because detecting exploitative behavior is a critical component of disengaging from maladaptive social exchanges.

To characterize how social learning under uncertainty unfolds in the social domain, it would be particularly useful to examine populations that are distinctly sensitive to uncertainty. For example, individual differences in aversion to uncertainty map well onto certain pathological disorders, such as anxiety (Boelen & Reijntjes, 2009; Carleton et al., 2012). Individuals with high levels of trait anxiety show observable difficulty learning the causal statistics of volatile reward environments in the nonsocial domain, which suggests that they have impairments in appropriately adjusting their learning in highly uncertain settings (Aylward et al., 2019; Browning et al., 2015). Although little is known about the scope of these learning impairments in social environments where uncertainty may be even greater, individuals with anxiety have reported difficulty maintaining healthy social relationships (Barrera & Norton,

2009; Eng, Heimberg, Hart, Schneier, & Liebowitz, 2001; Rubin & Burgess, 2001). This suggests that examining learning dynamics in populations that vary in uncertainty sensitivity may provide a more holistic picture of the relevant cognitive systems that support social learning processes.

In the current study, we merged empirical and computational approaches to test the joint impact of context (social vs. nonsocial) and uncertainty sensitivity (healthy individuals vs. individuals with anxiety) on adaptive social choice. To parametrize rewards and losses in the social domain, we used an incentive-compatible version of the well-vetted trust-game paradigm, which we optimized for examining evolving reward-learning dynamics with a Bayesian reinforcement-learning (BRL) model. Critically, because the reward structure of the task gradually fluctuated over the course of the experiment, subjects were required to continually adjust their learning rate as the task progressed. To specifically compare learning across social and nonsocial domains, we used a matched slot-machine game in which all aspects of the trust game were preserved. Given our hypothesis that uncertainty may be exacerbated in social exchanges, we predicted that (a) people would be quicker to learn the reward contingencies of the trust game compared with the slot-machine game; (b) asymmetrical learning profiles would emerge in which losses will be overweighed relative to rewards, and these effects would be amplified in social contexts; and (c) individuals with trait anxiety (i.e., those who have trouble with processing uncertainty) will exhibit dampened learning effects, especially for negative social information.

Method

Subjects

We conducted a power analysis using G*Power (Version 3; Faul, Erdfelder, Lang, & Buchner, 2007), which revealed that a sample size of 129 would be necessary to detect a medium effect size (f^2) of 0.15 ($\alpha = .95$, power = .95). However, because we wanted to ensure we had sufficient power to detect individual differences in anxiety, we estimated that a final sample of approximately 400 would be adequate for ensuring we had a sizeable group of individuals with clinically significant anxiety symptoms. In addition, our aim was to collect data from a diverse and representative clinical population. Recent studies suggest that online samples are not only more representative of the population at large (Berinsky, Huber, & Lenz, 2012; Horton, Rand, & Zeckhauser, 2011) but also are more anonymous, which increases the likelihood that subjects feel comfortable disclosing anxiety symptoms (Gillan & Daw, 2016; Shapiro, Chandler, & Mueller, 2013).

Accordingly, we recruited a sample of 412 subjects from Amazon Mechanical Turk (MTurk; mean age = 34.61 years, $SD = 9.41$; 53.1% female). Our sample was restricted to the United States to prevent systematic error due to English-comprehension skills and cross-cultural differences in economic decision making (Yamagishi & Yamagishi, 1994). Using our model Akaike information criterion (AIC) exclusion criteria (detailed in the Supplemental Material available online), we excluded 58 subjects from all analyses (a 14% attrition rate) who either demonstrated poorer learning than chance ($n = 20$) or simply clicked through the entire experiment by indicating the same response on all trials on one or both tasks ($n = 38$). Our final sample thus consisted of 354 subjects. Within this final sample, 97 subjects (~27.4%) reported clinically significant symptoms of generalized anxiety disorder (see the Supplemental Material for scoring details).

Experimental design

All subjects completed both a repeated trust game and a slot-machine paradigm (presentation was counterbalanced across subjects; Fig. 1). For the trust game, subjects were told that they would be paired with three other online players on MTurk and would be making real decisions that would be realized at the end of the session (one trial was randomly selected for payout). After passing a comprehension check of all task instructions, subjects experienced a 1-min delay while the server supposedly paired them with other online players through MTurk. In reality, the online players followed a preprogrammed algorithm to maximize either stable or noisy behavior that allowed us to probe subjects' learning dynamics.

Subjects were not given any information about the other players beforehand and could see only their generated MTurk usernames. At the start of every trial, subjects clicked the "Select Player" button, which paired them with one of the players for the current trial (Fig. 1a). To standardize learning opportunities across all three players, we paired subjects with each player once every three trials in a pseudorandom order. On each trial, subjects were endowed with \$1.00 and asked how much of their endowment they wanted to invest. They indicated their choice using a slider bar that moved in \$0.10 increments (a minimum investment of \$0.10 was required to ensure learning on all trials). Subjects were told that all investments would be quadrupled when sent to the other player. Once the money was invested, the other player could decide how much money (if any) to return to the subject. After subjects indicated their desired investment, there was a 3- to 5-s jittered delay while the other player made a decision before seeing what the other player returned monetarily.

Unbeknownst to the subjects, the players responded according to preprogrammed algorithms (see Fig. 1b, which displays the proportion of the quadrupled investment returned on each trial). For example, if the subject invested \$0.50 on a particular trial and the proportion of return was set to .45, then the subject would receive \$0.90 back: $\$0.50 \times 4$ (investment quadrupled) \times .45. Subjects played a total of 28 rounds with each online player over the course of the session.

We chose to leverage slower, drifting change points in our design to create a more realistic social experience. On average, the trust-game paradigm took subjects approximately 25 min to complete. Many important social interactions, such as a job interview, occur on a similar timescale and often involve slow drifts in attitude and behavior in the interviewer for reasons unbeknownst to the interviewee. We therefore constructed the task to simulate the dynamics of real-world social exchanges. To leverage these subtle changes in reward dynamics, we designed the task so the amount of money a particular player returned slowly changed over the course of the experiment. For example, one player (denoted by the dark blue line in Fig. 1b) started out as a trustworthy reciprocator by always returning roughly half of the money. However, over time, this player gradually began to return less and less. During these trials, the reward contingencies drifted from approximately half of the investment being returned to the subjects to the other players keeping all of the money for themselves. The drifting reward rate (money returned) required subjects to continually track the relevant reward statistics of each player and learn when to change formerly optimal decision strategies. The three player types incrementally altered their behavior at particular change points in the task, and these points marked transitions between stable and drifting trial blocks. In addition, we added a probabilistic 4% uniform boundary to the proportion of money returned to add a margin of noise around the generated feedback on each trial. This added uncertainty was intended to prevent subjects from suspecting computer-generated responses.

A critical feature of our task was that the summed returns of all three player types were constructed to be exactly monetarily equivalent over the course of the game, assuming equal investments across players. In other words, all players had exactly the same overall reward rate and differed only in their starting points and temporal trajectories. All player return rates were also exactly matched to the slot machines over the course of the game. Notably, when the proportion of return for a player was set above .25 (indicated by the dashed gray line in Fig. 1b), subjects always maximized their earnings by investing the entire \$1.00. Conversely, when the proportion of return was set below .25, subjects

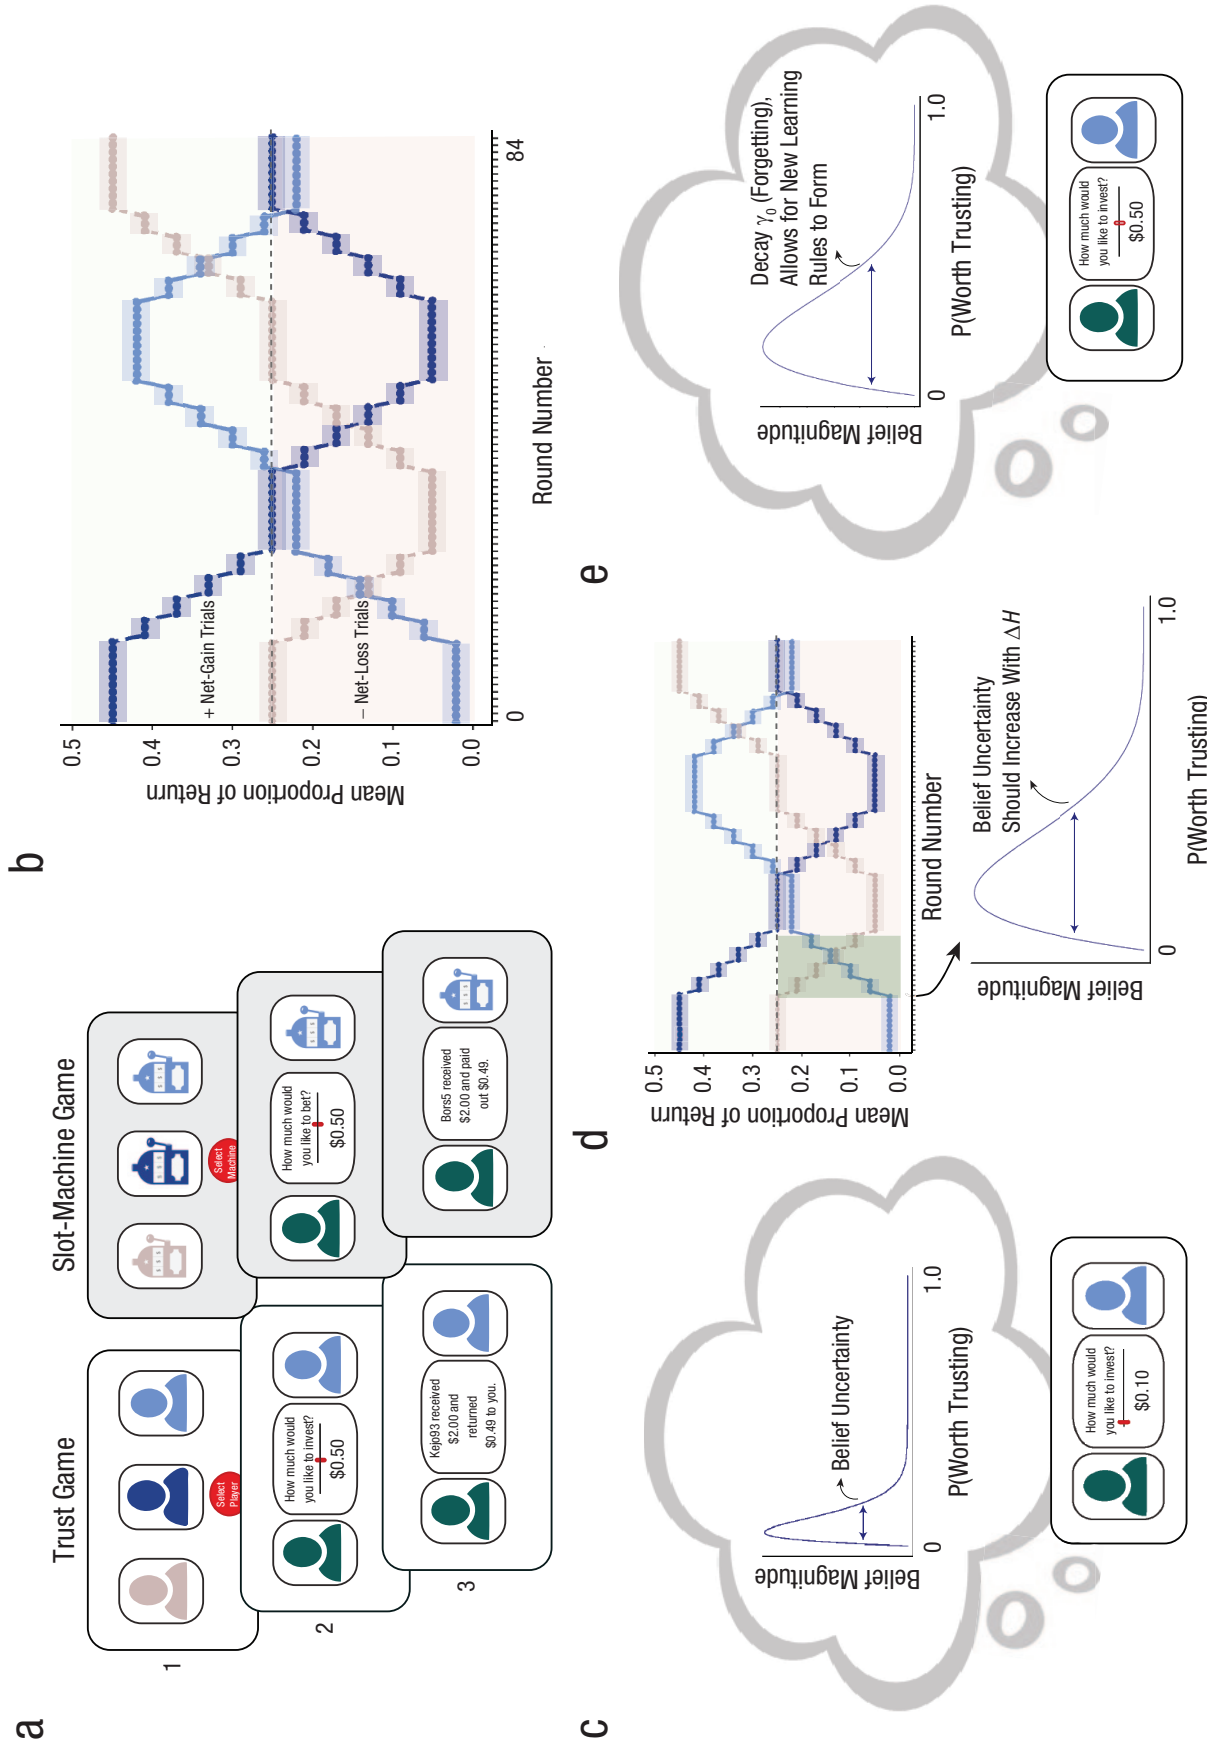


Fig. 1. Experimental design and method. The trial structures of the trust game and the matched slot-machine game are shown in (a). Participants began each trial by selecting one of the three players or three machines, respectively. They were then asked how much they wanted to invest or to bet, and afterward they were told how much the other player or machine had returned to them. The preprogrammed algorithm underlying monetary returns for each of the three players is shown in (b). The dashed horizontal 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. Assuming equal investments, the summed returns of all players were exactly monetarily equivalent over the course of the game. In our computational model, a beta distribution (c) represents the subject's current beliefs (y-axis) of the probability that it is worth investing the full \$1.00 in the other player (x-axis). Here, the subject exhibits a strong belief toward not investing (as indicated by the short width of the double-headed arrow) and therefore consistently opts to give the minimum of \$0.10. When there are change points in the task (as indicated by the green box), there is more uncertainty about the other player's strategy (d), and thus belief uncertainty should increase with entropy (ΔH). The role of negative decay ($\gamma_{0, neg}$) in increasing uncertainty of one's beliefs by downweighting past negative outcomes (e) allows new evidence to be more informative when one updates the posterior distribution. In this case, ($\gamma_{0, neg}$) allows the subject to incrementally adjust his or her belief that it is worth investing in the other player.

always incurred a net loss and hence minimized their losses by investing the required minimum of \$0.10.

Computational model

In our computational analyses, we compared three distinct models that differed in their psychological relevance. Our primary model of interest was a six-parameter dynamic BRL (DBRL) 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 (Franklin & Frank, 2015). The DBRL model represented subjects' current beliefs about the best strategy to implement on each trial to maximize payoffs. In Figure 1c, subjects' beliefs are summarized in a beta distribution, which includes both the mean belief and the uncertainty about this belief, updated separately for positive or negative outcomes (i.e., whether or not a player returned a sufficient portion of the subject's investment).

Because our task involved change points (i.e., when players begin to change their behavior), subjects should downweight past outcomes over the course of the task (Fig. 1d). For example, if previously untrustworthy players shift their behavior toward increasing reciprocity (i.e., change from untrustworthy to trustworthy), the model accounts for this change by decaying the influence of past outcome history, which in turn increases the overall uncertainty in the posterior distribution and allows more recent feedback to be more informative than outcomes in the distant past (Fig. 1e). A decay parameter γ was fitted for each subject to estimate the degree of learning flexibility (low γ = more decay of past outcomes), separately for positive and negative outcomes (γ_{pos} and γ_{neg} , respectively).

Further, the decay should increase when there is evidence that the other player might be changing strategy. To accommodate for this dynamic, we included a parameter that increases decay when the uncertainty—quantified as entropy H in the posterior distribution—about the other player's strategy increases (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 (i.e., trust being reciprocated or not), using a logit transform to maintain a range of 0 to 1 (see the Supplemental Material for additional equation details):

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

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

We predicted that the DBRL model would best capture learning in healthy subjects but would be a poorer fit to subjects with anxiety because those with anxiety would not efficiently use the uncertainty embedded in the task to appropriately adjust their learning. To account for these potential differences, we compared the fit of the DBRL model with a simplified BRL model that was equivalent in all respects, except that it did not include the decay rate (γ_1) and entropy (ΔH) parameters. Mechanistically, the simplified BRL model does not dynamically alter its uncertainty and learning with change points. Additionally, we compared the fit of both models with the fit of a standard reinforcement-learning (RL) model (see Table S2 in the Supplemental Material for all parameter details). While computationally more simple, Q-learning is generally insensitive to changes in task-level uncertainty and is also limited in its ability to quantitatively index learning flexibility (Daw, 2014). Therefore, we expected the standard Q-learning model to be a poorer fit across all subjects compared with both the dynamic and general BRL models.

For all models, subject choices on a given trial were modeled using inverse-temperature and bias parameters (Table S1). The inverse-temperature parameter computed explore/exploit trade-offs in relation to learned decision rules, whereas the bias parameter captured individual-specific choice benchmarks for investing.

Results

Behavioral results for healthy controls

Because we were interested in examining differences between social and nonsocial contexts, we first examined mean investments across the trust game and slot-machine game in healthy controls. Overall, subjects invested more money in the slot-machine game compared with the trust game, $t(256) = -6.42$, $p < .001$, $d = 0.13$, 95% confidence interval (CI) = $[-0.079, -0.042]$ (see Fig. 2a). Moreover, despite the fact that all players in the trust game (and all machines in the slot-machine game) were matched to return equivalent amounts of money across the task, healthy subjects asymmetrically invested in the players—repeated measures analysis of variance (ANOVA) on trust-game player type: $F(2, 512) = 23.52$ (Greenhouse-Geisser corrected), $p < .001$, $d = 0.61$. Bonferroni-corrected pairwise comparisons indicated that subjects invested more money in the trustworthy-start player, compared with both the untrustworthy-start player (mean difference = \$0.032, $p < .001$, 95% CI = $[\$0.013, \$0.051]$) and the neutral-start player (mean difference = \$0.048, $p < .001$, 95% CI = $[\$0.031, \$0.064]$). This suggests that the initially trustworthy player

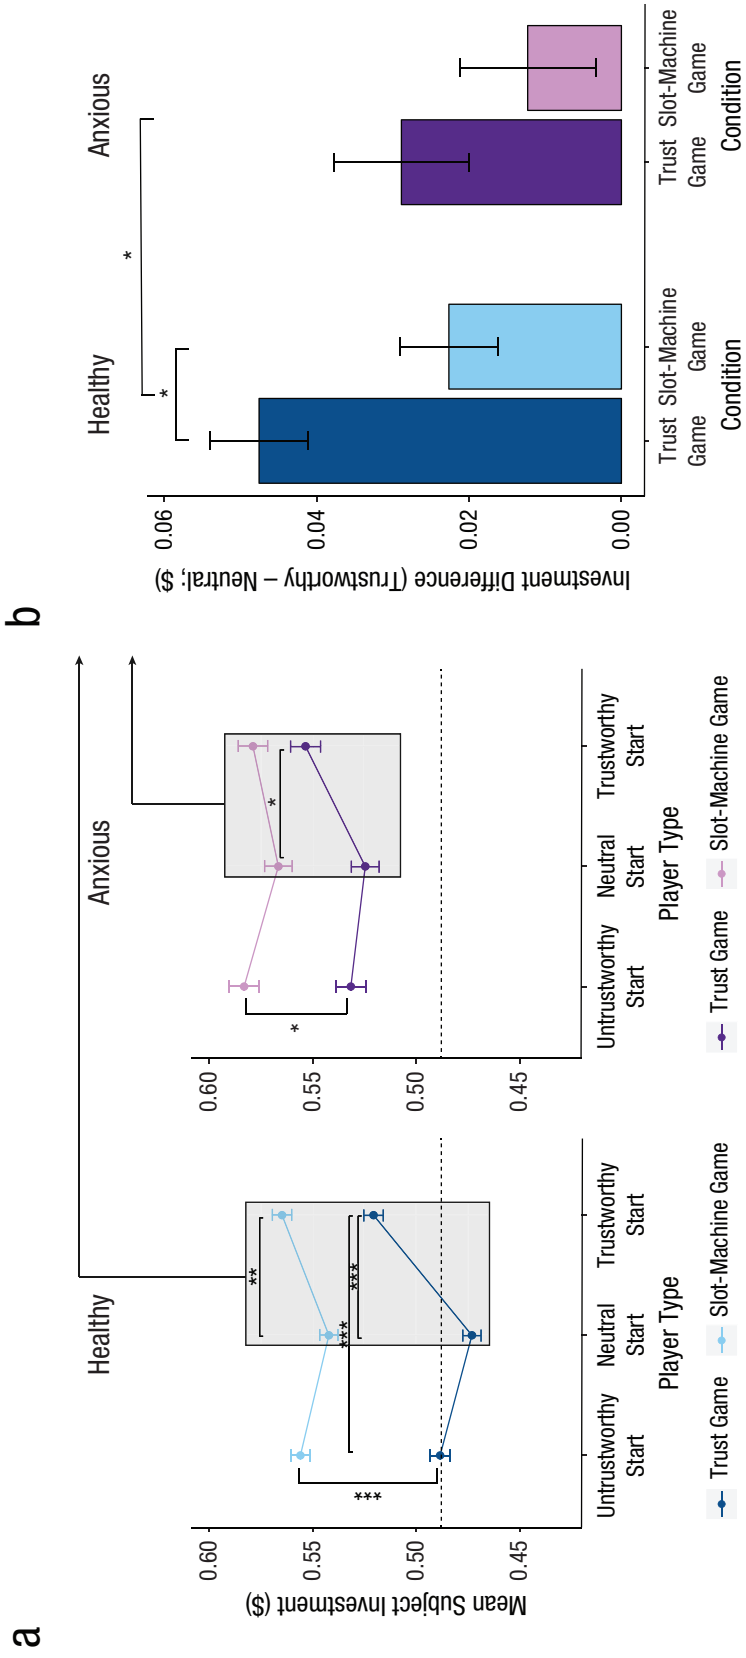


Fig. 2. Behavioral results. Subject investments across player types in the trust game and slot-machine game (a) are illustrated separately for healthy subjects and subjects with anxiety. The dotted line indicates optimal mean investment, given full knowledge of the task structure. Results highlighted by gray boxes are further analyzed in (b), which shows the difference in investments between the trustworthy-start and neutral-start players in each task, separately for healthy subjects and subjects with anxiety. In (a) significant differences between player types and between investments are indicated by asterisks (* $p < .05$, ** $p < .01$, *** $p < .001$); in (b), significant differences between conditions are indicated by asterisks (* $p < .05$). Error bars correspond to standard errors of the mean.

managed to confer a positive first impression that biased all subsequent decisions over the course of the task.

When comparing these asymmetrical investments across social and nonsocial domains, we observed that although a similar pattern emerged when playing with the slot machines—repeated measures ANOVA on slot-machine type, $F(2, 512) = 5.04, p = .007, d = 0.28$ —subjects were significantly more biased by first impressions when the task was social—repeated measures ANOVA on Task \times Player Type, $F(2, 512) = 4.42$ (Greenhouse-Geisser corrected), $p = .014, d = 0.11$. We further confirmed this by computing subject-level differences in monetary investments across tasks (i.e., comparing individual differences in how much subjects invested between each player type). After applying Bonferroni correction for multiple comparisons, we observed a significant effect of investment differences across tasks—(trust-game trustworthy start – trust-game neutral start > slot-machine trustworthy start – slot-machine neutral start), $t(256) = 2.74, p = .007, 95\% \text{ CI for the mean difference} = [0.007, 0.043]$ —revealing that subjects are more biased and susceptible to first impressions in the social domain (Fig. 2b).

Behavioral results for subjects with anxiety

We next examined whether the observed patterns of behavior seen in healthy controls differed from those in subjects with anxiety. Results revealed that just as observed with healthy controls, subjects with anxiety invested significantly more money in the slot-machine game compared with the trust game, $t(96) = -2.53, p = .013, d = 0.13, 95\% \text{ CI for the mean difference} = [-0.07, -0.008]$, and there was no significant difference in mean investments between healthy controls and subjects with anxiety, $F(1, 352) = 3.071, p = .081$, although this effect was trending. Furthermore, a pattern of asymmetrical investment in player types similar to that seen in healthy controls was also observed in subjects with anxiety—repeated measures ANOVA on trust-game player type, $F(2, 192) = 4.718$ (Greenhouse-Geisser corrected), $p = .011, d = 0.44$ —so that the trustworthy-start player received the most money (Fig. 2a). However, unlike with healthy controls, these asymmetries did not extend to the nonsocial domain—repeated measures ANOVA on slot-machine type: $F(2, 192) = 1.49, p = .228$.

A group comparison on overall earnings (i.e., remaining investment not sent to partner + amount returned to subject) revealed that subjects with anxiety earned significantly less money than healthy controls in both the trust game and the slot-machine game—mixed design ANOVA on overall earnings as a function of Group (subjects with anxiety vs. healthy subjects) \times

Condition (trust game vs. slot-machine game): main effect of anxiety, $F(1, 352) = 8.106, p = .005$, generalized $\eta^2 = .018$; main effect of condition, $F(1, 352) = 6.53, p = .011$, generalized $\eta^2 = .003$. To further probe these differences, we examined mean investments across task blocks. Breaking out investments across net-gain (positive valence) and net-loss (negative valence) trials (Fig. 3a) revealed that, compared with healthy controls, subjects with anxiety uniquely overinvested in social partners, particularly during loss blocks—mixed design ANOVA on investments as a function of Group (subjects with anxiety vs. healthy subjects) \times Valence (gain vs. loss): main effect of valence, $F(1, 352) = 713.28, p < .001$, generalized $\eta^2 = .35$; Anxiety \times Valence interaction, $F(1, 352) = 6.30, p = .012$, generalized $\eta^2 = .005$. These effects were particularly prevalent when reward dynamics were downward trending (see Fig. 3b).

More specifically, in the trust game, subjects with anxiety gave significantly more money during negative-valence blocks for both the neutral-start player—mixed-design Group \times Valence ANOVA on investments of the neutral-start player: main effect of anxiety, $F(1, 352) = 4.60, p = .033$, generalized $\eta^2 = .009$; Anxiety \times Valence interaction, $F(1, 352) = 4.54, p = .034$, generalized $\eta^2 = .004$ —and the trustworthy-start player—mixed-design Group \times Valence ANOVA on investments of the trustworthy-start player; Anxiety \times Valence interaction, $F(1, 352) = 7.71, p = .006$, generalized $\eta^2 = .007$ —compared with healthy controls. This suggests that individuals with anxiety are slower to learn the statistics of negative outcomes relative to healthy controls. Conversely, these effects were observed to a lesser degree in the slot-machine game, in which subjects with anxiety invested significantly more (compared with healthy controls) only during negative-valence blocks for the positive-start machine—mixed-design ANOVA with Group \times Valence, $F(1, 352) = 4.14, p = .043$, generalized $\eta^2 = .004$ —although this effect was trending for the neutral-start machine, $F(1, 352) = 2.95, p = .087$.

Modeling results

To further probe these learning differences, we next examined trial-by-trial learning effects using our computational models. First, we wanted to determine which of our models (DBRL, BRL, or standard RL) best fitted the data. Accordingly, using Bayesian model selection, we compared the relative fit of our DBRL, BRL, and RL models in both games, finding the DBRL model to be the winning model across both the healthy group and the group with anxiety (protected exceedance probability, or $\text{pxp} > .99$). However, when using pairwise comparisons, we observed no clear model-fit difference between the DBRL and BRL models for the group with anxiety (slot-machine game: $\text{pxp} = .54$, trust game:

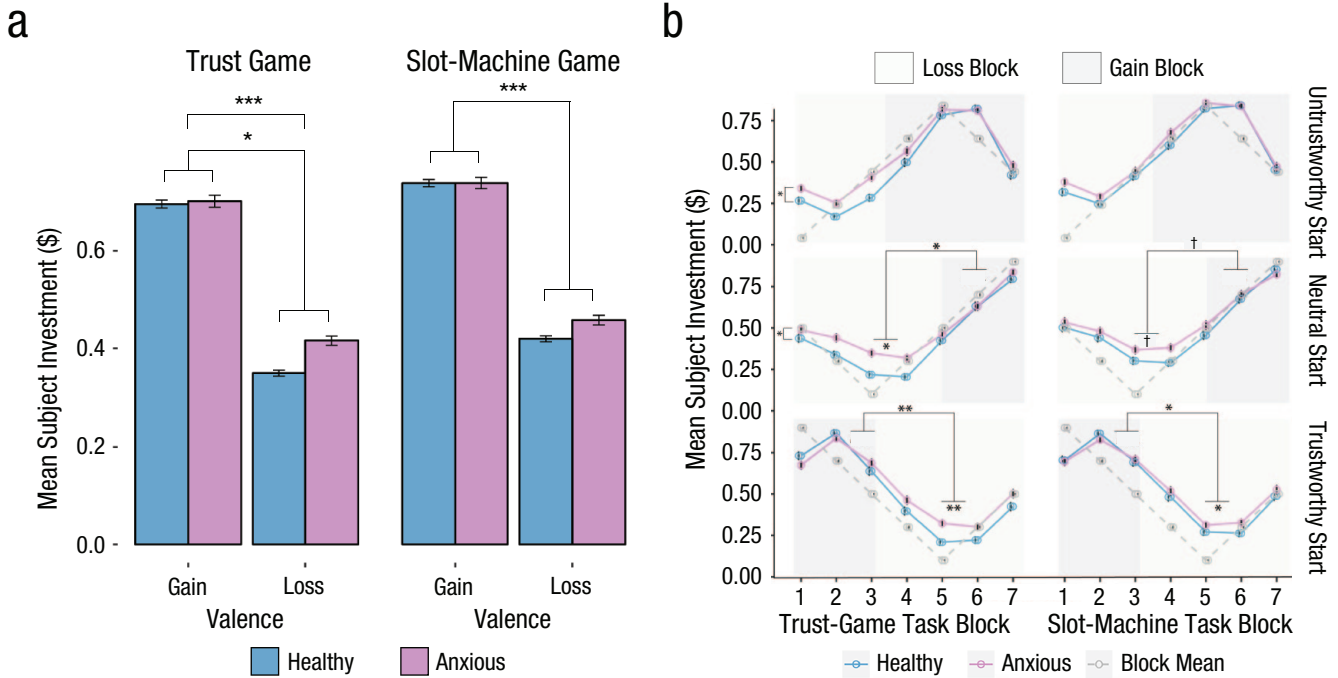


Fig. 3. Learning differences. Mean investments of healthy subjects and subjects with anxiety (a) are shown for gain and loss blocks in the trust game and slot-machine game, collapsed across player type. Learning curves for healthy subjects and subjects with anxiety are illustrated in (b), separately for each player type and each game. The gray dashed line in each graph corresponds to the rescaled proportion of return that the online player was set to per block. The lighter gray shading indicates negative-valence task blocks, in which subjects always lost money. The darker gray shading indicates positive-valence task blocks, in which subjects always earned additional money. Significant and marginally significant differences between block types in (a) and between task blocks in (b) are indicated by symbols ($^{\dagger}p < .10$, $*p < .05$, $**p < .01$, $***p < .001$). Error bars correspond to standard errors of the mean.

$pxp = .56$), suggesting that although healthy subjects were clearly best fitted by the DBRL model (slot-machine game: $pxp > .99$, trust game: $pxp > .99$), learning in subjects with anxiety could be explained equally well by either learning model (Fig. 4a). Because the key difference between the DBRL and BRL models was that the DBRL model uniquely incorporated changes in environmental uncertainty to govern the decay rate (i.e., the level of forgetting of past outcomes), our model-comparison results suggest that healthy subjects, who were clearly better fitted by the DBRL model, were likely using changes in task-level uncertainty to effectively guide behavior, whereas subjects with anxiety were less likely to exhibit these effects.

Learning differences between healthy controls and subjects with anxiety

Our next critical question was whether healthy controls and individuals with anxiety learn differently about the statistics of the environment when there is greater uncertainty. We thus compared the decay rate—which assesses an individual’s flexibility in learning—between our healthy group and our group with anxiety. We examined whether these groups would exhibit differences in how much they flexibly adapted to new

positive, as opposed to new negative, feedback from the player ($\gamma_{0_{\text{pos}}} - \gamma_{0_{\text{neg}}}$). Because decay allows for flexibility in updating, weighting past rewards relative to losses ($\gamma_{0_{\text{pos}}} > \gamma_{0_{\text{neg}}}$) should bias subjects toward consistently overinvesting. (See Fig. S7 in the Supplemental Material for more information on decay-rate difference and optimal investment.) Conversely, perseverating on past losses relative to rewards produces a bias toward underinvesting ($\gamma_{0_{\text{neg}}} > \gamma_{0_{\text{pos}}}$).

Both healthy subjects and subjects with anxiety showed a general bias toward weighting rewards more heavily than losses in the slot-machine game compared with the trust game ($\gamma_{0_{\text{pos}}} > \gamma_{0_{\text{neg}}}$)—mixed-design ANOVA (group: healthy subjects vs. subjects with anxiety) \times (condition: trust game vs. slot-machine game) on decay-rate difference: main effect of condition, $F(1, 352) = 12.94$, $p < .001$, generalized $\eta^2 = .006$ (see Fig. 4b). This resulted in subjects’ overinvesting in slot machines that had previously reaped monetary windfalls. However, only healthy subjects—not those with anxiety—selectively adjusted their learning in social contexts by demonstrating a greater likelihood of weighting losses more heavily than rewards ($\gamma_{0_{\text{neg}}} > \gamma_{0_{\text{pos}}}$). These results reveal that social context selectively influences the differential impact of positive and negative feedback on reward learning in healthy subjects but not in subjects with anxiety—mixed-design

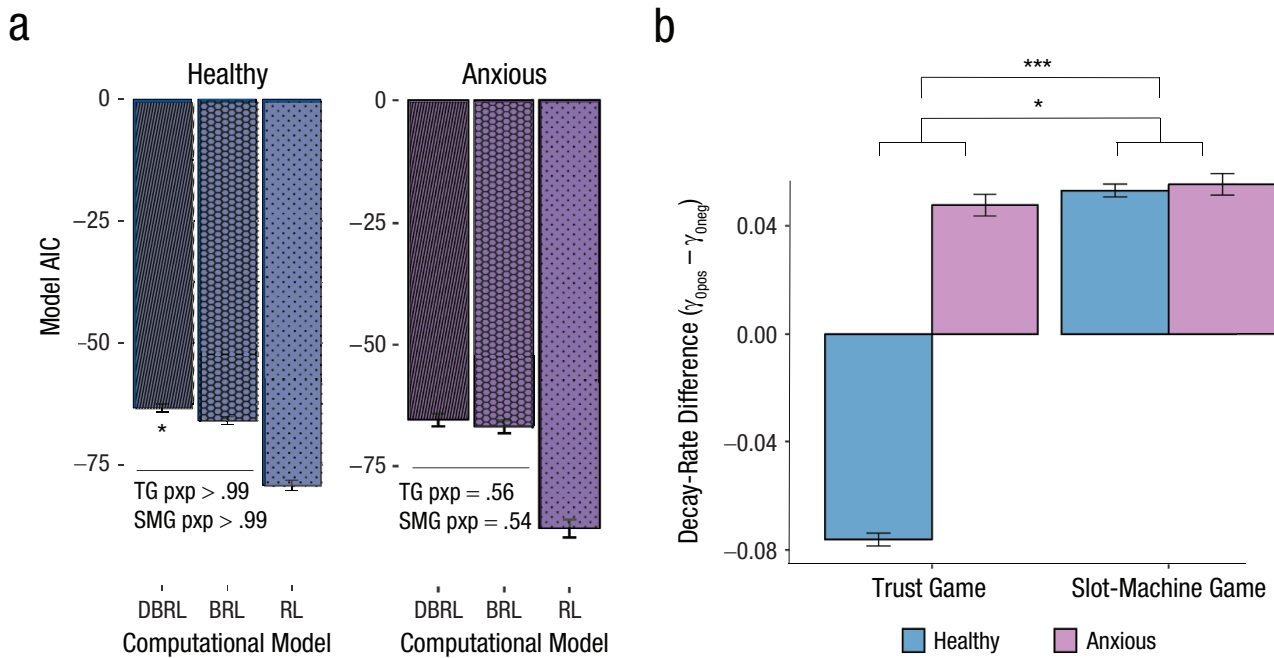


Fig. 4. Model comparisons. Akaike information criteria (AICs) are shown (a) for the Bayesian comparisons among the reinforcement-learning (RL), Bayesian RL (BRL), and dynamic BRL (DBRL) models, separately for healthy subjects and subjects with anxiety. For each group, the protected exceedance probability (pxp) is shown for the trust game (TG) and slot-machine game (SMG). Error bars correspond to standard errors of the mean, and the asterisk indicates the best-fitting model. In (b), mean estimates of decay-rate difference ($\gamma_{0_{\text{pos}}} - \gamma_{0_{\text{neg}}}$) are shown for each game, separately for healthy subjects and subjects with anxiety. Decay-rate differences index the extent to which individuals exhibit asymmetries in learning from past positive compared with past negative experiences. Error bars in (b) show standard errors. Significant differences between individuals and between games in (b) are indicated by asterisks ($*p < .05$, $***p < .001$).

Group \times Condition ANOVA: Group \times Condition interaction, $F(1, 352) = 4.14$, $p = .043$, generalized $\eta^2 = .005$.

Post hoc pairwise comparisons further revealed that there was a significant difference in decay rate between subjects with anxiety and healthy controls in the trust game, $t(352) = -2.57$, $p = .011$, whereas no differences emerged in the slot-machine game, $t(352) = -0.045$, $p = .96$. In other words, subjects with anxiety were generally susceptible to heavily weighting rewards, specifically when the decision involved other people. This learning pattern can be observed in the trajectories displayed in Figure 3a, which shows that subjects with anxiety were slower to learn when to stop investing in an exploitative social partner, relative to healthy subjects. We did not observe any significant differences in the extent to which subjects differed in their sensitivity to changes in uncertainty ($\gamma_{1_{\text{pos}}}$ and $\gamma_{1_{\text{neg}}}$), but as noted above, the inclusion of these parameters improved model fit only in healthy subjects.

Discussion

Learning under uncertainty is a daily endeavor, yet little is known about how this relationship unfolds in the social domain, where uncertainty is likely to be heightened

because of the noisy and ambiguous nature of social interactions. Previous work has illustrated that social and nonsocial reward learning are governed by largely overlapping neural circuitry, suggesting a domain-general account of social learning (Behrens et al., 2008). However, in the current study, we directly compared social and nonsocial learning under uncertainty, finding that healthy individuals exhibit distinct learning profiles across contexts; specifically, positive first impressions unduly biased subsequent learning—a finding uniquely observed in the social domain. Moreover, these healthy individuals learned asymmetrically from rewards and losses (i.e., weighting rewards more heavily than losses) in nonsocial contexts, resulting in consistent overinvesting during negative-prediction-error trials (i.e., when a slot machine resulted in a monetary loss). Effectively, healthy subjects kept betting on a previously rewarding slot machine, even though the evidence suggested that the effort was no longer worth it monetarily. Conversely, in the social domain, healthy individuals changed their learning pattern entirely so that they were more likely to weight losses (defections) more heavily than rewards (reciprocations). This suggests that healthy individuals were able to successfully recognize exploitative social behavior,

which led to a timely termination of the relationship. Given that both tasks were matched in their reward dynamics, these findings demonstrate that the structure of social environments recruits specific priors and computations that selectively modulate learning to reshape the way we process social information.

Our study also provides evidence that the ability to selectively adjust learning across contexts to avoid social exploitation is biased by one's sensitivity to uncertainty. Previous work in the nonsocial domain shows that individuals with trait anxiety have particular difficulty learning the statistics of volatile environments; however, there is no consensus about the nature of this aberrant learning (Aylward et al., 2019; Browning et al., 2015). Here, we showed that individuals with anxiety exhibit learning differences that are uniquely exacerbated by uncertainty in social contexts; as a result, they consistently overinvest in exploitative partners. The fact that individuals with anxiety consistently overinvested during loss blocks (i.e., negative-prediction-error trials in which a player routinely defects) could also indicate use of an alternative decision policy—one in which subjects were strategically forgoing monetary gains to promote trust and cooperation in exploitative social partners (Chang & Smith, 2015). Although this possibility may represent learning differences at the level of decision making rather than at the level of uncertainty perception, this specific pattern of behavioral rigidity highlights the importance of continued work to examine the link between uncertainty sensitivity and learning in anxiety disorders.

Our results also suggest that when people make adaptive decisions, computational dynamics of social learning under uncertainty likely involve the joint combination of uncertainty sensitivity and the ability to update beliefs in a flexible manner. We show that associative learning, indexed through classic RL, was a poorer fit to subject-specific data compared with both types of BRL models. The most sophisticated dynamic model allowed uncertainty to change dynamically with concomitant effects on learning. This model provided the best fit to social behavior, particularly in healthy subjects. However, the evidence favoring the more sophisticated DBRL model was weaker in subjects with anxiety, who were generally fitted equally well by both types of Bayesian models. The reduced fit of the dynamic model, which uniquely incorporated fluctuations in uncertainty into belief updating, provides some evidence that individuals with anxiety are less sensitive to environmental uncertainty in their behavior.

Together, these findings provide the first evidence we are aware of that learning under uncertainty uniquely unfolds across social and nonsocial contexts while also highlighting a candidate mechanism for how this process occurs. Future research should further explore how

uncertainty perception affects downstream learning and decision-making.

Transparency

Action Editor: Timothy J. Pleskac

Editor: D. Stephen Lindsay

Author Contributions

A. Lamba, O. FeldmanHall, and M. J. Frank designed and developed the study concept. Testing and data collection were performed by A. Lamba. A. Lamba analyzed and interpreted the data under the supervision of O. FeldmanHall and M. J. Frank. All authors jointly wrote and approved the final version of the manuscript for submission.

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Funding


Funding was provided by the Centers of Biomedical Research Excellence (COBRE).

Open Practices

All data from this study have been made publicly available via the Open Science Framework and can be accessed at <https://osf.io/ea67f/>. The design and analysis plans were not preregistered. The complete Open Practices Disclosure for this article can be found at <http://journals.sagepub.com/doi/suppl/10.1177/0956797620910993>. This article has received the badge for Open Data. More information about the Open Practices badges can be found at <http://www.psychologicalscience.org/publications/badges>.



ORCID iD

Oriel FeldmanHall  <https://orcid.org/0000-0002-0726-3861>

Supplemental Material

Additional supporting information can be found at <http://journals.sagepub.com/doi/suppl/10.1177/0956797620910993>

References

- Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning, 47*, 235–256. doi:10.1023/a:1013689704352
- Aylward, J., Valton, V., Ahn, W.-Y., Bond, R. L., Dayan, P., Roiser, J. P., & Robinson, O. J. (2019). Altered learning under uncertainty in unmedicated mood and anxiety disorders. *Nature Human Behaviour, 3*, 1116–1123. doi:10.1038/s41562-019-0628-0
- Barrera, T. L., & Norton, P. J. (2009). Quality of life impairment in generalized anxiety disorder, social phobia, and panic disorder. *Journal of Anxiety Disorders, 23*, 1086–1090. doi:10.1016/j.janxdis.2009.07.011
- Behrens, T. E., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. (2008). Associative learning of social value. *Nature, 456*, 245–249. doi:10.1038/nature07538

- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*, 1214–1221. doi:10.1038/nn1954
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis*, *20*, 351–368.
- Boelen, P. A., & Reijntjes, A. (2009). Intolerance of uncertainty and social anxiety. *Journal of Anxiety Disorders*, *23*, 130–135. doi:10.1016/j.janxdis.2008.04.007
- 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*, 590–596. doi:10.1038/nn.3961
- Carleton, R. N., Mulvogue, M. K., Thibodeau, M. A., McCabe, R. E., Antony, M. M., & Asmundson, G. J. (2012). Increasingly certain about uncertainty: Intolerance of uncertainty across anxiety and depression. *Journal of Anxiety Disorders*, *26*, 468–479. doi:10.1016/j.janxdis.2012.01.011
- Chang, L. J., & Smith, A. C. (2015). Social emotions and psychological games. *Current Opinion in Behavioral Sciences*, *5*, 133–140.
- Collins, A. G., & Frank, M. J. (2014). Opponent actor learning (OpAL): Modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive. *Psychological Review*, *121*, 337–366. doi:10.1037/a0037015
- Daw, N. D. (2014). Advanced reinforcement learning. In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics: Decision making and the brain* (2nd ed., pp. 299–320). San Diego, CA: Academic Press.
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, *18*, 185–196. doi:10.1016/j.conb.2008.08.003
- Diaconescu, A. O., Mathys, C., Weber, L. A., Daunizeau, J., Kasper, L., Lomakina, E. I., & Stephan, K. E. (2014). Inferring on the intentions of others by hierarchical Bayesian learning. *PLOS Computational Biology*, *10*(9), Article e1003810. doi:10.1371/journal.pcbi.1003810
- Diaconescu, A. O., Mathys, C., Weber, L. A. E., Kasper, L., Mauer, J., & Stephan, K. E. (2017). Hierarchical prediction errors in midbrain and septum during social learning. *Social Cognitive and Affective Neuroscience*, *12*, 618–634. doi:10.1093/scan/nsw171
- Eng, W., Heimberg, R. G., Hart, T. A., Schneier, F. R., & Liebowitz, M. R. (2001). Attachment in individuals with social anxiety disorder: The relationship among adult attachment styles, social anxiety, and depression. *Emotion*, *1*, 365–380.
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175–191.
- FeldmanHall, O., & Shenhav, A. (2019). Resolving uncertainty in a social world. *Nature Human Behaviour*, *3*, 426–435. 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*, Article e12029. doi:10.7554/eLife.12029
- Gershman, S. J. (2016). Empirical priors for reinforcement learning models. *Journal of Mathematical Psychology*, *71*, 1–6. doi:10.1016/j.jmp.2016.01.006
- Gillan, C. M., & Daw, N. D. (2016). Taking psychiatry research online. *Neuron*, *91*, 19–23.
- Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics*, *14*, 399–425.
- Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A Bayesian foundation for individual learning under uncertainty. *Frontiers in Human Neuroscience*, *5*, Article 39. doi:10.3389/fnhum.2011.00039
- McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, *84*, 870–881. doi:10.1016/j.neuron.2014.10.013
- Montague, P. R., Dayan, P., Person, C., & Sejnowski, T. J. (1995). Bee foraging in uncertain environments using predictive Hebbian learning. *Nature*, *377*, 725–728. doi:10.1038/377725a0
- Nassar, M. R., Wilson, R. C., Heasley, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *The Journal of Neuroscience*, *30*, 12366–12378. doi:10.1523/jneurosci.0822-10.2010
- Niv, Y., Duff, M. O., & Dayan, P. (2005). Dopamine, uncertainty and TD learning. *Behavioral and Brain Functions*, *1*, Article 6. doi:10.1186/1744-9081-1-6
- Niv, Y., Eddlund, J. A., Dayan, P., & Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *The Journal of Neuroscience*, *32*, 551–562. doi:10.1523/JNEUROSCI.5498-10.2012
- O'Reilly, J. X. (2013). Making predictions in a changing world—inference, uncertainty, and learning. *Frontiers in Neuroscience*, *7*, Article 105. doi:10.3389/fnins.2013.00105
- Payzan-LeNestour, E., & Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLOS Computational Biology*, *7*(1), Article e1001048. doi:10.1371/journal.pcbi.1001048
- Platt, M. L., & Huettel, S. A. (2008). Risky business: The neuroeconomics of decision making under uncertainty. *Nature Neuroscience*, *11*, 398–403. doi:10.1038/nn2062
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, *9*, 545–556. doi:10.1038/nrn2357
- Rubin, K. H., & Burgess, K. B. (2001). Social withdrawal and anxiety. In M. W. Vasey & M. R. Dadds (Eds.), *The developmental psychopathology of anxiety* (pp. 407–434). New York, NY: Oxford University Press.
- Rushworth, M. F., & Behrens, T. E. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience*, *11*, 389–397. doi:10.1038/nn2066
- Shapiro, D. N., Chandler, J., & Mueller, P. A. (2013). Using Mechanical Turk to study clinical populations. *Clinical Psychological Science*, *1*, 213–220.

- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, England: MIT Press.
- Wehmeyer, M. L. (2013). *The Oxford handbook of positive psychology and disability*. New York, NY: Oxford University Press.
- Weinstein, A., & Littman, M. L. (2012, June). *Bandit-based planning and learning in continuous-action Markov decision processes*. Paper presented at the Twenty-Second International Conference on Automated Planning and Scheduling, Atibaia, São Paulo, Brazil.
- Yamagishi, T., & Yamagishi, M. (1994). Trust and commitment in the United States and Japan. *Motivation and Emotion, 18*, 129–166. doi:10.1007/bf02249397
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron, 46*, 681–692. doi:10.1016/j.neuron.2005.04.026