Reward and punishment reversal learning in major depressive disorder

Dahlia Mukherjee^{1,2*}, Alexandre L. S. Filipowicz^{2,3,4*}, Khoi Vo⁶, Theodore D. Satterthwaite⁵, & Joseph W. Kable^{2,4}

¹Department of Psychiatry, The Pennsylvania State University
Departments of ²Psychology, ³Neuroscience, ⁴Computational Neuroscience Initiative, and
⁵Psychiatry, University of Pennsylvania

⁶Department of Psychology & Neuroscience, Duke University
*co-first authors

Corresponding author:

Dahlia Mukherjee
Assistant Professor of Psychiatry and Behavioral Health,
Mood Disorders Program
Penn State Milton S. Hershey Medical Center
Department of Psychiatry, H073
500 University Dr, Rm C5613A
Hershey, PA 17033
Phone: 717 531 0003 x 283161
Fax : 717 531 6491
dmukherjee@pennstatehealth.psu.edu

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Abstract

Depression has been associated with impaired reward and punishment processing, but the specific nature of these deficits is less understood and still widely debated. We analyzed reinforcement-based decision-making in individuals diagnosed with major depressive disorder (MDD) to identify the specific decision mechanisms contributing to poorer performance. Individuals with MDD ($n = 64$) and matched healthy controls ($n = 64$) performed a probabilistic reversal learning task in which they used feedback to identify which of two stimuli had the highest probability of reward (reward condition) or lowest probability of punishment (punishment condition). Learning differences were characterized using a hierarchical Bayesian reinforcement learning model. While both groups showed reinforcement learning-like behavior, depressed individuals made fewer optimal choices and adjusted more slowly to reversals in both the reward and punishment conditions. Our computational modeling analysis found that depressed individuals showed lower learning rates and, to a lesser extent, lower value sensitivity in both the reward and punishment conditions. Learning rates also predicted depression more accurately than simple performance metrics. These results demonstrate that depression is characterized by a hyposensitivity to positive outcomes, which influences the rate at which depressed individuals learn from feedback, but not a hypersensitivity to negative outcomes as has previously been suggested. Additionally, we demonstrate that computational modeling provides a more precise characterization of the dynamics contributing to these learning deficits, and offers stronger insights into the mechanistic processes affected by depression.

Keywords: major depressive disorder, reinforcement learning, reversal learning, decision-making, computational psychiatry

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Major depressive disorder (MDD) is characterized by anhedonia and depressed mood. These symptoms may be driven by abnormal responses to affective stimuli, such as rewards and punishments (Chen, Takahashi, Nakagawa, Inoue, & Kusumi, 2015; Eshel & Roiser, 2010). Given that depressive disorders are ranked as the leading cause of disability worldwide, with lifetime prevalence close to 20% (Moussavi et al., 2007), better understanding how responses to rewards and punishments are altered in MDD could have a sizable impact by leading to more targeted treatment and prevention.

Although there is widespread evidence that individuals with MDD differ in their responses to rewards and punishments compared to healthy individuals, the relative severity and relationship between reward and punishment impairments is less clear (Chen et al., 2015). One prominent idea is that depression is associated with hyposensitivity to reward and hypersensitivity to punishment (Eshel & Roiser, 2010). However, this conclusion is often based on results from different paradigms in the two domains. For example, in the reward domain, one line of research has examined how receiving asymmetric rewards biases responding in a perceptual judgment task (Pizzagalli, Jahn, & O'Shea, 2005). Compared to non-depressed participants, nonclinical depressed participants and individuals with MDD show lower levels of response bias (Pizzagalli, Iosifescu, Hallett, Ratner, & Fava, 2008; Pizzagalli et al., 2005). This decreased reward bias persists during remission and predicts worse outcomes during treatment (Pechtel, Dutra, Goetz, & Pizzagalli, 2013).

In contrast, depression has also been linked, albeit inconsistently, with hypersensitivity to errors in cognitive tasks (Beats, Sahakian, & Levy, 1996; Chen et al., 2015; Elliot, Sahakian, Herrod, Robins, & Paykel, 1996; Eshel & Roiser, 2010). Depressed individuals tend to overreact to the errors they commit in these tasks, which impairs their subsequent performance (Beats et

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

al., 1996; Elliot et al., 1996; Eshel & Roiser, 2010; Steffens, Wagner, Levy, Horn, & Krishnan, 2001).

Sensitivity to different forms of feedback in depression has also been studied in instrumental learning tasks, where participants must learn to select the actions that lead to the most favorable outcomes on the basis of probabilistic feedback. For example, learning tendencies can be characterized by how likely participants are to repeat an action after positive feedback ('win-stay') or to switch their action after negative feedback ('lose-switch'). Some studies have found that depressed individuals show higher proportions of 'lose-switch' responses than healthy controls, suggesting a heightened sensitivity to negative feedback (Dombrovski et al., 2015; Dombrovski, Szanto, Clark, Reynolds, & Siegle, 2013; Murphy, Michael, Robbins, & Sahakian, 2003). However, findings have been mixed, with other studies failing to find depression-related differences in 'lose-switch' responses (Chase et al., 2010; Dombrovski et al., 2010; Gradin et al., 2011).

Here we use probabilistic reversal learning tasks to examine responses to rewards and punishments in depression under otherwise structurally identical task conditions. For the most part, previous studies using this paradigm have provided either cognitive (i.e., "correct" versus "incorrect") or reward feedback (i.e., presence versus absence of reward), and therefore do not distinguish positive and negative feedback from rewards and punishments. Here we set up two incentivized probabilistic reversal learning tasks where participants had to learn either to maximize reward (i.e., winning money versus nothing) or to minimize punishment (i.e., losing money versus nothing). In another report using these data, we have shown that depressed individuals perform significantly worse than healthy controls in both reward and punishment learning, and that these differences are larger than on several other decision-making tasks

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

(Mukherjee, Lee, Kazinka, Satterwaite, & Kable, 2020). Here we aim to examine whether similar or different mechanisms account for impaired performance in learning from rewards versus punishments: whether depressed individuals are hyposensitive to reward and/or hypersensitive to punishment, both, or neither.

To examine the reasons for impaired performance in reward and punishment learning, we turn to computational modeling, as it allows us to distinguish between a number of different potential causes for impairment on these tasks (Huys, Maia, & Frank, 2016; Montague, Dolan, Friston, & Dayan, 2012; Wang & Krystal, 2014). Reinforcement learning (RL) models provide a framework to characterize how feedback is used to learn in different environments (Sutton & Barto, 1998) and have proven a popular class of models for studying reward processing in depressed individuals (Chen et al., 2015; Huys, Daw, & Dayan, 2015; Huys, Pizzagalli, Bogdan, & Dayan, 2013). In addition, these models can dissociate two different sources of learning deficits: changes in learning rate, or the extent to which subjects adjust their behavior in response to feedback, and changes in value sensitivity, or the difference in desirability or subjective value of the outcomes. In the reward domain, early modeling attempts have attributed learning deficits in depression to lower learning rates (Chase et al., 2010; Chen et al., 2015), though more recent studies have argued that, rather than affecting learning, depression is associated with a lowered reward sensitivity (Huys et al., 2015; Huys et al., 2013).

Here we used hierarchical Bayesian modeling to identify whether the learning deficits we observed in both reward and punishment learning environments could be attributed to changes in learning-rate, sensitivity, or both, and also whether similar or different changes occur in the reward and punishment contexts. Moreover, we sought to rule out other potential biases that could affect performance, such as perseverative behaviors or biases towards particular items or

actions, irrespective of the task feedback. Finally, we test the benefit of fit RL parameters over-and-above behavioral metrics in classifying individuals with MDD (Huys et al., 2016; Montague et al., 2012; Wang & Krystal, 2014).

To preview our results, we find that the sources of learning deficits in MDD were similar in both the reward and punishment environments. In both tasks, poor performance in our MDD group was linked to both reduced learning rates and, to a lesser extent, reduced value sensitivity. We also find that learning rates were the strongest predictor of depression in both environments, and predicted better than any other model parameter or behavioral metric. Together these findings provide more robust insights into the specific feedback processing mechanisms affected by depression.

Methods & Materials

Participants

Between October 2012 and January 2014, 128 participants (64 diagnosed with MDD and 64 healthy controls) were recruited for the study. MDD participants were recruited through flyers in the Department of Psychiatry and Behavioral Health, Counseling and Psychological Services and the Hospital of the University of Pennsylvania, and through referrals from treatment research studies in the Department of Psychiatry. Healthy control participants were recruited from the staff and advanced students of the university through flyers posted in the Departments of Psychology and Psychiatry, the Law School, the Graduate Student Office, and the Hospital. Participants likely to meet study criteria based on an initial phone screen were invited for a diagnostic interview. MDD participants were enrolled if (1) they met the diagnostic criteria for current MDD episode, (2) had no history of substance abuse/dependence in the past 6 months, and (3) had no history of bipolar disorder and/or psychotic episodes. Potential comorbidities

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

beyond bipolar, substance use or psychotic disorders were not assessed. Diagnostic criteria were determined based on the Structured Clinical Interview for *Diagnostic and Statistical Manual of Mental Disorders, Fourth Edition* (SCID/DSM-IV; APA, 1994). Fifty-six percent (36 out of 64) of the MDD participants were referred directly from another clinical study or clinical service based on an MDD diagnosis. Within the MDD group, 10 were on medication only, 14 were receiving therapy only, 19 were receiving both medication and therapy, and 19 were in no treatment for depression (treatment data from 2 individuals were missing). Inclusion criteria for controls included absence of current or past psychiatric illness, as assessed by the SCID, and the absence of any psychotropic medications.

The study was approved by the Institutional Review Board at the University of Pennsylvania and all participants provided written informed consent. The clinical interview and study procedures were conducted by a master's level trained clinical psychologist (DM). Data from the reward and punishment reversal learning tasks were collected as part of a larger study investigating value-based decision-making in individuals diagnosed with MDD (Mukherjee et al., 2020). Participants were paid \$15/hr and received an additional incentive based on their choices in one randomly selected decision task out of the eight administered. The components of the study were administered as follows: administration of the structured clinical interview, performance on the eight decision-making tasks (including the reward and punishment probabilistic reversal learning tasks), performance on the similarities and matrix reasoning subtests of the Wechsler Abbreviated Intelligence Scale (WASI; Wechsler, 1999), and completion of self-report measures.

Measures

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Clinical Measures. Inclusion and exclusion diagnostic criteria were based on the SCID. Depression severity was assessed with the Beck Depression Inventory (BDI; Beck, Steer, & Brown, 1996). Current medication and psychotherapy treatment were self-reported.

Self-Report Measures. In addition to the BDI, the participants also completed the Rosenberg Self-Esteem Scale (Rosenberg, 1965), the Cognitive Behavioral Avoidance Scale (Ottenbreit & Dobson, 2004), the Behavioral Inhibition and Behavioral Activation Scale (Carver & White, 1994), the Snaith Hamilton Pleasure Scale (Snaith et al., 1995), the Beck Anxiety Inventory (Aaron Beck, Epstein, Brown, & Steer, 1988) and the Depression, Anxiety and Stress Scale (Lovibond & Lovibond, 1995).

Cognitive Measure. Cognitive ability was assessed with the Wechsler Abbreviated Scale of Intelligence - Second Edition (WASI-II). For the purpose of efficiency, two subtests (matrix reasoning and similarities) of the WASI were administered to each subject to obtain a full-scale IQ score.

Reward Probabilistic Reversal Learning Task. In this task, subjects chose between two distinct fractal stimuli (Figure 1a), which were positioned randomly at one of two locations (left and right of the central white dot) on screen. On each trial, participants pressed a keyboard button to choose one of the fractals. Positive feedback was provided if a fractal was rewarded (picture of a quarter); otherwise, neutral feedback was provided (a red dot). The fractals were probabilistically rewarded, with the richer fractal rewarded 70% of the time and the poorer fractal rewarded 30% of the time. Participants were not informed of the specific underlying reward structure of the task but were informed that on any given trial, one fractal had a higher likelihood of delivering a reward and that this association would reverse periodically throughout the task. All participants completed 4 practice trials before doing a full, 90-trial run. Unsignalled

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

reversals occurred every 30 trials (i.e., the fractal with higher reward probability now had the lower reward probability and vice versa). Each reward had a monetary value of \$0.25. At the end of the task, the screen displayed the total number of quarters the participant won.

Punishment Probabilistic Reversal Learning Task. This task was similar to the reward probabilistic reversal learning task except the goal for the participant was to try to avoid choosing the fractal leading to punishment (signified by a red cross overlaying a quarter; Figure 1b). The richer fractal was punished 30% of the time and the poorer fractal was punished 70% of the time. Participants were informed that one fractal image led to more losses than the other and that this would switch periodically. The participant started with \$22.50 and lost \$0.25 each time they received punishment feedback. Like the reward task, subjects did 4 practice trials before completing the full 90-trials run. On completion of the task, the screen displayed the total number of quarters the participant lost.

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Figure 1. Reversal learning task schematic. (a,b) At the start of each trial in both conditions, subjects chose one of two fractals that they thought would either give them the highest chance of reward (reward condition) or lowest chance of punishment (punishment condition). (a) In the reward condition, a quarter appeared if the subject's choice resulted in a reward. (b) In the punishment condition, a quarter with a red 'X' appeared if the subject's choice resulted in punishment. A red circle appeared for unrewarded or unpunished trials in the reward and punishment conditions respectively. (c) At the start of the task, one fractal had a 70% probability of reward/punishment whereas the other had a 30% probability. Reward/punishment probabilities switched every 30 trials. The fractal initially assigned the higher probability was counterbalanced across subjects.

Performance Analysis

Behavioral Metrics. We analyzed several basic measures of performance, including the proportion of rich choices and the proportion of win-stay and lose-shift choices. Rich choices were choices of the fractal with the higher probability of more positive outcomes (i.e., reward, for the reward learning task, or no punishment, for the punishment learning task). Win-stay choices were when a participant repeated the same choice (same fractal image as the previous trial) when their previous choice led to a positive outcome (reward or no punishment). Lose-shift choices were when a participant switched fractal choices when their previous choice led to a negative outcome (no reward or punishment).

Linear Regression Model. In order to test whether participant choice behavior was consistent with reinforcement learning, we fit a logistic regression to choice data (Lau & Glimcher, 2008). Behavior consistent with reinforcement learning should demonstrate an exponential decline of the influence of past rewards.

We tested this assumption using a logistic regression to estimate weights for rewards received (R) and choices (c) made on previous trials (t). This regression estimates the probability of choosing one fractal image, $\Pr(c_t = f_1)$, versus the other, $\Pr(c_t = f_2)$. Since there are only two options, we assumed symmetric weights for the two options, and the model for 10 previous rewards and one previous choice reduced to the following:

$$\log\left(\frac{\Pr(c_t = f_1)}{\Pr(c_t = f_2)}\right) = \sum_{i=t-1}^{t-10} \beta_i (R_{f1,i} - R_{f2,i}) + \beta_{last\ choice} (c_{f1,t-1} - c_{f2,t-1}) + \beta_0 \quad (1)$$

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Here β_i corresponds to the influence of reward received from up to 10 trials in the past on a participant's current fractal choice, $\beta_{last\ choice}$ captures the influence of a participant's last fractal choice (c_{f1} for fractal 1 and c_{f2} for fractal 2), and the β_0 intercept captures biases the subjects may have towards one of the fractals (values near 0 indicating no bias).

Reinforcement Learning Model. Choice data from all participants was fit using a standard reinforcement learning model (Sutton & Barto, 1998; Watkins & Dayan, 1992). The model used the sequence of choices and outcomes to estimate the value of each option for every trial. The values Q for any fractal f on any trial t are updated as trial by trial feedback is received using a standard update rule:

$$Q_{t+1}(f) = Q_t(f) + \alpha_f[R_t - Q_t(f)] \quad (2)$$

where R corresponds to the reward/outcome obtained (coded 1 for reward and 0 for no reward in the reward task, and 1 for no punishment and 0 for punishment in the punishment task) after choosing the fractal f , and α_f corresponds to a subject-specific learning rate which governs the extent to which the model updates its value estimate for the chosen fractal after receiving feedback. The extent to which differences in values between the two options influenced decisions were captured by a standard inverse temperature parameter β_f , which could range between 0 (no influence of fractal values) to $+\infty$ (strong influence of fractal values). While traditionally associated with explore/exploit tradeoffs, this parameter is mathematically equivalent to other formulations of value sensitivity (Huys et al., 2013; Ruppel, Stankevicius, Huys, Steele, & Seriès, 2018).

The model also included additional features that do not improve task performance but could nonetheless influence subject behavior. As highlighted in the results, the choices of both

depressed and healthy individuals were influenced by the previous actions they took. We modeled these tendencies by including a term to capture learning occurring at the level of actions, rather than at the level of fractals. Expected values for actions a were tracked using a very similar form to eq 1.

$$Q_{t+1}(a) = Q_t(a) + [R_t - Q_t(a)] \quad (3)$$

For reasons highlighted in the model fitting section, we omitted a subject-specific learning rate component, effectively fixing the learning rate for actions to a value of one, and modulated by an action sensitivity parameter (β_a). This parameter was allowed to vary between $-\infty$ and $+\infty$, as some subjects seemed to be *repelled* by the action value.

The model also included four parameters meant to capture response biases in subject choices. Two parameters captured perseverative tendencies towards repeating the same fractal choice (ρ_f) or same action (ρ_a) as the previous trial. Positive values indicate tendencies towards repeating the same fractal choice or action from the previous trial, regardless of outcome, whereas negative values correspond to tendencies against repeating the same fractal or action. Two final parameters captured subject specific tendencies towards choosing a particular fractal (π_f) or particular action (π_a) more than the other.

Together these parameters were used to compute the log odds, v , on any trial t that a subject would make a particular choice (eq. 2), with positive log odds corresponding to choosing fractal 1 (f_1), and the inverse logit corresponding to the probability that a subject would choose fractal 1 (eq. 3)

$$v_t = \beta_f[Q_t(f_1) - Q_t(f_2)] + \beta_a[Q_t(a = f_1) - Q_t(a = f_2)] + \rho_f + \rho_a + \pi_f + \pi_a \quad (4)$$

$$\Pr(f_1) = \text{logit}^{-1}(v_t) \quad (5)$$

In total, the model contains seven free parameters $\alpha_f, \beta_f, \beta_a, \rho_f, \rho_a, \pi_{f1}, \pi_a$, which were fit to subjects' data separately for the punishment and reward conditions.

Model fitting and parameter comparisons. The reinforcement learning model was fit as a multilevel hierarchical Bayesian model implemented in the Stan statistical modeling language (Carpenter et al., 2017) using R code (Stan Development Team, 2018). Hierarchical models have been shown to provide more precise subject and group level parameters estimates for different models of behavior (Huys et al., 2011; Huys et al., 2013). A subject i 's parameters were estimated as a joint multivariate distributions $\Pr(\mathbf{h}_i, \boldsymbol{\theta})$, with each component of $\mathbf{h}$ corresponding to a subject specific parameter distribution (i.e., $\mathbf{h}_i \in \{\alpha_{f,i}, \beta_{f,i}, \dots, \pi_{a,i}\}$). This multivariate parameter distribution was assumed to be generated from a group level prior distribution $\Pr(\boldsymbol{\theta})$. Using this prior, we sought the posterior multivariate parameter distributions for each subject, given their data $\mathbf{D}_i$:

$$\Pr(\mathbf{h}_i, \boldsymbol{\theta} | \mathbf{D}_i) \propto \Pr(\mathbf{D}_i | \mathbf{h}_i, \boldsymbol{\theta}) \Pr(\mathbf{h}_i, \boldsymbol{\theta}) \quad (6)$$

Which can also be written as

$$\Pr(\mathbf{h}_i, \boldsymbol{\theta} | \mathbf{D}_i) \propto \Pr(\mathbf{D}_i | \mathbf{h}_i, \boldsymbol{\theta}) \Pr(\mathbf{h}_i | \boldsymbol{\theta}) \Pr(\boldsymbol{\theta}) \quad (7)$$

Learning rate parameters were assumed to be generated from a beta prior, fractal sensitivity parameters generated from a gamma prior, and all other parameters generated from gaussian priors. We sampled from this posterior distribution using the Hamiltonian MCMC algorithm implemented in Stan. Six simultaneous chains were run with an 80 000 iteration warm-up (burn-

in) period, and 170 000 posterior draws. Model convergence was assessed by ensuring that $\hat{R}$ statistics were all near 1 (no lower than .99 and no greater than 1.01; Gelman & Rubin, 1992), effective sample size for each parameter were all greater than 1/100th the number of total samples, and that no divergences occurred during sampling. Initial model parametrizations resulted in a high proportion of divergent samples, which indicate biases in the MCMC sampling and violations of ergodicity (Betancourt & Girolami, 2015). We reduced these divergences by both reducing our sampling step size (i.e., adapt delta was set to .99) and implementing non-centered versions of our model (Betancourt & Girolami, 2015). These divergences can also be caused by models with highly correlated parameters, or parameters that cannot be readily identified given the available data. Initial attempts to fit our model with a learning rate parameter on the action values (e.g., adding a term α_a to eq. 2, similar to the role of α_f in eq. 1) resulted in a high number of divergent samples, indicating that this parametrization provides a poor explanation of our subject choice data. These divergences were eliminated by removing the action learning parameter from the model and assuming a fixed action learning rate of 1. Individual subject parameters used for simulations and to predict each subject's depression status were generated using the mean values from each subject's parameter distributions, (i.e., $E[\alpha_{f,i}], E[\beta_{f,i}], \dots, E[\pi_{a,i}]$).

Subject classifiers. Binary classifiers were trained on behavioral metrics and fitted model parameters to predict out of sample subject categories (depressed or healthy control). To account for multicollinearity between parameters and performance metrics, L1 penalized logistic regression (LASSO), tuned using leave one out cross validation, was used to identify the model parameters and behavioral metrics best able to classify depressed individuals from controls in each task condition. A logistic regression with combinations of the variables identified by the

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

LASSO regressions was then used to classify and identify depressed individuals. These classifiers were trained using a repeated 10-fold cross validation scheme (3 repetitions) using the caret package in R (Kuhn, 2008). The mean area under the receiver operating characteristic curve (AUC) for each repetition of each fold was used to assess out-of-sample classifier performance, where values of 1 indicates perfect classification accuracy and .5 indicate chance performance. Differences in model AUC values were compared using the DeLong method (DeLong, DeLong, & Clarke-Pearson, 1988) implemented using the pROC package (Robin et al., 2011).

Results

Demographic Characteristics

The two groups did not differ with respect to gender, education, ethnicity, or age (independent t-tests, all $ps > .3$; Table 1). The two groups also did not differ in cognitive ability ($p = .35$). Participants in the MDD sample were moderately to severely depressed, with a mean score of 30.03 (SD=10.46) on the BDI-II, which was significantly higher than control subjects ($M = 2.9$, $SD = 4.32$, $p < .001$). Within the MDD group, 45% and 52% were on current medication or in treatment for depression, respectively.

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Table 1

Demographic and Clinical Characteristics of MDD and Healthy Control Groups

Note. ** = $p < .001$

	Depressed		Control		Significance
	N	%	N	%	<i>p</i> value
Gender					.3
Male	27	57.8	33	51.6	
Female	37	42.2	31	48.4	
Ethnicity					.76
African	34	53.1	33	51.6	
Caucasian	24	37.5	23	35.9	
Asian	4	6.2	7	10.9	
Other	2	3.1	1	1.6	
Education					.31
No High	4	6.3	1	1.6	
High	26	41.3	23	35.9	
Associate'	10	15.9	12	18.8	
Bachelor's	12	19	16	25	
Master's	11	17.5	9	14.1	
Doctoral	0	0	3	4.7	
Depression Treatment					<.001**
No	19	29.7	60	95	
Medication Only	10	15.6	0	0	
Therapy Only	14	21.9	0	0	
Both	19	29.7	0	0	
Missing	2	3.1			
	M	SD	M	SD	<i>p</i> value
Age	40.45	13.48	38.53	11.73	.4
WASI	101.32	14.63	103.7	14.1	.35
BDI	30.03	10.46	2.9	4.32	<.001**

Depressed individuals showed poorer reversal learning than healthy controls

Depressed individuals performed worse than healthy controls on both reward and punishment reversal learning. Depressed individuals made significantly fewer rich choices in both the reward (mean proportion rich fractal choices [SD]: MDD = .59 [.10], Controls = .63 [.11]; $p = .033$) and the punishment conditions (MDD = .59 [.10], Controls = .64 [.11]; $p = .003$; Fig. 2 – see below).

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Figure 2. Individuals with MDD (red) made fewer rich choices than controls (blue) in both task conditions. (a) Differences in mean proportion correct responses between MDD (red) and healthy controls (blue) across all trials in the reward and punishment task conditions. (b) Mean proportion correct responses as a function of trial since the last reversal. (c-d) Identical statistics reported in (a) and (b) computed from simulations of subject behavior (100 independent simulations per subject using their fit parameter estimates). Bars and errorbars in (a, c) and lines and shading in (b,d) correspond to means and one standard error of the mean respectively. *: $p < .05$, **: $p < .01$. Grey areas in b) indicate trials after reversal on which rank sum tests $p < .05$.

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

We examined several aspects of performance to determine whether behavior in both reversal learning tasks was consistent with reinforcement learning models, and whether the behavioral deficits in depressed individuals might be captured by these models. We first examined rich choices after a reversal. Consistent with reinforcement learning, both groups gradually adjusted their responses after a reversal, with the depressed group appearing to learn the new contingencies more slowly in both the reward and punishment conditions (Fig. 2b).

Next we examined whether choices were more likely to be repeated if they were rewarded, given that RL performance should show a high proportion of win-stay behavior. In both the reward and punishment conditions, each group had proportions of win-stay choices well above 50% (one sample t-test, all $p < 8.2 \times 10^{-11}$), indicating that both groups were showing RL-like behavior. However, the proportion of win-stay choices was significantly lower for the depressed group in both the reward (mean proportion win-stay choices [SD]: MDD = .74 [.15], Controls = .83 [.17]; $p = .005$) and punishment conditions (MDD = .71 [.22], Controls = .83 [.17]; $p = .001$). In contrast, the proportion of lose-shift responses did not differ between groups in either the reward (MDD = .51 [.15], Controls = .54 [.18]; $p = .252$) or punishment conditions (MDD = .50 [.17], Controls = .54 [.18]; $p = .165$; Fig. 3). The higher proportion of win-stay choices in both groups in both tasks is consistent with reinforcement learning. However, depressed subjects made significantly fewer win-stay responses, suggesting lower sensitivity to learning from positive outcomes.

Figure 3. Individuals with MDD (red) made fewer ‘win stay’ but similar ‘lose shift’ choices compared to controls (blue) in both task conditions. (a-b) Proportion ‘win stay’ and ‘lose switch’ between subject groups in the reward condition (a) and punishment condition (b). (c-d) Proportion ‘win stay’ and ‘lose shift’ responses from simulated behavior using individual subject parameter estimates (100 simulations per subject). Bars and errorbars correspond to means and one standard error of the mean respectively. *: $p < .05$, **: $p < .01$.

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Looking at win-stay and lose-shift behavior also revealed an influence of action learning on task performance. In both the reward and punishment learning tasks, depressed and controls were significantly more likely to make win-stay responses when the specific motor action repeated across trials ($ps < .05$; Figure 4a-b). There was also a tendency for healthy controls to make more lose-shift response when the motor response for the previously selected fractal repeated across trials in the reward task, and a similar tendency in depressed individuals in the punishment task (all $ps < .05$).

Figure 4. Both MDD (red) and healthy controls (blue) show action learning biases. (a,b) Subject proportion ‘win-stay’ (WS) and ‘lose-shift’ (LS) responses as a function of whether the fractals were on the same (Same) or different (Diff) side as the previous trial in both the reward (a) and punishment (b) conditions. (c-d) Metrics from a,b for simulated behavior using individual subject parameter estimates (100 independent simulations per subject). Bars and errorbars correspond to means and one standard error of the mean respectively. $\cdot$: $p < .1$, $*$: $p < .05$, $**$: $p < .01$.

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Finally, we fit a logistic regression model to estimate the influence of outcomes received in the past on the current choice. Behavior consistent with reinforcement learning should show an exponential decrease in the influence of outcomes received from previous trials. As is evident from Figure 5, both groups showed this exponential decay, highlighting that both groups used RL-like learning processes. Additionally, controls showed a steeper decay of influence from previous outcomes, with choices depending on outcomes received up to two trials back, whereas depressed individuals were influenced by outcomes received up to three trials back. Such a steeper decay is consistent with higher learning rates in healthy controls.

Figure 5. Both subject groups show learning dynamics consistent with reinforcement learning dynamics. Beta weights from a logistic regression (Eq. 1) measuring the influence of reward 1-10 trials back (R-1 – R-10), the last fractal choice, and a bias towards either fractal in both the reward and punishment conditions.

Depressed individuals had lower learning rates and lower value sensitivity than controls

Since subject behavior in both groups showed tendencies consistent with reinforcement learning, we fit a reinforcement learning model to each subject's behavior and compared parameter values between groups to identify the learning components affected in depressed individuals. The main learning components in the model were the learning rate and value sensitivity for the fractals (α_f and β_f), which capture how much participants incorporate trial-by-trial feedback, and how sensitive they are to value differences between the different fractals. Our model also included parameters that could capture biases in subject performance, including action learning (β_a), perseverative biases (ρ_f and ρ_a), and choice biases (π_f and π_a). These parameters were fit to subject data using a hierarchical Bayesian model with group-level priors over each parameter, providing more precise subject-by-subject estimates (see methods). Although the bias parameters do not aid performance on this task, they each captured subject tendencies across a range of performance metrics, including their propensity to repeat or favor a choice or an action (Fig. 6).

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Figure 6. Spearman correlations between RL model parameters (y – axis) and behavioral variables (x – axis) in the Reward and Punishment conditions for both depressed individuals (MDD - red) and healthy controls (blue). Behavioral metrics considered were a) proportion win-stay responses (Win-Stay), b) mean proportion rich fractal choice (Prop. Rich Fractal), c) difference between proportion of win-stay responses when a rewarded fractal was repeated on the same side (WS Same) or different side (WS Diff), d) proportion of trials on which subjects repeated the same fractal choice (Repeat Choice) or same action (Repeat Action) as the previous trial, and e) biases in the proportion of time subjects chose fractal 1 (Prop. Fractal 1 Choice), or chose the fractal on the left (Prop. Left Action). Colored tiles with numbers indicate $p < .05$ with shading corresponding to strength of correlation (darker = stronger).

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

These model fits revealed that depressed individuals had both lower learning rates and slightly reduced value sensitivities. In both the reward and the punishment conditions, depressed participants had lower fractal learning rates (Wilcoxon rank sum test for differences in α_f - Reward: $p=.005$, Punishment: $p=.0007$). The two subject groups also differed in their expected fractal value sensitivity (β_f), with a trending sensitivity difference observed in the reward condition ($p=.085$), and a more pronounced difference in the punishment condition ($p=.04$). Beyond these major differences, controls also differed from depressed participants in their fractal bias parameter (ρ_f) in the punishment condition ($p=.01$), which, consistent with their behavior, captured a slight but non-significant tendency to choose the second fractal more than the first fractal (mean proportion second fractal choices in punishment condition [SE]: Controls=.54 [.01], MDD=.51 [.02]; Wilcoxon rank sum test $p=.15$). The two groups did not differ on any other parameters in either condition (all $ps>.148$; Fig. 7).

Reward Parameter Fits

Punish Parameter Fits

Figure 7. Individuals with major depressive disorder (MDD; red) have lower learning rates and value sensitivity than healthy controls (blue). Histograms and respective cumulative probability plots of individual parameter values fits in the reward and punishment conditions. p -values correspond to Wilcoxon rank-signed tests between subject groups.

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

These model fits captured all of the aspects of task performance and the differences in task performance between the two groups described above. Model simulations produced using each subject's specific parameter fits recapitulate overall learning differences between the depressed and control groups, as well as differences in learning dynamics after reversals (Fig. 2c-d). Model simulations also capture the difference in proportion of win-stay choices, and the lack of difference in lose-shift choices, between depressed individuals and healthy controls (Fig. 3c-d). In addition, the action learning parameter (β_a) in the model allowed it to capture differences in the proportion of win-stay choices when the motor action repeated (WS-Same) compared to when it did not (WS-Diff; Fig. 6).

Model parameters classify depressed versus controls better than simpler metrics of performance

Finally, we examined whether computational model fits could improve the ability to correctly identify depressed individuals from controls, over and above more basic measures of task performance. To account for multicollinearity across parameters and behavioral metrics, LASSO regression was used to identify the model parameters and behavioral variables best able to classify depressed individuals from controls in each task condition. When examining the parameters alone, consistent with the trends in group parameter differences, learning rate (α_f) and value sensitivity for fractals (β_f) contributed most to classification in both the reward and punishment conditions, along with fractal bias (π_f) in the punishment condition. When examining the behavioral metrics, proportion rich fractal choices, proportion win-stay responses, and action repeats (e.g., action perseveration) were the strongest predictors in the reward condition, whereas only proportion rich fractal choices was found as a predictor in the

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

punishment condition. When all parameters and behavioral variables were added in the same LASSO model, learning rate, sensitivity for actions (β_a), proportion rich fractal choices, and fractal and action repetitions were found to be predictors in the reward condition, whereas learning rate, fractal bias, and proportion rich fractal choices were the best predictors in the punishment condition.

In the reward condition, out-of-sample classification accuracy was highest for a model trained using all of the LASSO identified parameters and behavioral metrics in the reward condition (“full” model), and provided better classification accuracy than any of the individual parameters or behavioral variables on their own (AUC for full model=.69, Delong test used to compare AUCs from the full model with other models: all $ps<.03$; Fig. 8a). In the punishment condition, the full model provided better classification accuracy than any of the individual parameters or behavioral variables ($ps<.05$), except for the learning rate ($p=.121$; Fig. 8b).

In both the reward and punishment conditions, the learning rate and proportion rich responses were the most consistent predictors of depression. We ran a final logistic classifier in both conditions that only included these two parameters to assess whether they could perform at the level of each of the full classifier models. Indeed, these reduced classifiers performed as well as the full model classifiers (AUC – Reward Condition: full model=.69, reduced model=.69; $p=.464$; Punishment Condition: full model=.70, reduced model=.71; $p=.730$; Fig. 8a,b).

Figure 8. Learning rate improves out of sample depression classification accuracy beyond simple performance metrics in both task conditions. (a,b) Out-of-sample area under the receiver operating characteristic curve (AUC) for logistic regression models with different sets of LASSO identified parameter and behavioral metrics for the reward (a) and punishment conditions (b). Bars and errorbars correspond to median and bootstrapped standard error estimates of AUC values (2000 bootstrapped samples). *: $p < .05$, **: $p < .01$.

Discussion

In this study we sought a mechanistic understanding as to why depressed individuals learn less effectively from rewards and punishments than controls. Overall, we found that depressed individuals learn from feedback at a slower rate than healthy individuals and appear less sensitive to the value of outcomes, as evidenced by lower rates of win-stay responses, shallower learning curves after reversals, results from the linear regression fits, and lower learning rate and value sensitivity parameter values in our computational modeling analysis. These differences held irrespective of whether the task involved learning from rewards or punishments, suggesting that the performance differences in both task conditions arise from similar mechanisms. Furthermore, our modeling approaches help rule out other potential biases that could have influenced subject behavior. While some perseverative, choice, and action biases were apparent at the individual subject level, these did not vary systematically between subject groups.

Our findings are consistent with previous findings showing reduced responsiveness to reward in depressed individuals (Eshel & Roiser, 2010; Henriques & Davidson, 2000; Henriques, Glowacki, & Davidson, 1994; Pizzagalli et al., 2008, 2005; Rupprechter et al., 2018). Previous work has shown that rewarding feedback (e.g., rewarded trials in our reward condition) is treated in a similar way as lack of punishment in aversion learning tasks (e.g., un-punished trials in the punishment condition), both at the level of behavior and in the brain (Erdeniz & Done, 2019; Palminteri et al., 2012). Our detailed analysis lends support to the idea that depression results in a global reduction in sensitivity to positive outcomes, regardless of the context in which the positive outcome is delivered (e.g., even if lack of punishment serves as the positive outcome).

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

However, our results do not support the hypothesis of a global increase in sensitivity to punishment: in fact, depressed individuals showed poorer overall performance in learning to avoid punishment. Poorer performance on punishment learning is consistent with the deficits observed in depressed individuals on the Iowa Gambling Task, as disadvantageous choices on that task are characterized by large punishments (Hegedűs et al., 2018; Must, Horvath, Nemeth, & Janka, 2013). In our study, depressed individuals also did not differ from controls in their level of ‘lose-switch’ responses. Evidence for increased ‘lose-switch’ responses in depressed individuals in learning paradigms has been inconsistent, with some studies showing heightened responses in depressed subject groups (Dombrovski et al., 2013; Murphy et al., 2003) and others not (Chase et al., 2010; Chen et al., 2015; Dombrovski et al., 2010; Gradin et al., 2011). Our results suggest that, while hypersensitivity to negative feedback may be found in cognitive tasks (Beats et al., 1996; Elliot et al., 1996; Eshel & Roiser, 2010), this does not reliably translate to a general hypersensitivity to punishment during learning tasks.

Several previous studies using probabilistic learning have failed to find overall performance differences between depressed and non-depressed individuals (Chase et al., 2010; Dombrovski et al., 2010; Dombrovski et al., 2015, 2013; Moutoussis et al., 2018). We think the most likely explanation for the discrepancy in results is statistical power, as our study has twice the sample size of any of the previous ones (Chase et al., 2010; Dombrovski et al., 2010; Dombrovski et al., 2015, 2013; Moutoussis et al., 2018; Rupprechter et al., 2018), and in several of these previous studies the depressed groups showed similar non-significant trends towards poorer performance and lower learning rates (Chase et al., 2010; Dombrovski et al., 2010). Of note, though, the one previous study that examined probabilistic learning for both rewards and punishments reported similar behavior in depressed individuals and healthy controls

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

(Dombrovski et al., 2010). Their task, however, did not involve any reversals and used an easier probabilistic discrimination (80% vs 20%, compared to 70% vs 30%), and therefore may not have been as sensitive to the learning deficits we observe here.

Our analysis also provides details into the feedback processing mechanisms affected by depression. Learning rates were consistently lower in individuals with MDD in both task conditions, indicating difficulty learning from task feedback in MDD. However, more recent work has proposed that, rather than decreased learning rates, reward processing deficits arise from lower value sensitivity, a factor not often examined in many RL analyses of MDD behavior (Huys et al., 2015; Huys et al., 2013). Consistent with this claim, in addition to learning rate differences, we show that value sensitivity is also reduced in MDD, particularly in punishment learning. However, our out of sample prediction analysis suggests that these sensitivity differences contribute less to observed deficits than learning rates.

We note two important caveats regarding these findings on learning rate and value sensitivity. First, in our task, changes in value sensitivity have the same effect in an RL model as changes in the exploration-exploitation tradeoff. Therefore, the reduced value sensitivity we see in depressed individuals could be interpreted as either a greater tendency towards exploratory choices or a reduced subjective desirability of positive versus negative outcomes. Second, the task conditions in which depressed individuals show lower reward sensitivity rather than lower learning rate (Huys et al., 2013) involve a qualitatively different environment, where participants make perceptual judgements that are implicitly influenced by rewards observed on previous trials (Pizzagalli et al., 2008). It is possible that in explicit learning environments, such as reversal learning tasks, the influence on learning rate becomes more apparent. Indeed, consistent with our findings, a more recent study with a task set up similar to our reward condition also reports lower

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

learning rate and value sensitivity in depressed individuals, suggesting that deficits in both computations may be at play (Rupprechter et al., 2018).

Finally, although not measured directly in the current study, our results suggest neurobiological implications. Reinforcement learning has been linked to dopaminergic signals in the striatum (Daw, Gershman, Seymour, Dayan, & Dolan, 2011; Donoso, Collins, & Koehlin, 2014; McGuire, Nassar, Gold, & Kable, 2014; O'Doherty et al., 2004; Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006). A number of studies have demonstrated that striatal responses to reward, including those received during reversal learning tasks, are attenuated in depressed individuals (Gradin et al., 2011; Kumar et al., 2008; Pizzagalli et al., 2009; Remijnse et al., 2009; Robinson, Cools, Carlisi, Sahakian, & Drevets, 2012). Another study found that reduced ventral striatal responsiveness to unexpected rewards predicted the severity of depression across both unipolar and bipolar depressed groups (Satterthwaite et al., 2015). These results are consistent with the interpretation that MDD blunts responsiveness to positive outcomes, which could subsequently influence how these feedback signals are used for learning. A key priority for future research should be to more closely link these neural and behavioral effects of depression, as well as to explore these effects across a wider range of disorders where blunted value responsiveness is a key component, including across mood disorders (unipolar and bipolar depression) and psychotic disorders that share the common symptom of anhedonia (Gold, Waltz, Prentice, Morris, & Heerey, 2008; Waltz, Frank, Robinson, & Gold, 2007) .

The current results need to be appraised in light of certain limitations. Although we did screen out individuals with bipolar, substance abuse and/or psychotic symptoms, we did not assess other potential co-morbidities. Beyond the screen for substance use disorder, we did not collect detailed information on alcohol use or smoking, which may be associated with learning

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

performance. The MDD group reported depression ranging between moderate and severe severity levels, and were matched in IQ to our healthy controls, so our conclusions are most applicable to depressed individuals with these characteristics. Finally, our study was not designed to evaluate potential effects of therapy or drug treatments.

This last caveat suggests an important area for future research. The current study does not address whether impaired reward and punishment learning is a risk factor for developing depression or a consequence of the disorder or its treatment. Future studies could address this question through longitudinal or treatment studies. Whether individuals who exhibit poorer learning from rewards and punishments are more likely to develop depression, and/or whether these deficits increase with the onset of depressive symptoms and improve in remission, are both critical questions for understanding the association between depression and reward and punishment learning.

References

- Beats, B., Sahakian, B., & Levy, R. (1996). Cognitive performance in tests sensitive to frontal lobe dysfunction in the elderly depressed. *Psychological Medicine*, 26(3), 591–603.
<https://doi.org/10.1017/S0033291700035662>
- Beck, A, Steer, R., & Brown, G. (1996). Beck depression inventory-II. *San Antonio*, 78(2), 490–498.
- Beck, Aaron, Epstein, N., Brown, G., & Steer, R. (1988). An inventory for measuring clinical anxiety: Psychometric properties. *Journal of Consulting and Clinical Psychology*, 56(6), 893–897.
- Betancourt, M., & Girolami, M. (2015). Hamiltonian Monte Carlo for Hierarchical Models. *Current Trends in Bayesian Methodology with Applications*, 79–101.
<https://doi.org/10.1201/b18502-5>
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1). <https://doi.org/10.18637/jss.v076.i01>
- Carver, C. S., & White, T. L. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS Scales. *Journal of Personality and Social Psychology*, 67(2), 319.
- Chase, H. W., Frank, M. J., Michael, A., Bullmore, E. T., Sahakian, B. J., & Robbins, T. W. (2010). Approach and avoidance learning in patients with major depression and healthy controls: Relation to anhedonia. *Psychological Medicine*, 40(3), 433–440.
<https://doi.org/10.1017/S0033291709990468>
- Chen, C., Takahashi, T., Nakagawa, S., Inoue, T., & Kusumi, I. (2015). Reinforcement learning

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

- in depression: A review of computational research. *Neuroscience and Biobehavioral Reviews*, 55, 247–267. <https://doi.org/10.1016/j.neubiorev.2015.05.005>
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204–1215. <https://doi.org/10.1016/j.neuron.2011.02.027>
- DeLong, E. R., DeLong, D., & Clarke-Pearson, D. (1988). Comparing the Areas under Two or More Correlated Receiver Operating Characteristic Curves : A Nonparametric Approach. *Biometrics*, 44(3), 837–845.
- Dombrovski, A., Clark, L., Siegle, G., Butters, M., Ichikawa, N., Sahakian, B., & Szanto, K. (2010). Reward/Punishment Reversal Learning in Older Suicide Attempters. *The American Journal of Psychiatry*, 167(6), 699–707.
- Dombrovski, A. Y., Szanto, K., Clark, L., Aizenstein, H. J., Chase, H. W., Reynolds, C. F., & Siegle, G. J. (2015). Corticostriathalamic Reward Prediction Error Signals and Executive Control in Late-Life Depression. *Psychological Medicine*, 45(7), 1413–1424. <https://doi.org/10.1017/S0033291714002517>
- Dombrovski, A. Y., Szanto, K., Clark, L., Reynolds, C. F., & Siegle, G. J. (2013). Reward signals, attempted suicide, and impulsivity in late-life depression. *JAMA Psychiatry*, 70(10), 1020–1030. <https://doi.org/10.1001/jamapsychiatry.2013.75>
- Donoso, M., Collins, A. G. E., & Koechlin, E. (2014). Foundations of human reasoning in the prefrontal cortex. *Science*, 344(6191), 1481–1486. <https://doi.org/10.1126/science.1252254>
- Elliot, R., Sahakian, B., Herrod, J., Robins, T., & Paykel, E. (1996). Neuropsychological impairments in unipolar depression: The influence of perceived failure on subsequent performance. *Psychological Medicine*, 26(5), 975–989.

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

<https://doi.org/10.1017/S0033291700035303>

Erdeniz, B., & Done, J. (2019). Common and Distinct Functional Brain Networks for Intuitive and Deliberate Decision Making. *Brain Sciences*, 9(7), 174.

<https://doi.org/10.3390/brainsci9070174>

Eshel, N., & Roiser, J. P. (2010). Reward and punishment processing in depression. *Biological Psychiatry*, 68(2), 118–124. <https://doi.org/10.1016/j.biopsych.2010.01.027>

Gelman, A., & Rubin, D. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4), 457–511.

Gold, J. M., Waltz, J. A., Prentice, K. J., Morris, S. E., & Heerey, E. A. (2008). Reward processing in schizophrenia: A deficit in the representation of value. *Schizophrenia Bulletin*, 34(5), 835–847. <https://doi.org/10.1093/schbul/sbn068>

Gradin, V. B., Kumar, P., Waiter, G., Ahearn, T., Stickle, C., Milders, M., ... Steele, J. D. (2011). Expected value and prediction error abnormalities in depression and schizophrenia. *Brain*, 134(6), 1751–1764. <https://doi.org/10.1093/brain/awr059>

Hegedűs, K. M., Szkaliczki, A., Gál, B. I., Andó, B., Janka, Z., & Álmos, P. Z. (2018). Decision-making performance of depressed patients within 72 h following a suicide attempt. *Journal of Affective Disorders*, 235(February), 583–588. <https://doi.org/10.1016/j.jad.2018.04.082>

Henriques, J. B., & Davidson, R. J. (2000). Decreased responsiveness to reward in depression. *Cognition and Emotion*, 14(5), 711–724. <https://doi.org/10.1080/02699930050117684>

Henriques, Jeffrey B., Glowacki, J. M., & Davidson, R. J. (1994). Reward Fails to Alter Response Bias in Depression. *Journal of Abnormal Psychology*, 103(3), 460–466.

<https://doi.org/10.1037/0021-843X.103.3.460>

Huys, Q. J. M., Cools, R., Gölzer, M., Friedel, E., Heinz, A., Dolan, R. J., & Dayan, P. (2011).

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Disentangling the roles of approach, activation and valence in instrumental and pavlovian responding. *PLoS Computational Biology*, 7(4).

<https://doi.org/10.1371/journal.pcbi.1002028>

Huys, Q. J. M., Daw, N. D., & Dayan, P. (2015). Depression: A Decision-Theoretic Analysis.

Annual Review of Neuroscience, 38(1), 1–23. <https://doi.org/10.1146/annurev-neuro-071714-033928>

Huys, Q. J. M., Maia, T. V., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience*, 19(3), 404–413.

<https://doi.org/10.1038/nn.4238>

Huys, Q. J., Pizzagalli, D. A., Bogdan, R., & Dayan, P. (2013). Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. *Biology of Mood & Anxiety Disorders*, 3(1). <https://doi.org/10.1186/2045-5380-3-12>

Kuhn, M. (2008). Introduction. *Journal of Statistical Software*, 28(5), 1–26.

Kumar, P., Waiter, G., Ahearn, T., Milders, M., Reid, I., & Steele, J. D. (2008). Abnormal temporal difference reward-learning signals in major depression. *Brain*, 131(8), 2084–2093. <https://doi.org/10.1093/brain/awn136>

Lau, B., & Glimcher, P. W. (2008). Value Representations in the Primate Striatum during Matching Behavior. *Neuron*, 58(3), 451–463. <https://doi.org/10.1016/j.neuron.2008.02.021>

Lovibond, P., & Lovibond, S. (1995). The structure of negative emotional states: Comparison of the Depression Anxiety Stress Scales (DASS) with the Beck Depression and Anxiety Inventories. *Behaviour Research and Therapy*, 33(3), 335–343.

McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

influences on learning rate in a dynamic environment. *Neuron*, 84(4), 870–881.

<https://doi.org/10.1016/j.neuron.2014.10.013>

Montague, P. R., Dolan, R., Friston, K., & Dayan, P. (2012). Computational Psychiatry. *Trend in Cognitive Sciences*, 16(1), 72–80. <https://doi.org/10.1016/j.tics.2011.11.018> Trends

Moussavi, S., Chatterji, S., Verdes, E., Tandon, A., Patel, V., & Utsun, B. (2007). Depression, chronic diseases, and decrements in health: results from the Wolrd Health Surveys. *The Lancet*, 370(9590), 851–858.

Moutoussis, M., Rutledge, R. B., Prabhu, G., Hrynkiewicz, L., Lam, J., Ousdal, O. T., ... Dolan, R. J. (2018). Neural activity and fundamental learning, motivated by monetary loss and reward, are intact in mild to moderate major depressive disorder. *PLoS ONE*, 13(8), 1–20. <https://doi.org/10.1371/journal.pone.0201451>

Mukherjee, D., Lee, S., Kazinka, R., Satterwaite, T., & Kable, Joseph, W. (2020). The Depressed Decision-Maker:HowValue Based Decision Making Differs in Major Depressive Disorder. *PsyArxiv*. <https://doi.org/10.31234/osf.io/7gnb9>

Murphy, F. C., Michael, A., Robbins, T. W., & Sahakian, B. J. (2003). Neuropsychological impairment in patients with major depressive disorder: The effects of feedback on task performance. *Psychological Medicine*, 33(3), 455–467. <https://doi.org/10.1017/S0033291702007018>

Must, A., Horvath, S., Nemeth, V. L., & Janka, Z. (2013). The Iowa Gambling Task in depression - what have we learned about sub-optimal decision-making strategies? *Frontiers in Psychology*, 4, 1–6. <https://doi.org/10.3389/fpsyg.2013.00732>

O’Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science (New*

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

- York, N.Y.), 304(5669), 452–454. <https://doi.org/10.1126/science.1094285>
- Ottenbreit, N. D., & Dobson, K. S. (2004). Avoidance and depression: the construction of the Cognitive–Behavioral Avoidance Scale. *Behaviour Research and Therapy*, 42(3), 293–313.
- Palminteri, S., Justo, D., Jauffret, C., Pavlicek, B., Dauta, A., Delmaire, C., ... Pessiglione, M. (2012). Critical Roles for Anterior Insula and Dorsal Striatum in Punishment-Based Avoidance Learning. *Neuron*, 76(5), 998–1009. <https://doi.org/10.1016/j.neuron.2012.10.017>
- Pechtel, P., Dutra, S. J., Goetz, E. L., & Pizzagalli, D. A. (2013). Blunted reward responsiveness in remitted depression. *Journal of Psychiatric Research*, 47(12), 1864–1869. <https://doi.org/10.1016/j.jpsychires.2013.08.011>
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, 442(7106), 1042–1045. <https://doi.org/10.1038/nature05051>
- Pizzagalli, D. A., Iosifescu, D., Hallett, L. A., Ratner, K. G., & Fava, M. (2008). Reduced hedonic capacity in major depressive disorder: Evidence from a probabilistic reward task. *Journal of Psychiatric Research*, 43(1), 76–87. <https://doi.org/10.1016/j.jpsychires.2008.03.001>
- Pizzagalli, D. A., Jahn, A. L., & O'Shea, J. P. (2005). Toward an objective characterization of an anhedonic phenotype: A signal-detection approach. *Biological Psychiatry*, 57(4), 319–327. <https://doi.org/10.1016/j.biopsych.2004.11.026>
- Pizzagalli, D., Holmes, A., Dillon, D., Goetz, E., Birk, J., Bogdan, R., ... Fava, M. (2009). Reduced Caudate and Nucleus Accumbens Response to Rewards in Unmedicated Individuals With Major Depressive Disorder. *The American Journal of Psychiatry*, 166(6),

702–710.

Remijnse, P. L., Nielen, M. M. A., Van Balkom, A. J. L. M., Hendriks, G. J., Hoogendijk, W. J., Uylings, H. B. M., & Veltman, D. J. (2009). Differential frontal-striatal and paralimbic activity during reversal learning in major depressive disorder and obsessive-compulsive disorder. *Psychological Medicine*, 39(9), 1503–1518.

<https://doi.org/10.1017/S0033291708005072>

Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., & Mueller, M. (2011). pROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics*, 8, 12–77. <https://doi.org/10.1007/s00134-009-1641-y>

Robinson, O., Cools, R., Carlisi, C., Sahakian, B., & Drevets, W. (2012). Ventral Striatum Response During Reward and Punishment Reversal Learning in Unmedicated Major Depressive Disorder. *The American Journal of Psychiatry*, 169(2), 152–159.

Rosenberg, M. (1965). Rosenberg self-esteem scale (RSE). *Acceptance and Commitment Therapy. Measures Package*, 61(52).

Rupprechter, S., Stankevicius, A., Huys, Q. J. M., Steele, J. D., & Seriès, P. (2018). Major Depression Impairs the Use of Reward Values for Decision-Making. *Scientific Reports*, 8(1), 1–8. <https://doi.org/10.1038/s41598-018-31730-w>

Satterthwaite, T. D., Kable, J. W., Vandekar, L., Katchmar, N., Bassett, D. S., Baldassano, C. F., ... Wolf, D. H. (2015). Common and Dissociable Dysfunction of the Reward System in Bipolar and Unipolar Depression. *Neuropsychopharmacology*, 40(9), 2258–2268. <https://doi.org/10.1038/npp.2015.75>

Snaith, R. P., Hamilton, M., Morley, S., Humayan, A., Hargreaves, D., & Trigwell, P. (1995). A scale for the assessment of hedonic tone. The Snaith-Hamilton Pleasure Scale. *British*

REVERSAL LEARNING IN MAJOR DEPRESSIVE DISORDER

Journal of Psychiatry, 167(1), 99–103.

Steffens, D. C., Wagner, H. R., Levy, R. M., Horn, K. A., & Krishnan, K. R. R. (2001).

Performance feedback deficit in geriatric depression. *Biological Psychiatry*, 50(5), 358–363. [https://doi.org/10.1016/S0006-3223\(01\)01165-9](https://doi.org/10.1016/S0006-3223(01)01165-9)

Sutton, R., & Barto, A. (1998). *Introduction to reinforcement learning*. Cambridge, MA: MIT Press.

Team, S. D. (2018). RStan: the R interface for Stan.

Waltz, J. A., Frank, M. J., Robinson, B. M., & Gold, J. M. (2007). Selective Reinforcement Learning Deficits in Schizophrenia Support Predictions from Computational Models of Striatal-Cortical Dysfunction. *Biological Psychiatry*, 62(7), 756–764.

<https://doi.org/10.1016/j.biopsych.2006.09.042>

Wang, X.-J., & Krystal, J. (2014). Computational Psychiatry. *Neuron*, 84, 638–654.

<https://doi.org/10.1024/1661-4747/a000295>

Watkins, C. J. C. H., & Dayan, P. (1992). Technical Note: Q-Learning. *Machine Learning*, 8(3–4), 279–292. <https://doi.org/10.1007/BF00992698>