


A comparison of ‘pruning’ during multi-step planning in depressed and healthy individuals

Paul Faulkner¹ , Quentin J. M. Huys^{2,3}, Daniel Renz⁴, Neir Eshel⁵, Stephen Pilling⁶, Peter Dayan⁷ and Jonathan P. Roiser⁸

Original Article

Cite this article: Faulkner P, Huys QJM, Renz D, Eshel N, Pilling S, Dayan P, Roiser JP (2021). A comparison of ‘pruning’ during multi-step planning in depressed and healthy individuals. *Psychological Medicine* 1–9. <https://doi.org/10.1017/S0033291721000799>

Received: 14 May 2020

Revised: 23 December 2020

Accepted: 18 February 2021

Key words:

Decision-making; depression; pruning; serotonin; unmedicated

Author for correspondence:

Paul Faulkner,

E-mail: paul.faulkner@roehampton.ac.uk

¹Department of Psychology, University of Roehampton, London, UK; ²Division of Psychiatry, University College London, London, UK; ³Max Planck Centre for Computational Psychiatry and Ageing Research, University College London, London, UK; ⁴Translational Neuromodeling Unit, Institute for Biomedical Engineering, University of Zurich and ETH Zurich, Zurich, Switzerland; ⁵Department of Psychiatry and Behavioral Sciences, Stanford University, Palo Alto, California, USA; ⁶Division of Psychology and Language Sciences, University College London, London, UK; ⁷Department of Computational Neuroscience, Max Planck Institute for Biological Cybernetics, Tübingen, Germany and ⁸Institute of Cognitive Neuroscience, University College London, London, UK

Abstract

Background. Real-life decisions are often complex because they involve making sequential choices that constrain future options. We have previously shown that to render such multi-step decisions manageable, people ‘prune’ (i.e. selectively disregard) branches of decision trees that contain negative outcomes. We have theorized that sub-optimal pruning contributes to depression by promoting an oversampling of branches that result in unsavoury outcomes, which results in a negatively-biased valuation of the world. However, no study has tested this theory in depressed individuals.

Methods. Thirty unmedicated depressed and 31 healthy participants were administered a sequential reinforcement-based decision-making task to determine pruning behaviours, and completed measures of depression and anxiety. Computational, Bayesian and frequentist analyses examined group differences in task performance and relationships between pruning and depressive symptoms.

Results. Consistent with prior findings, participants robustly pruned branches of decision trees that began with large losses, regardless of the potential utility of those branches. However, there was no group difference in pruning behaviours. Further, there was no relationship between pruning and levels of depression/anxiety.

Conclusions. We found no evidence that sub-optimal pruning is evident in depression. Future research could determine whether maladaptive pruning behaviours are observable in specific sub-groups of depressed patients (e.g. in treatment-resistant individuals), or whether misuse of other heuristics may contribute to depression.

Introduction

Major depressive disorder is a leading contributor to disability worldwide, affecting more than 300 million people at any time (World Health Organization, 2017). Although people with depression experience impairments in decision-making that contribute to their disorder (Clark et al., 1992; Eshel & Roiser, 2010; Husain & Roiser, 2018; Pulcu, Thomas, Trotter, & McFarquhar, 2015), our knowledge of maladaptive decision-making processes in depression remains incomplete. Understanding such depression-related impairments may aid our attempts to improve the quality of life for those that struggle with this disorder.

Although people with anhedonic forms of depression exhibit dysfunctional evaluation of immediately-available, yet superficially appetitive outcomes, those with other forms of depression (e.g. those associated with helplessness) may exhibit maladaptive assessment of outcomes that occur only after extended sequences of decisions (Dayan & Huys, 2008; Huys et al., 2015b). Addressing the substantial computational challenge associated with planning during multi-step decision-making lies at the heart of modern reinforcement learning. Such multi-step decision-making requires searching a decision-tree of options in which the first choice constrains future choices. These searches are difficult and necessitate the use of certain heuristics.

Using the multi-step decision-making task in Fig. 1, we investigated various heuristics that healthy individuals use to approximate the best possible overall outcome without exhausting cognitive resources. One particular heuristic we identified is the ‘pruning’ of branches of a decision-tree that contain very negative outcomes in a Pavlovian manner (Huys et al., 2012; Huys, Daw, & Dayan, 2015a). Specifically, healthy individuals selected the optimal sequence of decisions when potential paths did not contain a large loss, although they were dramatically

© The Author(s), 2021. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

CAMBRIDGE
UNIVERSITY PRESS

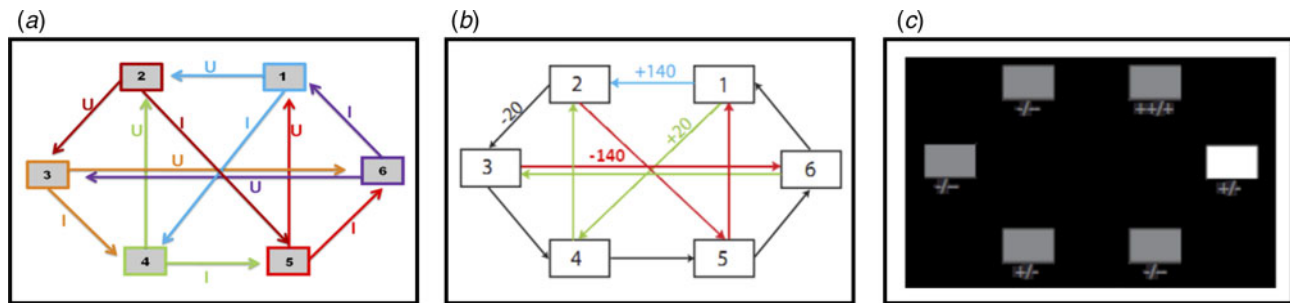


Fig. 1. (a) Deterministic transition matrix presented to participants during the training phase to aid learning. (b) Deterministic reward matrix. Note that this was never presented to participants; they must instead learn the reward structure through trial and error. (c) Task as presented to participants. The white box denotes the state that the participant is currently in. Symbols below each state denote the deterministic reward achieved by transitioning away from that state; ‘+’ = +140 points; ‘+’ = +20 points; ‘-’ = -20 points; ‘-’ = -140 points.

impaired at doing so when the optimal sequence began with a large loss at the first decision step (Fig. 2).

Pruning of decision branches that contain a large loss can reflect the rational allocation of computational resources. Importantly, such pruning may also promote emotional well-being; curtailing a tree search in the face of a large punishment avoids aversive outcomes and the consequent negatively-biased valuation of one’s surroundings (Dayan & Huys, 2008). Conversely, an oversampling of branches associated with unsavoury outcomes (i.e. a consistent selection of branches or options that typically result in punishments) can contribute to a more negatively-biased valuation of one’s surroundings and can thus promote depression (Dayan & Huys, 2008; Huys *et al.*, 2015a). Indeed, research indicates that depression may be associated with an oversampling of aversive information and outcomes (e.g. Garrett *et al.* 2014; Joormann, Hertel, Brozovich, & Gotlib, 2005; Joormann & Gotlib, 2008). For example, Garrett *et al.* (2014) report that when constructing beliefs about the likelihood of future positive and negative events occurring, non-depressed individuals exhibit a bias towards the processing of information that supports a positively-skewed view of future positive events occurring, while ignoring information to the contrary. However, the authors also report that clinically-depressed individuals consider such information as much as they consider information that supports a negatively-skewed view of future events, and as such overevaluate aversive information. In addition, Joormann and Gotlib (2008) report that negatively-valenced but irrelevant words interfere with the ability to update working memory (indexed by response latencies on a recognition task) in depressed participants more than in healthy individuals. These results further indicate that depression may be associated with an overevaluation of aversive information, and that people with depression might sub-optimally prune decision trees by oversampling branches associated with negative outcomes.

Pavlovian pruning of decision trees has been hypothesized to depend on the brain’s serotonin system (Dayan & Huys, 2008). Specifically, reductions in central serotonin are proposed to result in a decrease in Pavlovian behavioural inhibition, and an increased choice of options that result in negative outcomes due to decreases in pruning (Dayan & Huys, 2008). This hypothesis was based on observations that serotonin influences behavioural inhibition in response to losses (Cools, Nakamura, & Daw, 2011; Crockett, Clark, & Robbins, 2009). Because depression is associated with dysfunction in a number of neurotransmitter systems, including the serotonin system (Anderson, 2000; Cannon *et al.*, 2007; Parsey *et al.*, 2003), this theory provides

one possible mechanism by which depressed individuals exhibit negatively-biased valuations of themselves and their surroundings (Clifford and Hemsley, 1987).

The current study aimed to determine whether depression is associated with sub-optimal pruning of decision trees. We compared the performance of depressed and healthy participants on a sequential decision-making task designed to reveal pruning behaviours. Standard, Bayesian and computational analyses were utilized. It was predicted that (1) compared to healthy individuals, depressed individuals would exhibit less pruning of decision-tree branches that begin with a large punishment, and (2) that in depressed individuals, those with the highest levels of depression would demonstrate the lowest pruning.

Methods

Participants

This study employed a between-subjects design. Thirty depressed individuals were recruited via the Camden and Islington NHS Foundation Trust Psychological Therapy Services, while 31 healthy controls were recruited via online and print advertisements. All participants gave written informed consent after receiving a detailed explanation of the study (approved by the London – Queen Square NHS Research Ethics committee). Exclusion criteria [assessed by the Mini International Neuropsychiatric Inventory (MINI); Sheehan *et al.*, 1998] for the healthy control participants included past or present major depressive disorder, bipolar disorder, psychosis, anxiety disorders, substance/alcohol dependence or recent (<6 months) abuse, any neurological disorder and not being a native English speaker. Depressed participants were subject to the same exclusion criteria, except they had to endorse depressive symptoms for a minimum of 10 days within the last 2 weeks (number of past depressive episodes was not important for inclusion), and could also have a diagnosis of an anxiety disorder, or historical substance/alcohol dependence that was restricted to a depressive episode. Exclusion criteria for depressed participants included the use of antidepressants within the last month. All participants provided written informed consent and were compensated £30 for participation, as well as an additional £0–20 depending on task performance.

Procedure

All participants were initially contacted over the telephone (depressed participants were contacted after a referral from the

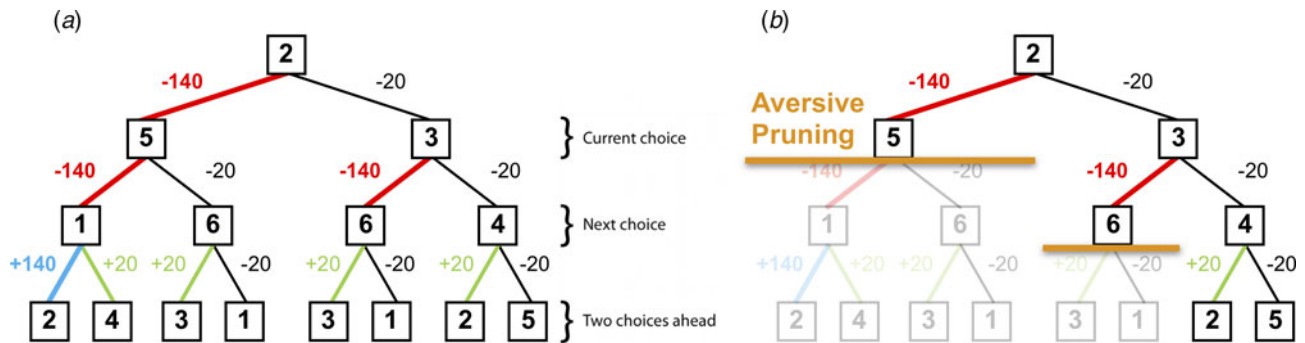


Fig. 2. (a) A typical decision-tree and financial outcomes up to a depth of 3 starting from state 2. Numbers in each box denote the state number. (b) Same decision-tree starting, aversively pruned due to a large negative outcome at the first step. Note that this aversive pruning also avoids the large positive transition, but almost halves the computational load.

Camden and Islington NHS Foundation Trust Psychological Therapy Services, while non-depressed participants responded to the advert). At this point, participants were questioned to determine whether they had ever experienced any symptoms of a psychiatric disorder. Importantly, the specifics of such symptoms or disorders were not discussed at this stage; detailed information pertaining to this was to be obtained in-person via administration of the MINI interview, as instructed by the London – Queen Square Research Ethics Committee. However, individuals who at this stage endorsed previously or currently experiencing such symptoms/disorders were not invited to attend the testing session. The testing session took place at the Institute of Cognitive Neuroscience, University College London. Participants completed one testing session each, in which they initially underwent full screening to determine eligibility (via administration of the MINI), and then completed a battery of questionnaires before being trained on and performing the sequential decision-making task.

Questionnaire measures

The Beck Depression Inventory (BDI; Beck, Steer, & Brown, 1996) was administered to quantify the severity of depression. This self-report questionnaire consists of 21 items, each of which is scored on a 0–3-point scale. A total score of 0–9 indicates minimal depression, while total scores of 10–18, 19–29 and 30–63 indicate ‘mild’, ‘moderate’ and ‘severe’ depression, respectively.

The State/Trait Anxiety Inventory (STAI; Spielberger, Gorsuch, Lushene, Vagg, & Jacobs, 1983) is a 40-item self-rating anxiety measure. Participants score each item as either 1 (‘do not agree at all’), 2 (‘agree somewhat’), 3 (‘agree moderately’) or 4 (‘very much agree’). Scores from questions 1–20 and 21–40 are then separately summed to respectively determine state and trait anxiety.

The Wechsler Test of Adult Reading (WTAR; Wechsler, 2001) was used to quantify verbal IQ. Participants were required to read a list of 50 increasingly-uncommon words, and received one point for each correct pronunciation.

Decision-making task

This task is described in Huys et al. (2012), and is presented in Fig. 1. Participants were required to initially complete a training phase, during which they learned how to transition throughout a matrix of six states by referring to the schematic of the transition

matrix presented in Fig. 1a. Importantly, as in Huys et al. (2012), participants were only allowed to proceed to the task phase after demonstrating successful learning of this transition matrix by completing a test without the aid of the schematic.

The task phase began with a further, shorter, training phase designed to teach participants the deterministic financial outcomes associated with each transition. Participants were not presented with a schematic of the action-reward matrix (see Fig. 1b), but learned via trial and error. To help them, the values of the deterministic financial outcomes associated with each transition out of a state were depicted symbolically below each state (but without being identified with the choice options ‘U’ or ‘I’), both during this final training phase and throughout the entire task; ‘++’ denotes a £1.40 gain, ‘+’ denotes a 20 pence gain, ‘--’ denotes a £1.40 loss, and ‘-’ denotes a 20 pence loss. Upon completion of this final training, participants completed 48 trials of the task, each of varying length (2–8 moves). On each trial, participants began in a random state and were instructed to complete a sequence of transitions of a pre-specified length (2–8 moves) to maximize financial gain. On 50% of the trials, transitions were made immediately after each key press, followed by the presentation of the financial outcome for that transition. On the remaining trials (termed ‘plan-ahead’ trials), participants were instructed to plan ahead the remaining (2–4) moves and complete the full sequence of transitions; the transitions and resultant financial outcomes were only presented after the final key press had been made. A schematic of the decision-tree when starting in state 2 can be seen in Fig. 2a.

Model-based statistical analyses

As in Huys et al. (2012) the first 24 of the 48 trials were considered an extension of the reward matrix training, and data from these trials were not analysed. However, we also performed an analysis of data from all 48 trials; whilst doing so did not change the overall significance (or lack thereof) of our findings, these are reported in the online Supplementary Materials. A set of eight increasingly complex models was fit to the data using a Bayesian model comparison approach; these models are fully defined in Huys et al. (2012). Briefly, each successive model had an extra parameter to explain the data, and was assessed according to its Bayesian Information Criterion (BIC_{int}) which is based on the likelihood that the model can accurately explain the data and penalizes the model for its extra complexity.

The first model was a simple ‘Lookahead’ model, which assumed that subjects evaluate the entire decision-tree to choose the optimal sequence of transitions. That is, for a trial of length d , this model assumed that participants considered all 2^d possible sequences of transitions and chose the financially most beneficial sequence. Specifically, the Q -value of each action (a) in the present state (s) was given by the sum of the immediate reward $R(a,s)$ and the value of the optimal action from the next state $s' = T(a,s)$:

$$Q_{\text{Lookahead}}(a, s) = R(a, s) \max_{a'} Q_{\text{Lookahead}}(a', T(a, s))$$

This equation is iterated until the end of the tree has been reached. Because a search of the entire tree is unlikely for depths >3 due to the large computational demands, more complex models were fit to the data which allowed for testing of hypotheses pertaining to pruning. The first of these, termed the ‘Discount’ model, built upon the previous model by including a ‘discount’ factor that assumed participants exhibit a general tendency to fail to evaluate all 2^d (i.e. up to $2^8 = 256$) sequences. Specifically, this model assumed that evaluation stopped at any stage along a path with probability γ . This is equivalent to the standard model of exponential discounting in economics, and implies that the deeper into the tree an outcome, the less likely it is to be included in the calculation:

$$Q_{\text{Discount}}(a, s) = R(a, s) + (1 - \gamma) \max_{a'} Q_d(a', T(a, s))$$

where, at each step, the next transition is weighted by the probability $(1-\gamma)$ that is encountered.

The next model, termed the ‘Pruning’ model, is central to this study’s hypotheses. This model splits the above γ parameter that quantifies the tendency to stop a tree-search into two separate parameters. The first of these is a ‘general pruning’ parameter (γ_G) that quantifies the proportion of trials on which tree searches were curtailed due to a general failure to look ahead (identical to the ‘discount’ parameter in the previous model). The second is termed a ‘specific pruning’ parameter (γ_S) that quantifies the proportion of trials on which the tree search was stopped specifically when the next transition would incur a large loss (see Fig. 2b for an example):

$$Q_{\text{Pruning}}(a, s) = R(a, s) + (1 - x) \max_{a'} Q_p(a', T(a, s))$$

$$x = \begin{cases} \gamma_S & \text{if } R(a, s) \text{ is the large negative reinforcement} \\ \gamma_G & \text{else} \end{cases}$$

The next model, termed the ‘Pruning and Pavlovian’ model, accounted for ‘Learned Pavlovian’ attraction/repulsion to states that are associated with future financial rewards that are not achievable because too few moves remain to obtain them. Specifically, it accounted for such learning due to the addition of a second state-action value which depends on the long-term average value of the states, which is itself learned by standard temporal difference learning after multiple exposures: ←

$$Q_{\text{Pruning and Pavlovian}}(a, s) = Q_{\text{Pruning}}(a, s) + \omega V(T(a, s))$$

where V is the value that is learned by standard temporal

difference learning:

$$V(s) \leftarrow V(s) + \varepsilon(V(s') + r_t - V(s))$$

where $V(s')$ is set to zero at the final transition.

Finally, to distinguish the effect of pruning from the effect of loss aversion (i.e. the notion that a loss of a given amount is more aversive than a reward of the same amount is appetitive), we replicated the above four models but relaxed how they treated the different outcomes. In the original model, preferences for financial outcomes were assumed to be proportional (e.g. a loss of £1.40 was assumed to exactly cancel out a gain of £1.40). In the new models, we fitted separate parameters for each of the financial outcomes, so that individuals could weight outcomes in a non-proportional manner. These four models are termed ‘rho’ (i.e. ‘Lookahead rho’, ‘Discount rho’, ‘Pruning rho’ and ‘Pruning and Pavlovian rho’), with each having three additional parameters (a parameter for each outcome, but no overall scaling parameter as in the original models). In principle, this allowed participants to be attracted to a reward and repelled from a loss, and *vice versa*. If pruning is observable above and beyond an individual’s simple preferences for rewards and losses, the differential sensitivities to rewards and punishments cannot, by themselves, account for the pruning effects in the above four (i.e. non-‘rho’) models.

Group comparisons and psychometric correlation analyses

Once the best-fitting model was identified, its parameter estimates were extracted and compared between groups. Frequentist analyses were performed using the Statistical Package for Social Scientists version 26 (SPSS Inc., Chicago, Illinois, USA). Bayesian analyses were also performed using JASP [JASP Team (2019), version 0.11.1] because they provide Bayes Factors, which depict a ratio of the probability of the evidence for one hypothesis (i.e. the null) relative to another (i.e. the experimental). Comparing evidence in this way allows one to demonstrate support for the null hypothesis, as opposed to simply failing to reject the null as when using frequentist approaches (Wetzels et al., 2011).

To determine the effects of depression status on parameter estimates from the most parsimonious model, frequentist and Bayesian independent samples t tests were performed, with the relevant parameter estimate added as the dependent variable. To examine relationships between parameter estimates and psychometric questionnaire data, frequentist and Bayesian bivariate correlation analyses were performed.

For the frequentist analyses, a significance threshold of $\alpha = 0.05$ (two-tailed) was adopted. For the Bayesian analyses, on the basis of Jeffreys (1961), we considered Bayes Factors (BF_{10} ; NB: not logarithmically transformed) smaller than 1/100 to be extreme evidence for the null hypothesis, a BF_{10} between 1/100 and 1/30 to be very strong evidence for the null, a BF_{10} between 1/30 and 1/10 to be strong evidence for the null, a BF_{10} between 1/10 and 1/3 to be moderate evidence for the null, and a BF_{10} between 1/3 and 1 to be not worth more than a bare mention. Conversely, we considered BF_{10} larger than 100 to be extreme evidence for the experimental hypothesis, a BF_{10} between 100 and 30 to be very strong evidence for the experimental hypothesis, a BF_{10} between 30 and 10 to be strong evidence for the experimental hypothesis, a BF_{10} between 10 and 3 to be moderate evidence

for the experimental hypothesis, and a BF_{10} between 3 and 1 to be not worth more than a bare mention.

Results

Participant characteristics

Compared to healthy controls, depressed participants self-reported higher depression scores on the BDI [$t(59) = 13.779$; $p < 0.001$, $BF_{10} = 1.11 \times 10^{+16}$], greater trait anxiety on the STAI [$t(59) = -7.601$; $p < 0.001$, $BF_{10} = 1.672 \times 10^{+7}$] and greater state anxiety on the STAI [$t(59) = -5.219$; $p < 0.001$, $BF_{10} = 4987.46$]. Frequentist t tests failed to reject the null hypothesis that depressed and healthy individuals did not differ in terms of age, years of education or IQ, although Bayesian analyses failed to provide support for or against the null hypothesis that these two groups did not differ in these characteristics; [age: $t(59) = 1.227$; $p = 0.225$, $BF_{10} = 0.496$; years of education: $t(59) = 1.491$; $p = 0.141$, $BF_{10} = 0.660$; IQ: $t(59) = 1.780$; $p = 0.083$, $BF_{10} = 1.058$]. A full description of participant characteristics is presented in Table 1.

Computational analyses

The ability of all eight models to explain participants' choices can be seen in Fig. 3. The inclusion of each extra parameter improved the predictive probabilities of the models, while the four models that incorporated the 'rho' parameter were all able to predict a higher proportion of participants' choices than the corresponding models that did not include this parameter (Fig. 3a). Importantly, the Pruning 'rho' model outperformed all others due to it achieving the lowest BIC_{int} score (Fig. 3b). As expected, the most complex model, the Pruning and Pavlovian 'rho' model, achieves the highest predictive probability (i.e. it is able to accurately predict the highest proportion of participants' choices, as shown in Fig. 3a). However, it is penalized for its added complexity (Fig. 3b). This means that the Pruning 'rho' model is considered the most parsimonious, and therefore the winning model.

Specifically, the winning Pruning 'rho' model included both γ_S and γ_G parameters, which suggests that loss-specific pruning had a robust influence on behaviour. The fraction of choices that was correctly predicted by this winning model can be seen in Fig. 4a–c. Because this Pruning 'rho' model outperforms all others, the γ_S and γ_G parameter values, as well as the reward sensitivities to each of the four transition types, were extracted from this model and compared between depressed and healthy participants.

Group comparisons

Frequentist and Bayesian analyses indicated that depressed and healthy individuals did not differ in terms of total money won on the task [$t(59) = 0.345$; $p = 0.731$, $BF_{10} = 0.274$]. Further, depressed and healthy individuals did not differ in terms of γ_S [$t(59) = 0.320$; $p = 0.750$, $BF_{10} = 0.272$]. However, while the frequentist analysis indicated that depressed and healthy individuals did not differ in terms of γ_G either, the Bayesian analysis failed to provide support for or against there being a group difference in this variable: [$t(59) = -1.123$; $p = 0.266$, $BF_{10} = 0.442$]. These data can be seen in Fig. 4d. There were no differences between the groups in terms of reward sensitivities for the +140 transitions [$t(59) = -0.105$; $p = 0.917$, $BF_{10} = 0.262$], +20 transitions [$t(59) = -0.073$; $p = 0.942$, $BF_{10} = 0.261$] or the -20 transitions [$t(59) =$

Table 1. Participant characteristics and parameter estimates from the winning pruning 'rho' model

Participant characteristics	Healthy individuals	Depressed individuals
<i>N</i>	31	30
Gender (M/F)	15/16	15/15
Age (years)	30.42 (10.80)	33.82 (10.44)
Education (years)	14.38 (2.94)	15.43 (2.51)
IQ	108.00 (14.79)	114.55 (7.37)
BDI	3.00 (4.61)	25.44 (7.48)
Trait anxiety	14.71 (8.16)	32.88 (9.74)
State anxiety	10.19 (7.53)	23.44 (11.39)
Number of depressive episodes	–	15.56 (13.42)
Current episode length (months)	–	3.64 (4.91)
No. of days depressed within past 2 weeks	–	10.63 (2.86)
Number of suicide attempts	–	0.55 (1.21)
Parameter estimates		
γ_S	0.630 (0.14)	0.617 (0.19)
γ_G	0.123 (0.04)	0.136 (0.05)
Rho -140	-2.600 (1.12)	-2.378 (0.97)
Rho -20	-0.006 (0.05)	0.011 (0.05)
Rho +20	1.510 (0.78)	1.520 (0.79)
Rho +140	7.252 (3.48)	7.340 (3.05)

All values expressed as means (s.d.), except for no. of subjects (*N*) and gender.

0.459; $p = 0.648$, $BF_{10} = 0.285$]. Again, while a frequentist analysis indicated that depressed and healthy individuals did not differ in their sensitivity to the -140 transitions, the results of the Bayesian analysis narrowly missed out on providing support for the null hypothesis that these two groups did not differ in this variable [$t(59) = -0.831$; $p = 0.409$, $BF_{10} = 0.348$]. These data can be seen in Fig. 4e.

In depressed participants, BDI scores did not correlate with γ_S ($r = 0.068$; $p = 0.748$, $BF_{10} = 0.261$), γ_G ($r = -0.083$; $p = 0.693$, $BF_{10} = 0.267$), or reward sensitivities to the +140 transitions ($r = 0.038$; $p = 0.857$, $BF_{10} = 0.252$) or the -140 transitions ($r = -0.137$; $p = 0.514$, $BF_{10} = 0.304$). Further, frequentist analyses indicated that BDI scores also did not correlate with the +20 transitions or -20 transitions, although the Bayesian analyses did not provide support for or against the null hypothesis that BDI scores did not correlate with the sensitivity to these two transition types (+20 transitions: $r = -0.226$; $p = 0.276$, $BF_{10} = 0.435$; -20 transitions: $r = 0.261$; $p = 0.207$, $BF_{10} = 0.528$).

In addition, in depressed participants, trait anxiety scores did not correlate with γ_S ($r = 0.083$; $p = 0.694$, $BF_{10} = 0.267$), γ_G ($r = 0.084$; $p = 0.688$, $BF_{10} = 0.268$), or reward sensitivities to the +140 transitions ($r = 0.156$; $p = 0.457$, $BF_{10} = 0.323$), the +20 transitions ($r = 0.309$; $p = 0.133$, $BF_{10} = 0.724$), the -20 transitions ($r = 0.066$; $p = 0.755$, $BF_{10} = 0.260$) or -140 transitions ($r = -0.070$; $p = 0.740$, $BF_{10} = 0.261$).

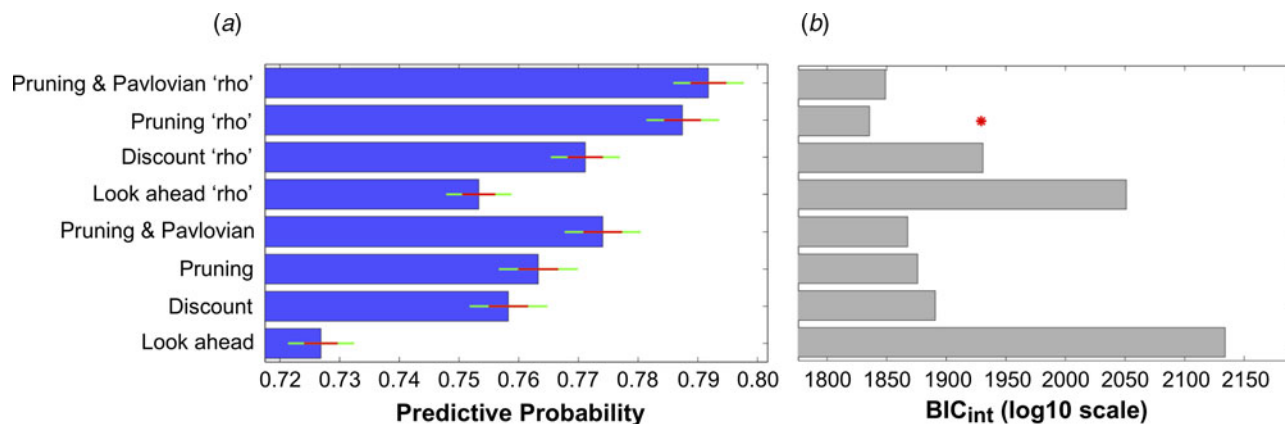


Fig. 3. (a) Mean predictive probabilities for all models. All models that include the 'rho' parameter fit the data better than the corresponding models that do not contain this parameter. (b) Model comparisons using each model's Bayesian Information Criterion (BIC_{int}). Despite the fact that the model that predicts the highest proportion of participants' choices is the 'Pruning and Pavlovian' model that contains the 'rho' parameter, this model is penalized due to its added complexity. The most parsimonious (i.e. 'winning') model is therefore the 'Pruning' model that includes the extra 'rho' parameter.

Finally, both frequentist and Bayesian analyses indicated that state anxiety scores did not correlate with γ_S ($r = -0.092$; $p = 0.662$, $BF_{10} = 0.272$) in depressed participants. However, while the frequentist analysis indicated that state anxiety did not correlate with γ_G , results of the Bayesian analysis narrowly missed out on providing support for there being no relationship between these two variables ($r = -0.177$; $p = 0.397$, $BF_{10} = 0.349$). Both frequentist and Bayesian analyses indicated that state anxiety was not related to sensitivities to the +140 transitions ($r = -0.048$; $p = 0.820$, $BF_{10} = 0.254$), +20 transitions ($r = -0.013$; $p = 0.950$, $BF_{10} = 0.249$), -20 transitions ($r = -0.004$; $BF_{10} = 0.248$, $p = 0.985$) or -140 transitions ($r = 0.068$; $p = 0.745$, $BF_{10} = 0.261$).

Discussion

We previously showed that healthy individuals prune decision trees to render complex sequential decisions manageable. Here, we tested the hypothesis that depressed individuals are less able to optimally use this heuristic, and that this inability may be related to their severity of depression. However, contrary to predictions, depressed and healthy individuals in this study did not differ in their pruning behaviours.

The finding that participants pruned branches of decision trees that began with a large loss, regardless of the potential utility of that branch, replicates our previous work (Huys et al., 2012, 2015b). The most parsimonious model in the current study was the same as the 'winning' model in Lally et al. (2017). However, the most parsimonious model in our original study (Huys et al., 2012) included a Pavlovian parameter which indicated that participants had a reflexive attraction/aversion to certain states, whereas including this Pavlovian parameter weakened model parsimony in the current study, meaning that participants did not display strong attractions/aversions to specific states in this dataset.

The addition of four separate parameters to the model (one for each of the four financial outcomes) allowed for the possibility that there was a difference between how participants weighted gains and losses. This is important to avoid confusing pruning with loss aversion. Replicating the findings of Huys et al. (2012), we found that these 'rho' parameters indeed improved the predictive probability of the models. However, as for the healthy individuals in Huys et al. (2012), the current participants did not exhibit typical loss aversion. Instead, the large gain (+140)

was roughly three times more appetitive than the large loss (-140) was aversive. Interestingly, sensitivity to -140 transition types was found to be significantly weaker, while sensitivity to the +140 transitions was significantly greater, in the 'winning' (Pruning 'rho') model than in the model that did not quantify participants' pruning behaviours (i.e. the 'lookahead' model; see online Supplementary Materials). Taken together, these results suggest that loss aversion certainly fails to explain away pruning, although pruning may uncover a form of risk seeking.

Importantly, the current findings failed to support our hypotheses that (a) depressed individuals sub-optimally prune decision trees, and that (b) in depressed individuals, those with the highest levels of depression would demonstrate the lowest pruning. Our latter hypothesis was based partly on the results of Huys et al. (2012), which indicated a relationship between pruning and level of depression. However, it must be noted that the specific relationship reported in Huys et al. (2012) was a significant *positive* correlation between specific pruning and *sub-clinical* depression scores on the BDI in *healthy* participants, not depressed individuals. Further, this relationship was also not replicated by our subsequent studies (i.e. Huys et al., 2015b). While it currently appears that no relationship between pruning and magnitude of depression exists in depressed individuals who share similar characteristics to those who participated in the current study, future research should attempt to determine the replicability of the initial finding of a relationship between pruning and sub-clinical depression in healthy individuals. Further, it is unlikely that we failed to observe sub-optimal pruning in the current depressed individuals because the current healthy participants also pruned sub-optimally, because the latter demonstrated specific pruning to a very similar magnitude as the healthy participants included in Lally et al. (2017) (pruning parameter estimate = ~ 0.6), and they actually pruned slightly more than those in Huys et al. (2012). While this does not explain why the current results do not support the theory put forward by Dayan and Huys (2008), there may be a number of factors that do.

First, levels of depression in our participants may not have been great enough to promote sub-optimal pruning. While the current depressed participants exhibited similar levels of depression to those in studies that report maladaptive decision-making in depression (Joormann & Gotlib, 2008; Kumar et al., 2018; McFarland & Klein, 2009; Ubl et al., 2015), they were all

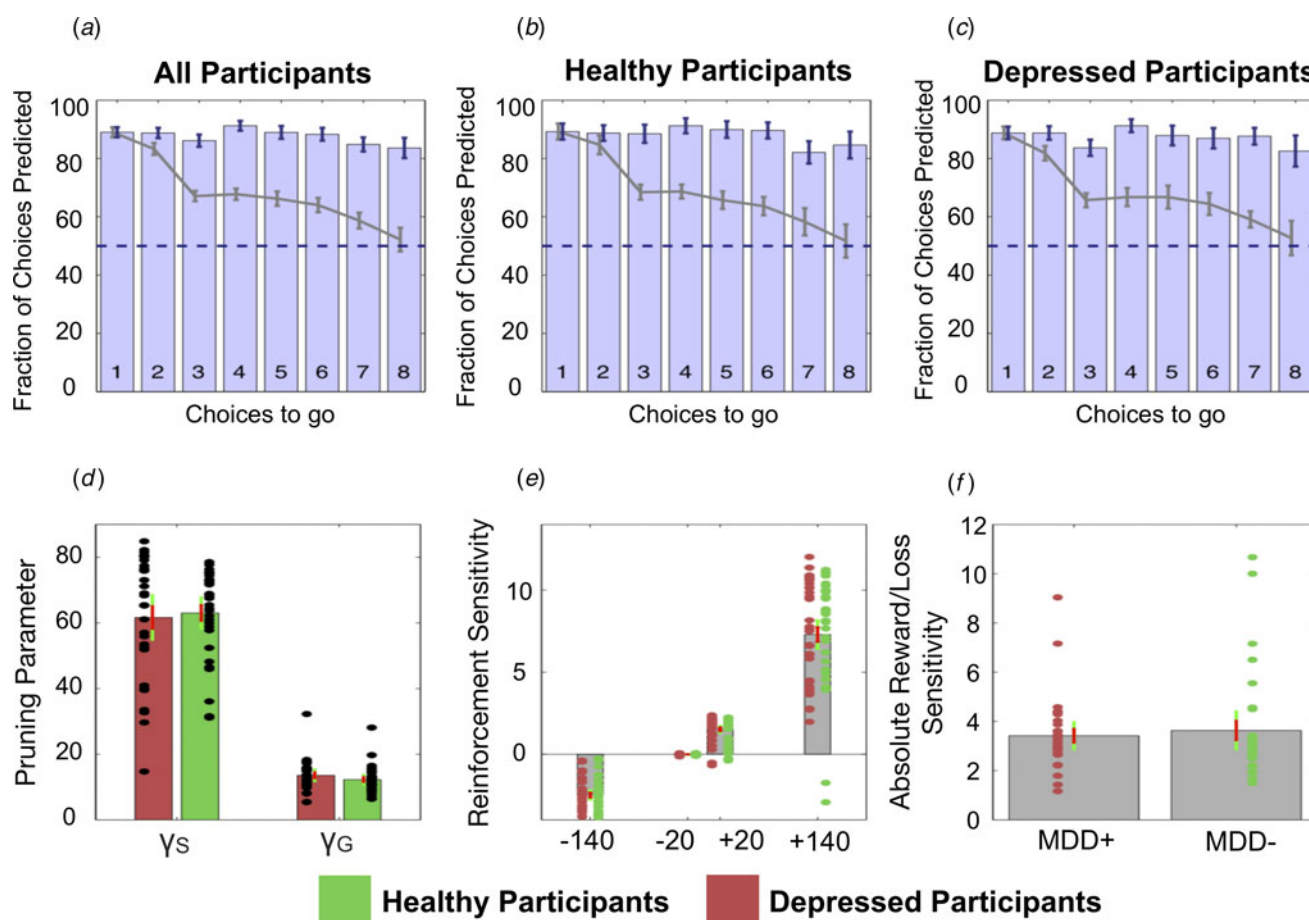


Fig. 4. Top: The fraction of choices correctly predicted by the best-fitting model (the Pruning 'rho' model). (a) All participants combined. (b) Healthy participants only. (c) Depressed participants only. Each bar depicts this as a function of the number of choices remaining on each trial. For example, the right most bar (i.e. bar '8') depicts the fraction of choices at a depth of 1 on eight-choice trials that were correctly predicted by this model; the third rightmost bar (i.e. bar '6') depicts both the fraction of choices that were correctly predicted by this model at (1) a depth of 1 on six-choice trials, (2) a depth of 2 on seven-choice trials and (3) a depth of 3 on eight-choice trials, and so on. Grey lines depict the full 'Lookahead' model. The blue dashed lines depict chance (i.e. 50%). The winning model correctly predicts choices of both depressed participants and healthy controls to roughly the same extent. Further, the full Lookahead model is only able to correctly predict decisions that are eight choices away in the sequence on roughly 50% of trials (i.e. at chance level). The winning model correctly predicts all choices to roughly the same extent, no matter how many choices are remaining. Note that these models include data from, and disregard differences between, trials in which transitions were displayed immediately after each button press and trials in which participants had to enter the entire sequence of transitions at once (i.e. so-called 'plan-ahead' trials.). Bottom: Parameters of the winning Pruning 'rho' model. (d) Specific and general pruning parameters. (e) Reinforcement sensitivity to each transition type. (f) Absolute ratio of reward (+140) to loss (−140) sensitivity. Red denotes depressed participants, green denotes healthy participants. Error bars denote 1 standard deviation above/below the mean (red) and 95% confidence intervals (green).

undergoing 'low-intensity treatments' due to being deemed to exhibit mild/intermediate levels of depression. Indeed, this may be one reason as to why our depressed patients failed to exhibit greater sensitivity to losses than healthy participants, or indeed loss aversion at all, as highly depressed patients (Baek et al., 2017; with mean BDI = 30.10) exhibit heightened aversion to losses, while patients with lower levels of depression do not (Charpentier, Aylward, Roiser, & Robinson, 2017; mean BDI = 16.96). Future studies could therefore examine the pruning behaviours of more severely depressed participants.

Second, Pavlovian pruning of decision trees has been hypothesized to depend on the brain's serotonin system (Dayan & Huys, 2008). Specifically, reductions in central serotonin are proposed to result in a decrease in Pavlovian behavioural inhibition, increasing the choosing of options that result in negative outcomes due to decreases in pruning (Dayan & Huys, 2008). Because serotonin is considered to influence behavioural inhibition in response to losses (Cools et al., 2011; Crockett et al., 2009), and because

depression is associated with dysfunction in a number of neurotransmitter systems including the serotonin system (Anderson, 2000; Cannon et al., 2007; Parsey et al., 2003), it has been argued that depressed participants may exhibit maladaptive pruning due to serotonergic dysfunction (e.g. Dayan & Huys, 2008; Huys et al. 2015a). However, there is no way of knowing whether the current depressed participants exhibit serotonergic dysfunction because no neurochemical measures were collected. Further, depression has been associated with altered dopaminergic functioning (Nestler & Carlezon, 2006), and dopamine is thought to influence decision-making in an opponent fashion to serotonin (Boureau & Dayan, 2010; Daw, Kakade, & Dayan, 2002; Dayan & Huys, 2009). For example, reductions in serotonin via acute tryptophan depletion can enhance the motivational influence of aversive stimuli on instrumental responding, while reductions in dopamine can diminish the influence of appetitive stimuli on such responding (Hebart & Gläscher, 2014). To truly determine the influence of serotonergic function on decision-tree pruning, future studies

may consider examining the effects of a serotonin challenge such as acute tryptophan depletion on performance on this task.

Third, while no group difference was observed in task performance, this does not mean that there are no group differences in the brain regions that are recruited during aversive pruning. We have recently shown that aversive pruning recruits the pregenual anterior cingulate cortex and subgenual anterior cingulate cortex (sgACC) (Lally et al., 2017). Interestingly, the sgACC is overactive in depression (Drevets, Savitz, & Trimble, 2008; Drevets et al., 1997; Mayberg et al., 1997), while the degree of sgACC reactivity to negatively-valenced stimuli can predict treatment response in depressed patients (Roiser, Elliott, & Sahakian, 2012). In addition, depression-related reductions in serotonin 1A receptor availability are greatest in the sgACC (Moses-Kolko et al., 2008). Therefore, comparing the neural mechanisms of pruning in depressed and healthy individuals may reveal group differences in pruning-related sgACC function.

While the above factors may explain why the current results do not support the theory put forward by Dayan and Huys (2008), the fact remains that on the basis of the current findings, it may simply be that depressed individuals may not experience or exhibit maladaptive pruning behaviours during multi-step planning, and that sub-optimal pruning may not promote depressive behaviours. However, pruning is only one heuristic that people use to render complex, sequential decisions manageable. Specifically, healthy participants solve problems by fragmenting deep sequences (i.e. sequences with a depth >3) into sub-sequences of shorter lengths (termed 'fragmenting'), and recall and re-use previous fragmented solutions on subsequent trials (termed 'memoization'), rather than always searching the tree anew each time (Huys et al., 2015b). Unlike the case for pruning (Dayan & Huys, 2008; Huys et al., 2012), neither theoretical nor empirical studies have suggested that altered use of fragmenting or memoization during planning may promote depression. However, future studies might compare the use of these two heuristics, along with the use of pruning, in highly-depressed and healthy individuals.

Our study has several limitations. First, the sample size was relatively limited, which reduced our ability to detect small effects, or to determine the effects of individual differences on task performance; this could be particularly pertinent as Huys et al. (2015b) report that individuals use certain heuristics to solve planning problems in an idiosyncratic fashion. Another potential limitation is the fact that our healthy and depressed individuals may have differed in age, years of education and IQ. While frequentist analyses found no evidence that these two groups differed in such characteristics, Bayesian analyses provided, at best, only anecdotal evidence that these two groups did not differ in this way. Whilst there were no relationships between any of these characteristics and task performance (see online Supplementary Materials), it is therefore possible, although unlikely, that our sampling method may have introduced extraneous variables into our dataset that influenced findings. A further potential limitation is that we did not investigate whether pruning abilities differed as a function of specific clinical depression subtype, such as anhedonic *v.* non-anhedonic forms of depression. However, we did quantify self-reported anhedonia using one item of the BDI, and examined whether it related to pruning behaviours (see online Supplementary Materials). Relationships between anhedonia and sensitivity to each of the +140, +20 and -140 transition types were observed, suggesting that patients with higher anhedonia were less sensitive to both rewards and

punishments, although these results remain preliminary due to our sample size. However, no relationships with pruning were detected. Future studies should determine the association between specific symptoms of depression and pruning behaviours in larger samples of depressed participants. Further, while we have reported that pruning is insensitive to the magnitude of the large loss (Huys et al., 2012), pruning becomes mathematically more disadvantageous as the magnitude of the large loss decreases (relative to the magnitude of the large reward). However, the current data do not indicate whether depressed individuals prune sub-optimally when pruning is more disadvantageous (i.e. when the large loss costs 70 points rather than 140), or whether over-pruning in these circumstances can promote depression. Finally, our calculated reward sensitivities in Fig. 4e indicate that, once pruning is taken into account, participants are three times more 'sensitive' to the large reward than to the large loss. However, pruning and risk seeking are rather entangled in the current task, and it would be interesting to combine it with a compatible, but independent measure of risk seeking.

In summary, we replicated previous findings that people prune decision trees to solve complex planning problems. However, we failed to provide support for the hypothesis that depressed individuals prune sub-optimally. Future research is needed to achieve a more complete understanding of whether misuse of certain heuristics in sequential decision-making can contribute to depression.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/S0033291721000799>.

Financial support. This work was supported by a Medical Research Council Ph.D. Studentship awarded to Dr Paul Faulkner. Professor Peter Dayan was funded by the Gatsby Charitable Foundation, the Max Planck Society and the Alexander von Humboldt Foundation. The funders had no role in the study design, data collection or analysis.

Conflict of interest. None.

Ethical standards. The authors assert that all procedures contributing to this work comply with the ethical standards of the relevant national and institutional committees on human experimentation and with the Helsinki Declaration of 1975, as revised in 2008.

References

- Anderson, I. M. (2000). Selective serotonin reuptake inhibitors versus tricyclic antidepressants: A meta-analysis of efficacy and tolerability. *Journal of Affective Disorders*, 58, 19–36.
- Baek, K., Kwon, J. H., Chae, J.-H., Chung, Y. A., Kralik, J. D., Min, J.-A., ... Jeong, J. (2017). Heightened loss aversion to risk and loss in depressed patients with a suicide attempt history. *Scientific Reports*, 7, 11228–11241.
- Beck, A. T., Steer, R., & Brown, G. (1996). *Manual for the Beck Depression Inventory-II*. San Antonio, TX: Psychological Corporation.
- Boureau, Y.-L., & Dayan, P. (2010). Opponency revisited: Competition and cooperation between dopamine and serotonin. *Neuropsychopharmacology*, 36, 74–97.
- Cannon, D. M., Ichise, M., Rollis, D., Klaver, J. M., Gandhi, S. K., Charney, D. S., ... Drevets, W. C. (2007). Elevated serotonin transporter binding in major depressive disorder assessed using positron emission tomography and [¹¹C]DASB; comparison with bipolar disorder. *Biological Psychiatry*, 62(8), 870–877.
- Charpentier, C. J., Aylward, J., Roiser, J. P., & Robinson, O. J. (2017). Enhanced risk aversion, but not loss aversion, in unmedicated pathological anxiety. *Biological Psychiatry*, 81(12), 1014–1022.
- Clark, G., Hops, H., Lewinsohn, P. M., Andrews, J., Seeley, J. R., & Williams, J. (1992). Cognitive-behavioural group treatment of adolescent depression: Prediction of outcome. *Behavior Therapy*, 23(3), 341–354.

- Clifford, P. I., & Hemsley, D. R. (1987). The influence of depression on the processing of personal attributes. *The British Journal of Psychiatry*, *150*(1), 98–103.
- Cools, R., Nakamura, K., & Daw, N. D. (2011). Serotonin and dopamine: Unifying affective, activational, and decision functions. *Neuropsychopharmacology*, *36*, 98–113.
- Crockett, M. J., Clark, L., & Robbins, T. W. (2009). Reconciling the role of serotonin in behavioural inhibition and aversion: Acute tryptophan depletion abolishes punishment-induced inhibition in humans. *Journal of Neuroscience*, *29*(38), 11993–11999.
- Daw, N. D., Kakade, S., & Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural Networks*, *15*(4–6), 603–616.
- Dayan, P., & Huys, Q. J. M. (2008). Serotonin, inhibition and negative mood. *PLoS Computational Biology*, *4*(2), 1–11.
- Dayan, P., & Huys, Q. J. M. (2009). Serotonin in affective control. *Annual Review of Neuroscience*, *32*, 95–126.
- Drevets, W. C., Price, J. L., Simpson, J. R., Jr., Todd, R. D., Reich, T., Vannier, M., & Raichle, M. E. (1997). Subgenual prefrontal cortex abnormalities in mood disorders. *Nature*, *386*, 824–827.
- Drevets, W. C., Savitz, J., & Trimble, M. (2008). The subgenual anterior cingulate cortex in mood disorders. *CNS Spectrums*, *13*, 663–681.
- Eshel, N., & Roiser, J. P. (2010). Reward and punishment processing in depression. *Biological Psychiatry*, *68*(2), 118–124.
- Garrett, N., Sharot, T., Faulkner, P., Korn, C. W., Roiser, J. P., & Dolan, R. J. (2014). Losing the rose tinted glasses: Neural substrates of unbiased belief updating in depression. *Frontiers in Human Neuroscience*, *8*(639), 1–9.
- Hebart, M. N., & Gläscher, J. (2014). Serotonin and dopamine differentially affect appetitive and aversive general Pavlovian-to-instrumental transfer. *Psychopharmacology*, *232*(2), 437–451.
- Husain, M., & Roiser, J. P. (2018). Neuroscience of apathy and anhedonia: A transdiagnostic approach. *Nature Reviews Neuroscience*, *19*(8), 470–484.
- Huys, Q. J. M., Daw, N. D., & Dayan, P. (2015a). Depression: A decision-theoretic analysis. *Annual Reviews of Neuroscience*, *38*, 1–23.
- Huys, Q. J. M., Eshel, N., O’Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: How the Pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, *8*(3), 1–13.
- Huys, Q. J. M., Lally, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., ... Roiser, J. P. (2015b). Interplay of approximate planning strategies. *PNAS*, *112*(10), 3098–3103.
- Jeffreys, H. (1961). *Theory of probability* (3rd ed.). Oxford: Oxford University Press, Clarendon Press.
- Joormann, J., & Gotlib, I. H. (2008). Updating the contents of working memory in depression: Interference from irrelevant negative material. *Journal of Abnormal Psychology*, *117*(1), 182–192.
- Joormann, J., Hertel, P. T., Brozovich, F., & Gotlib, I. H. (2005). Remembering the good, forgetting the bad: Intentional forgetting of emotional material in depression. *Journal of Abnormal Psychology*, *114*(4), 640–648.
- Kumar, P., Goer, F., Murray, L., Dillon, D. G., Beltzer, M. L., Cohen, A. L., ... Pizzagalli, D. A. (2018). Impaired reward prediction error encoding and striatal-midbrain connectivity in depression. *Neuropsychopharmacology*, *43*, 1581–1588.
- Lally, N., Huys, Q. J. M., Eshel, N., Faulkner, P., Dayan, P., & Roiser, J. P. (2017). The neural basis of aversive Pavlovian guidance during planning. *Journal of Neuroscience*, *37*(42), 10215–10229.
- Mayberg, H. S., Brannan, S. K., Mahurin, R. K., Jerabek, P. A., Brickman, J. S., Tekell, J. L., ... Fox, P. T. (1997). Cingulate function in depression: A potential predictor of treatment response. *Neuroreport*, *3*(8), 1057–1061.
- McFarland, B. R., & Klein, D. N. (2009). Emotional reactivity in depression: Diminished responsiveness to anticipated reward but not to anticipated punishment or to nonreward or avoidance. *Depression and Anxiety*, *26*(2), 117–122.
- Moses-Kolko, E. L., Wsiner, K. L., Price, J. C., Berga, S. L., Drevets, W. C., Hanusa, B. H., ... Meltzer, C. C. (2008). Serotonin 1A receptor reductions in postpartum depression: A positron emission tomography study. *Fertility and Sterility*, *89*(3), 685–692.
- Nestler, E. J., & Carlezon, W. A. (2006). The mesolimbic dopamine reward circuit in depression. *Biological Psychiatry*, *59*(12), 1151–1159.
- Parsey, R. V., Oquendo, M. A., Ogden, R. T., Olvet, D. M., Simpson, N., Huang, Y.-Y., ... Mann, J. J. (2003). Altered serotonin 1A binding in major depression: A [carbonyl-C-11]WAY100635 positron emission tomography. *Biological Psychiatry*, *59*(2), 106–113.
- Pulcu, E., Thomas, E. J., Trotter, P. D., & McFarquhar, M. (2015). Social-economic decision making in current and remitted major depression. *Psychological Medicine*, *45*(6), 1301–1313.
- Roiser, J. P., Elliott, R., & Sahakian, B. J. (2012). Cognitive mechanisms of treatment in depression. *Neuropsychopharmacology*, *37*, 117–136.
- Sheehan, D. V., Lecrubier, Y., Sheehan, K. H., Amorim, P., Janavs, J., Weiller, E., ... Dunbar, G. C. (1998). The Mini-International Neuropsychiatric Interview (M.I.N.I.): The development and validation of a structured diagnostic psychiatric interview for DSM-IV and ICD-10. *Journal of Clinical Psychiatry*, *59*(20), 34–57.
- Speilberger, C. D., Gorsuch, R. L., Lushene, P. R., Vagg, P. R., & Jacobs, G. A. (1983). *State-trait anxiety inventory for adults (STAI-AD) manual*. Palo Alto, CA: Consulting Psychologists Press.
- Uhl, B., Kuehner, C., Kirsch, P., Ruttorf, M., Diener, C., & Flor, H. (2015). Altered reward and loss processing and prediction error signalling in depression. *Social Cognitive and Affective Neuroscience*, *10*(8), 1102–1112.
- Wechsler, D. (2001). *Wechsler text of adult reading*. San Antonio, TX: Harcourt Assessment.
- Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., & Wagenmakers, E.-J. (2011). Statistical evidence in experimental psychology: An empirical comparison using 855 *t* tests. *Perspectives in Psychological Science*, *6*(3), 291–298.
- World Health Organization (2017). Depression and other common mental disorders: Global Health Estimates. Geneva. Licence: CC BY-NC-SA 3.0 IGO.