

Low-Serotonin Levels Increase Delayed Reward Discounting in Humans

Nicolas Schweighofer,¹ Mathieu Bertin,^{2,5} Kazuhiro Shishida,³ Yasumasa Okamoto,³ Saori C. Tanaka,² Shigeto Yamawaki,³ and Kenji Doya^{2,4}

¹Department of Biokinesiology and Physical Therapy, University of Southern California, Los Angeles, California 90089, ²Computational Neuroscience Laboratories, Advanced Telecommunications Research Institute International, Kyoto 619-0288, Japan, ³Department of Psychiatry and Neurosciences, Hiroshima University, Higashihiroshima 739-8526, Japan, ⁴Okinawa Institute of Science and Technology, Okinawa 904-0411, Japan, and ⁵Laboratoire d'Informatique de Paris 6, Université Paris 6 Pierre et Marie Curie, 75005 Paris, France

Previous animal experiments have shown that serotonin is involved in the control of impulsive choice, as characterized by high preference for small immediate rewards over larger delayed rewards. Previous human studies under serotonin manipulation, however, have been either inconclusive on the effect on impulsivity or have shown an effect in the speed of action–reward learning or the optimality of action choice. Here, we manipulated central serotonergic levels of healthy volunteers by dietary tryptophan depletion and loading. Subjects performed a “dynamic” delayed reward choice task that required a continuous update of the reward value estimates to maximize total gain. By using a computational model of delayed reward choice learning, we estimated the parameters governing the subjects’ reward choices in low-, normal, and high-serotonin conditions. We found an increase of proportion in small reward choices, together with an increase in the rate of discounting of delayed rewards in the low-serotonin condition compared with the control and high-serotonin conditions. There were no significant differences between conditions in the speed of learning of the estimated delayed reward values or in the variability of reward choice. Therefore, in line with previous animal experiments, our results show that low-serotonin levels steepen delayed reward discounting in humans. The combined results of our previous and current studies suggest that serotonin may adjust the rate of delayed reward discounting via the modulation of specific loops in parallel corticobasal ganglia circuits.

Key words: serotonin; delay reward discounting; reinforcement learning model; tryptophan; decision; learning

Introduction

The importance of resisting the draw of small immediate rewards when larger rewards are available after a delay was well illustrated in the “marshmallow test” experiment, in which 4-year-old children could receive one marshmallow immediately or two marshmallows if they waited a few minutes. Ten years later, those individuals who resisted the immediate temptation were more academically and socially competent than their more impulsive counterparts (Shoda et al., 1990). When choosing between delayed rewards, animals and humans compare the estimated values associated with each reward and choose the reward associated with the larger value (Platt, 2002). Critical to these choices is the steepness of the discounting of the estimated reward value: $V = R * G(D)$, where R is the reward, and $G(D)$ is a discounting function that decreases with delay D . A steep rate of discounting results in smaller values of larger delayed rewards and, as a consequence, in increased choices of the more immediate small rewards. Although no discounting promotes large reward acquisition,

discounting is necessary in foraging tasks, in which pursuing large but delayed rewards is at the expense of other rewards. Therefore, to maximize total gain in such tasks, the discount rate should be carefully adjusted (Schweighofer and Doya, 2003; Schweighofer et al., 2006).

The function of the central serotonergic system is not well understood. Serotonin has been implicated in controlling impulsive behavior (Buhot, 1997; Robbins, 2000). In animals, increased serotonin levels decrease impulsive choice (Poulos et al., 1996; Bizot et al., 1999), and decreased serotonin levels increase impulsive choice (Wogar et al., 1993; Bizot et al., 1999), presumably via an increase in the rate of discounting (Mobini et al., 2000a,b). Not all studies, however, have found such an effect of serotonin in animals (Winstanley et al., 2004; Cardinal, 2006) or in humans (Crean et al., 2002; Tanaka et al., 2007). Furthermore, two other roles have been proposed for the ascending serotonergic system in reward choice tasks: modulation of the speed of action–reward learning (Clarke et al., 2004; Chamberlain et al., 2006) and modulation of the optimality of action choice (Rogers et al., 1999).

Here, we hypothesized that central serotonin level has an effect on reward choices via modulation of the rate of delayed reward discounting in humans, with low-serotonin levels yielding an increase in small reward choices via steep discounting, and high serotonin yielding an increase in large delayed reward choices via shallow discounting. To test this hypothesis, we ma-

Received Nov. 7, 2007; revised March 19, 2008; accepted March 19, 2008.

This work was supported in part by the Core Research for Evolutional Science and Technology and by National Science Foundation Grant IIS 0535282 (N.S.). We thank Cheol Han and Linda Fetters for helpful discussions.

Correspondence should be addressed to Nicolas Schweighofer, Department of Biokinesiology and Physical Therapy, University of Southern California, 1450 East Alcazar Street, Los Angeles, CA 90089. E-mail: schweigh@usc.edu.

DOI:10.1523/JNEUROSCI.4982-07.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/284528-05\$15.00/0

nipulated central serotonergic levels by dietary tryptophan depletion and loading (Young and Gauthier, 1981; Young et al., 1985; Carpenter et al., 1998). Additionally, to simultaneously control for possible effects of serotonin on the flexibility of action–reward learning and optimality of action choice in addition to discounting, we modeled the subjects' behavior in a “dynamic” delayed reward discounting task with a reinforcement learning model (Sutton and Barto, 1998) and studied the effect of serotonin condition on the three meta-parameters of the model (Doya, 2002): the discount factor γ , which controls the delayed reward discounting; the learning rate α , which controls the speed of action–reward learning; and the inverse temperature β , which controls the variability in action choice and therefore controls the balance between “exploration” and “exploitation” (supplemental material, available at www.jneurosci.org).

Materials and Methods

Serotonin manipulation. We used a within-subject, double-blind, placebo-controlled, and counter-balanced design. Twenty male subjects participated in three experimental sessions, with a minimum interval of 1 week between sessions. In the morning of each session, the subjects consumed one of three amino acid drinks: one containing a standard amount of tryptophan (control condition; 2.3 g per 100 g amino acid mixture), one containing excess tryptophan (loading condition; 10.3 g), and one without tryptophan (depletion condition; 0 g) (for details, see supplemental material, available at www.jneurosci.org). Six hours after consumption, plasma-free tryptophan was significantly lower in the depletion condition ($p < 0.001$; multiple comparisons after repeated measures ANOVA) and higher in the loading condition ($p < 0.001$) compared with the control condition.

Behavioral task. Most measures of delay discounting in humans are questionnaire based; subjects are asked to make a number of choices between small immediate rewards and larger rewards weeks, months, or even years in the future after thinking about the consequences of each alternative (Frederick et al., 2002). Here, we developed a novel task to measure delay discounting, akin to the tasks used in animal experiments and to a task for humans developed by Reynolds and Schiffbauer (2004), in which the delays are experienced by the subjects. Unlike in questionnaire-based tasks, which suffer from large between-subject variability (Frederick et al., 2002), there is little between-subject variability in the discount rates in our task (Schweighofer et al., 2006).

We modified the typical animal-like task in three important aspects, however. First, in animal experiments, the intertrial intervals are adjusted such that all trials have a constant duration; therefore, the number of choice opportunities is fixed, regardless of the actual choices. Here, to avoid too-long idling time after immediate reward choice, the intertrial interval is fixed, and the trial duration is therefore variable. We fixed the total session time such that the subjects would not select small, more immediate rewards to shorten the session; as a result, the total number of trials in each session was variable and depended on the subject's choices. Second, an adjusting delay procedure is often used in animal experiments. Such procedure, however, seems to be unreliable, because previous choices influenced the next choices regardless of the delays (Cardinal et al., 2002). In our task, the delays are randomly generated at each trial from a large range of possible delays, and the discounting factor is obtained by considering all trials. Third, to test potential effects of serotonin on the learning rate and exploratory choices, we created a “dynamic” environment: the delays for the same number of black patches slowly became shorter or longer for identical stimuli.

Subjects were instructed to maximize their reward gain in this “dynamic” delayed reward choice task. At each trial, subjects must choose between pursuing a more immediate 10 Japanese yen reward or a more delayed 40 yen reward. The reward amount and delay were coded with two colored square mosaics, each composed of 100 small patches, displayed side by side on a computer monitor (Fig. 1). The mosaic color coded for the monetary reward amount, white indicated the 10 yen reward and yellow the 40 yen reward, and the number of black patches in

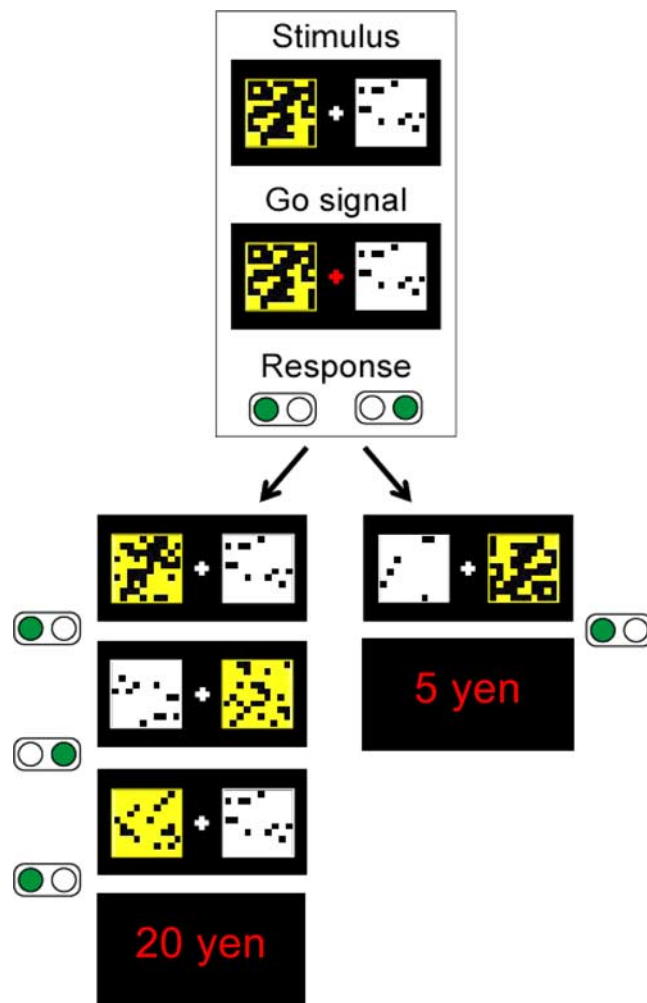


Figure 1. A trial of the delayed reward choice task. At the onset of the trial, the subject must select either a white or yellow mosaic after the fixation cross turns red (“Go” signal). Each button press (green disk) adds a number of colored patches to the selected mosaic. In the example shown here, if the white mosaic is selected, the subject receives 5 yen in two steps of 1.5 s each. If the yellow mosaic is selected, the subject receives 20 yen in four steps. The position of the squares (left or right) changed randomly at each step. At each trial, the initial numbers of black patches for both mosaics were randomly drawn from uniform distributions and indicated different delays. The intertrial interval, which corresponds to the reward display, was fixed and equal to one time step [adapted from the study by Schweighofer et al. (2006)].

each mosaic indicated the delays to the rewards (this implicit coding of the delay was designed to prevent explicit computation of the reward ratios). At each time step, the subjects were prompted to choose one of the two mosaics by pushing the corresponding mouse button. At the next time step, the stimulus selected in the previous step showed fewer black patches; the other stimulus was unchanged. The reward was displayed for one time step after filling of the corresponding square, and the next trial started at the next step. Because the numbers of initial black patches were randomly selected from uniform distributions, the delays varied at each trial (supplemental material, available at www.jneurosci.org).

The session was divided in 10 blocks, each lasting 210 s and separated by 15 s rest time from the next. In the first five blocks of the experiment, the number of patches removed per step was fixed. The first four blocks were used as warm-up blocks to obtain stable subject's responses and convergence of the value functions estimated by the model. Blocks six to nine corresponded to the “dynamic” environment, in which the number of patches removed per step was varied in small increments from trial to trial, so that the delays for the same number of black patches became shorter or longer for identical stimuli (Fig. S1, available at www.jneurosci.org as supplemental material). Therefore, to maximize the to-

tal gain, the subject needed to continuously update their reward value estimates in these blocks. In the last block, the environment was static again, and the parameters identical to those of the last trial of block nine. Although the subjects could not perceive the very small trial-by-trial changes in the distribution of parameters in the dynamic blocks, they slowly adapted to the new environment by modifying their reward choices (Fig. S1, available at www.jneurosci.org as supplemental material). Subjects were paid a sum corresponding to the total accumulated reward at the end of each session.

Reinforcement learning model. Reinforcement learning models offer a biologically plausible framework to model human delayed reward choices (Pagnoni et al., 2002; O'Doherty et al., 2004; Tanaka et al., 2004, 2007). A reinforcement learning agent aims at selecting the actions that maximize its reward gain by estimating discounted reward values (Sutton and Barto, 1998). In reinforcement learning, three meta-parameters control reward choice: the reward discounting factor $1 \geq \gamma > 0$, with a small value corresponding to steeper discounting; the speed of value update, or learning rate, $\alpha \geq 0$; and the inverse temperature $\beta > 0$, which controls the variability in action choice. Here, we used a reinforcement learning model called Q-learning (Sutton and Barto, 1998), which requires estimation of "action values." At each step, the action value $Q[s(t), a(t)]$ represents the expected sum of discounted future rewards by taking the action $a(t)$ (i.e., choosing the stimulus leading to the large or the small reward) at state $s(t)$ (i.e., for a particular estimated delay) and following the current policy in subsequent steps. For each subject, each serotonin condition, and each block, we systematically searched for the meta-parameters set $\{\alpha, \beta, \gamma\}$ that yielded the smallest "error" between the choices made by the model and those made by the subject (supplemental material, available at www.jneurosci.org).

Results

Effect of serotonin manipulation on small reward choices

A two-way repeated measure ANOVA (factors: tryptophan condition and block, both repeated measures; in this analysis, as well as in all repeated measure analyses that use block as a factor, only blocks 5–10 were included) showed a significant condition \times block interaction effect on the percentage of small reward choice ($F_{(10,190)} = 2.29$; $p = 0.015$). *Post hoc* comparisons (*t* tests; noncorrected) showed an increased percentage of small reward choice in the tryptophan depletion condition (mean \pm SE, $45.0 \pm 1.8\%$) compared with the loading condition ($41.0 \pm 2.3\%$; $p = 0.042$) and a trend compared with controls ($40.3 \pm 2.6\%$; $p = 0.082$) (Fig. 2a).

Serotonin modulates discounting, but not learning rate or variability of choice

The reinforcement learning model could accurately predict the subjects' behavior in all three conditions: the average proportion of correctly predicted reward choices for all subjects was 85.8% for the depletion condition (range, 78.8–91.0%), 86.0% for the control condition (range, 74.2–90.9%), and 85.1% for the loading condition (range, 79.0–93.0%).

A multivariate ANOVA (two factors: tryptophan condition and block) with the three meta-parameters as dependent variables showed a main effect of tryptophan condition on the discount factor γ ($F_{(2,342)} = 4.60$; $p = 0.011$) but not on the learning rate α ($F_{(2,342)} = 0.28$; $p = 0.76$) nor on the inverse temperature β ($F_{(2,342)} = 1.36$; $p = 0.26$). Because this multivariate analysis did not consider repeated measures, we performed separate two-way repeated measure ANOVAs (factors: tryptophan condition and block) for each meta-parameter. This analysis confirmed the




Figure 2. Effect of tryptophan conditions on the percentage of small reward choices (a), the discount parameter γ (b), the learning rate α (c), and the variability of action choice β (d). Note that the smaller discount parameters lead to steeper discounting. Asterisks indicate significant differences: * $p < 0.05$, ** $p < 0.01$.

main effect of tryptophan condition on the discount factor ($F_{(2,190)} = 5.37$; $p = 0.009$), with no interaction ($F_{(2,190)} = 0.664$; $p = 0.757$). The discount factor in the depletion condition ($\gamma = 0.822 \pm 0.014$) was smaller than in the control condition ($\gamma = 0.867 \pm 0.012$; noncorrected two-tail *t* test, $p = 0.009$) and smaller than in the loading condition ($\gamma = 0.868 \pm 0.014$; $p = 0.015$) (Fig. 2b). There was no main effect of tryptophan condition on the learning rate α ($F_{(2,190)} = 0.338$; $p = 0.715$) (Fig. 2c) and no condition \times block interaction ($F_{(2,190)} = 1.57$; $p = 0.118$). Similarly, there was no main effect of tryptophan condition on the inverse temperature β ($F_{(2,190)} = 1.50$; $p = 0.236$) (Fig. 2d) and no condition \times block interaction ($F_{(10,190)} = 0.867$; $p = 0.566$).

Fatigue caused by serotonin manipulation does not affect changes in discounting

In our experiment, serotonin manipulation had an effect on subjective fatigue (K. Shishida, Y. Okamoto, S. Asahi, G. Okada, K. Ueda, N. Shirao, A. Kinoshita, K. Onoda, H. Yamashita, N. Schweighofer, S. C. Tanaka, K. Doya, and S. Yamawaki, unpublished observations), as measured by the fatigue subscale of the Profile of Mood States questionnaire completed just before the second plasma measurement in each session (supplemental material, available at www.jneurosci.org). The fatigue score was significantly higher than in control condition for both the low- and high-serotonin conditions. There was no effect of the serotonergic manipulation on other aspects of moods (tension-anxiety, depression-dejection, anger-hostility, vigor-activity, and confusion-bewilderment) (Shishida, Okamoto, Asahi, Okada, Ueda, Shirao, Kinoshita, Onoda, Yamashita, Schweighofer, Tanaka, Doya, and Yamawaki, unpublished observations). Here, we verified that fatigue did not influence delayed reward discounting with a one-way ANOVA (tryptophan condition was the fixed factor, fatigue score the covariate, and the discounting parameter γ the dependent variable). In this model, tryptophan condition had a significant effect on γ ($F_{(2,356)} = 4.379$; $p = 0.013$), but fatigue did not ($p = 0.754$).

Discussion

Our results show low-serotonin levels cause impulsivity, as indicated by a higher level of small reward choices than for control serotonin levels, via an increase in delayed reward discounting. High-serotonergic levels, however, produced no detectable changes in reward choice and discounting compared with control levels. Furthermore, serotonin did not modulate the rate of update of the reward values nor the variability of reward choice. Although the observed differences in discount factor between the depletion and the control condition may appear small at first sight, the difference in discounted reward values becomes impor-

tant at longer delays: using the average discount factors, the values are two times larger after a 13 s delay in the loading condition than in the depletion condition. Nonetheless, the relatively mild manipulation of central serotonin used here (manipulation of central serotonin by dietary manipulation) compared with that used in animal experiments (lesions the dorsal raphe nucleus) is a likely explanation for the relatively smaller effect of low serotonin on small reward choice and discounting.

In our experiment, tryptophan loading had no effect on either reward choice or discounting. Because increased serotonergic levels has been shown to decrease impulsive choice in animals (Poulos et al., 1996; Bizot et al., 1999), a possibility is that tryptophan loading does not increase central serotonin sufficiently to yield a behavioral effect. Although other studies (Bjork et al., 1999; Luciana et al., 2001) have shown an effect of tryptophan loading compared with depletion in behavioral tasks, no placebo control was used in these studies. Therefore, there is not yet conclusive evidence that tryptophan loading induces any behavioral changes in humans. Use of the selective serotonin reuptake inhibitor (SSRI) citalopram would be an alternative method, because it has been shown to yield significant behavioral effects (Chamberlain et al., 2006).

Comparison with previous discounting studies

The results of our study contrast with those of a previous study, which found no effect of tryptophan depletion on delayed reward discounting (Crean et al., 2002). This difference may be attributable to differences in procedures, because Crean et al. (2002) used a questionnaire-based experiment in which the delays are not experienced. As we previously argued (Schweighofer et al., 2006), the decay rates observed in the present experiment were close to that observed in animal studies and similar to that reported in a related human experiment (Reynolds and Schiffbauer, 2004) but several orders of magnitude larger than that observed in other questionnaire-based human studies, such as that used by Crean et al. (2002). Similarly, in our previous study, which uses an experimental design similar to that of the present study (Tanaka et al., 2007), we did not demonstrate a significant behavioral effect of serotonin on reward discounting. The relatively small effect of tryptophan depletion on behavior, as observed in the present study (~4% increase in small reward choices in low-serotonin condition compared with other controls and loading), combined with a small number of trials in our previous study as a result of the limited functional magnetic resonance imaging (fMRI) experimental time, are likely explanations for this difference.

In all other studies aimed at studying the shape of human delayed reward discounting, discounting has been found to be hyperbolic. In our task, an exponential function better fits the data better than a hyperbolic function, and hyperbolic discounting does not maximize the theoretical gain (Schweighofer et al., 2006). An interesting feature of the hyperbolic discounting function is the prediction of the crossover of preference as the time to the small reward approaches. Our subjects, however, did not show such preference reversals, because only 0.58% of all trials for all subjects lead to reversals; this lack of reversal is predicted by exponential discounting.

A first potential confound in our study is a possible role of serotonin in modulating the perceived reward magnitude. In animals, serotonin depletion has been shown to only modulate delay discounting and not reward magnitude (Mobini et al., 2000b). However, in humans, tryptophan depletion has been shown to alter attention toward reward magnitude in other forms of choice behaviors (Rogers et al., 2003). A future experiment, in which

both magnitudes and delays are concurrently manipulated is thus needed to explore this possibility. Furthermore, our experiment would gain to be reproduced, because both the sample size and the effects are relatively small, and our findings are inconsistent with the results of Crean et al. (2002).

A second potential confound in our study is the possible effect of “effort” in influencing the decisions at each trial, because subjects were required to make a response at every time step. We do not believe, however, that the larger number of required responses to obtain large delayed rewards significantly affected decisions for two reasons. First, the effort required was constant during the whole experiment, because a button press was required at every time step, independent of the decisions. Second, in animal studies of effort-based decisions, rats are required to climb a 30 cm barrier (Rudebeck et al., 2006). Compared with such tests, the simple button press in our experiment requires little physical effort.

Possible neural mechanisms

Parallel corticobasal ganglia loops are involved in reward prediction at times scales similar to those of the present study (Tanaka et al., 2004). Additionally, in a previous fMRI study (Tanaka et al., 2007), we found that low-serotonin levels increased activity of the ventral part of the striatum, which correlated with more discounted reward values, and high-serotonin levels increased activity of the dorsal part of the striatum, which correlated with less discounted reward values. Therefore, serotonin may adjust delayed reward discounting for rewards available at short time scales, via the modulation of specific loops in a parallel corticobasal circuit. If imagery reward prediction is required, as in questionnaire-based human experiments, the frontal cortex would be additionally recruited (McClure et al., 2004). Serotonin depletion may have little or no effect on discounting of these imagery-evaluated delayed rewards.

Clinical implications of our study

Our findings, combined with those of our previous imaging experiment (Tanaka et al., 2007), suggest that serotonin plays a role in pathological impulsivity: we specifically predict that either low central serotonergic levels or abnormal serotonergic receptor distributions, or both, in the basal ganglia, can be a major cause of impulsive behavior.

In addition, serotonin seems to play a major role in depression, because SSRIs are known to effectively relieve symptoms; however, the therapeutic mechanisms of these drugs are still not well understood (Wong and Licinio, 2001). The increased rate of reward discounting that we found at low-serotonin levels may explain certain aspects of depressive behavior: when future rewards have values near zero, the optimal strategy is not to act. Future experiments using delayed reward paradigms could be designed to study impulsivity in patients with major depression.

References

- Bizot J, Le Bihan C, Puech AJ, Hamon M, Thiebot M (1999) Serotonin and tolerance to delay of reward in rats. *Psychopharmacology (Berl)* 146:400–412.
- Bjork JM, Dougherty DM, Moeller FG, Cherek DR, Swann AC (1999) The effects of tryptophan depletion and loading on laboratory aggression in men: time course and a food-restricted control. *Psychopharmacology (Berl)* 142:24–30.
- Buhot MC (1997) Serotonin receptors in cognitive behaviors. *Curr Opin Neurobiol* 7:243–254.
- Cardinal RN (2006) Neural systems implicated in delayed and probabilistic reinforcement. *Neural Netw* 19:1277–1301.

- Cardinal RN, Daw N, Robbins TW, Everitt BJ (2002) Local analysis of behaviour in the adjusting-delay task for assessing choice of delayed reinforcement. *Neural Netw* 15:617–634.
- Carpenter LL, Anderson GM, Pelton GH, Gudin JA, Kirwin PD, Price LH, Heninger GR, McDougle CJ (1998) Tryptophan depletion during continuous CSF sampling in healthy human subjects. *Neuropsychopharmacology* 19:26–35.
- Chamberlain SR, Muller U, Blackwell AD, Clark L, Robbins TW, Sahakian BJ (2006) Neurochemical modulation of response inhibition and probabilistic learning in humans. *Science* 311:861–863.
- Clarke HF, Dalley JW, Crofts HS, Robbins TW, Roberts AC (2004) Cognitive inflexibility after prefrontal serotonin depletion. *Science* 304:878–880.
- Crean J, Richards JB, de Wit H (2002) Effect of tryptophan depletion on impulsive behavior in men with or without a family history of alcoholism. *Behav Brain Res* 136:349–357.
- Doya K (2002) Metalearning and neuromodulation. *Neural Netw* 15:495–506.
- Frederick S, Loewenstein G, O'Donoghue T (2002) Time discounting and time preference: a critical review. *J Econ Lit* 40:351–401.
- Luciana M, Burgund ED, Berman M, Hanson KL (2001) Effects of tryptophan loading on verbal, spatial and affective working memory functions in healthy adults. *J Psychopharmacol* 15:219–230.
- McClure SM, Laibson DI, Loewenstein G, Cohen JD (2004) Separate neural systems value immediate and delayed monetary rewards. *Science* 306:503–507.
- Mobini S, Chiang TJ, Ho MY, Bradshaw CM, Szabadi E (2000a) Effects of central 5-hydroxytryptamine depletion on sensitivity to delayed and probabilistic reinforcement. *Psychopharmacology (Berl)* 152:390–397.
- Mobini S, Chiang TJ, Al-Ruwaitea AS, Ho MY, Bradshaw CM, Szabadi E (2000b) Effect of central 5-hydroxytryptamine depletion on intertemporal choice: a quantitative analysis. *Psychopharmacology (Berl)* 149:313–318.
- O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304:452–454.
- Pagnoni G, Zink CF, Montague PR, Berns GS (2002) Activity in human ventral striatum locked to errors of reward prediction. *Nat Neurosci* 5:97–98.
- Platt ML (2002) Neural correlates of decisions. *Curr Opin Neurobiol* 12:141–148.
- Poulos CX, Parker JL, Le AD (1996) Dexfenfluramine and 8-OH-DPAT modulate impulsivity in a delay-of-reward paradigm: implications for a correspondence with alcohol consumption. *Behav Pharmacol* 7:395–399.
- Reynolds B, Schiffbauer R (2004) Measuring state changes in human delay discounting: an experiential discounting task. *Behav Processes* 67:343–356.
- Robbins TW (2000) From arousal to cognition: the integrative position of the prefrontal cortex. *Prog Brain Res* 126:469–483.
- Rogers RD, Everitt BJ, Baldacchino A, Blackshaw AJ, Swainson R, Wynne K, Baker NB, Hunter J, Carthy T, Booker E, London M, Deakin JF, Sahakian BJ, Robbins TW (1999) Dissociable deficits in the decision-making cognition of chronic amphetamine abusers, opiate abusers, patients with focal damage to prefrontal cortex, and tryptophan-depleted normal volunteers: evidence for monoaminergic mechanisms. *Neuropsychopharmacology* 20:322–339.
- Rogers RD, Tunbridge EM, Bhagwagar Z, Drevets WC, Sahakian BJ, Carter CS (2003) Tryptophan depletion alters the decision-making of healthy volunteers through altered processing of reward cues. *Neuropsychopharmacology* 28:153–162.
- Rudebeck PH, Walton ME, Smyth AN, Bannerman DM, Rushworth MF (2006) Separate neural pathways process different decision costs. *Nat Neurosci* 9:1161–1168.
- Schweighofer N, Doya K (2003) Meta-learning in reinforcement learning. *Neural Netw* 16:5–9.
- Schweighofer N, Shishida K, Han CE, Okamoto Y, Tanaka SC, Yamawaki S, Doya K (2006) Humans can adopt optimal discounting strategy under real-time constraints. *PLoS Comput Biol* 11:1349–1356.
- Shoda Y, Mischel W, Peake PK (1990) Predicting adolescent cognitive and self-regulatory competencies from preschool delay of gratification. *Dev Psychol* 26:978–986.
- Sutton RS, Barto AG (1998) Reinforcement learning. Cambridge, MA: MIT.
- Tanaka SC, Doya K, Okada G, Ueda K, Okamoto Y, Yamawaki S (2004) Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nat Neurosci* 7:887–893.
- Tanaka SC, Schweighofer N, Asahi S, Shishida K, Okamoto Y, Yamawaki S, Doya K (2007) Serotonin differentially regulates short- and long-term prediction of rewards in the ventral and dorsal striatum. *PLoS ONE* 2:e1333.
- Winstanley CA, Dalley JW, Theobald DE, Robbins TW (2004) Fractionating impulsivity: contrasting effects of central 5-HT depletion on different measures of impulsive behavior. *Neuropsychopharmacology* 29:1331–1343.
- Wogar MA, Bradshaw CM, Szabadi E (1993) Effect of lesions of the ascending 5-hydroxytryptaminergic pathways on choice between delayed reinforcers. *Psychopharmacology (Berl)* 111:239–243.
- Wong ML, Licinio J (2001) Research and treatment approaches to depression. *Nat Rev Neurosci* 2:343–351.
- Young SN, Gauthier S (1981) Effect of tryptophan administration on tryptophan, 5-hydroxyindoleacetic acid and indoleacetic acid in human lumbar and cisternal cerebrospinal fluid. *J Neurol Neurosurg Psychiatry* 44:323–328.
- Young SN, Smith SE, Pihl RO, Ervin FR (1985) Tryptophan depletion causes a rapid lowering of mood in normal males. *Psychopharmacology (Berl)* 87:173–177.