Published in final edited form as:

Neural Netw. 2009 April ; 22(3): 294–304. doi:10.1016/j.neunet.2009.03.010.

Valuation of uncertain and delayed rewards in primate prefrontal cortex

Soyoun Kim¹, Jaewon Hwang², Hyojung Seo¹, and Daeyeol Lee¹

¹ Department of Neurobiology, Yale University School of Medicine

² Department of Brain and Cognitive Sciences, University of Rochester

Abstract

Humans and animals often must choose between rewards that differ in their qualities, magnitudes, immediacy, and likelihood, and must estimate these multiple reward parameters from their experience. However, the neural basis for such complex decision making is not well understood. To understand the role of the primate prefrontal cortex in determining the subjective value of delayed or uncertain reward, we examined the activity of individual prefrontal neurons during an inter-temporal choice task and a computer-simulated competitive game. Consistent with the findings from previous studies in humans and other animals, the monkey's behaviors during inter-temporal choice were well accounted for by a hyperbolic discount function. In addition, the activity of many neurons in the lateral prefrontal cortex reflected the signals related to the magnitude and delay of the reward expected from a particular action, and often encoded the difference in temporally discounted values that predicted the animal's choice. During a computerized matching pennies game, the animals approximated the optimal strategy, known as Nash equilibrium, using a reinforcement learning algorithm. We also found that many neurons in the lateral prefrontal cortex conveyed the signals related to the animal's previous choices and their outcomes, suggesting that this cortical area might play an important role in forming associations between actions and their outcomes. These results show that the primate lateral prefrontal cortex plays a central role in estimating the values of alternative actions based on multiple sources of information.

Keywords

game theory; inter-temporal choice; reinforcement learning; utility theory; temporal discounting

1. Introduction

Behavioral goals can be characterized in multiple dimensions. For example, two alternative rewards of different magnitudes might become available after unequal delays, and they might be forfeited according to different probabilities. Economic theories, such as the expected utility theory (von Neumann & Morgenstern, 1944), have shown that even for the outcomes with such diverse attributes, a numerical score or utility can be assigned to each option so that the decision maker's choices can be seen as a process of maximizing these quantities. A decision maker

Correspondence to: Daeyeol Lee, Ph.D., Department of Neurobiology, Yale University School of Medicine, 333 Cedar Street, New Haven, CT 06510, USA, phone: +1 203 785 3527, e-mail: E-mail: daeyeol.lee@yale.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

whose choices can be consistently described by such utility functions has been referred to as rational. Although actual choices of human decision makers do not always conform to this rationality criterion (Kahneman & Tversky, 1979; Rick & Loewenstein, 2008), the expected utility theory and its derivatives still provide a valuable framework to analyze the choice behaviors in a variety of contexts (Starmer, 2000). For example, the assumption that the utility of a reward expected from a given action decreases with its delay can account for the observation that people and animals are less likely to pursue a delayed reward than a more immediate reward of a similar magnitude (Frederick, Loewenstein, & O'Donoghue, 2002; Green & Myerson, 2004).

Inter-temporal choice refers to the problem of choosing among different rewards that may be delivered at different times. This occurs frequently in our daily lives, since rewards and other outcomes from many of our choices are often not immediately realized and might be delivered asynchronously. When the reward from a particular action is delayed, this decreases the decision maker's preference and tendency to take the action to acquire it. This is referred to as temporal discounting, and can be formalized mathematically by a discount function describing how the preference for a reward decreases with its delay. A large number of behavioral studies have demonstrated that the value of a delayed reward decreases more steeply when the delay is relatively short, indicating that a delayed reward might be discounted hyperbolically (Rachlin & Green, 1972; Ainslie & Herrnstein, 1981; Green, Fisher, Perlow, & Sherman, 1981; Green, Fry, & Myerson, 1994; Frederick et al., 2002; Green & Myerson, 2004; Kalenscher & Pennartz, 2008; Kim, Hwang, & Lee, 2008). Although the neural mechanisms involved in temporal discounting and inter-temporal choices are still poorly understood, previous neuroimaging studies have suggested that the fronto-cortico-striatal network may play an important role (Tanaka et al., 2004; McClure, Laibson, Loewenstein, & Cohen, 2004; McClure, Ericson, Laibson, Loewenstein, & Cohen, 2007; Kable & Glimcher, 2007; Berns, Laibson, & Loewenstein, 2007; Luhmann, Chun, Yi, Lee, & Wang, 2008; Weber & Huettel, 2008; Gregorios-Pippas, Tolber, & Schultz, 2009; Ballard & Knutson, 2009). In addition, single-neuron recording studies in non-human primates have shown that the information about the magnitude and timing of an upcoming reward is encoded in the activity of individual neurons in the dorsolateral prefrontal cortex (DLPFC; Leon & Shadlen, 1999; Roesch & Olson, 2005; Tsujimoto & Sawaguchi, 2005). Moreover, some DLPFC neurons encode the difference in the temporally discounted values of alternative choices, suggesting that the DLPFC might play a key role in inter-temporal choice (Kim et al., 2008). In the present study, we show that the activity of some DLPFC neurons might also reflect the temporally discounted value of the target chosen by the animal.

Although appropriately chosen utility functions can account for a variety of choice behaviors, this approach is limited when the decision makers change their preferences according to their experience. In other words, standard economic choice theory does not describe how the utility functions can be determined based on the previous experience of the decision maker. In the reinforcement learning theory (Sutton & Barto, 1998), value functions represent the weighted sum of rewards expected in the future, and analogous to the utility functions, determine the likelihood of choosing a given action in a particular context. In contrast to utility functions, however, value functions are modified when there is a discrepancy between the expected and actual outcomes. This discrepancy is referred to as reward prediction error, and the learning process driven by such errors can explain how the choices are modified by experience. The reinforcement learning theory has been enormously successful not only in describing the observed choice behaviors in humans and animals, but also in accounting for the neural activity observed in various brain areas during adaptive decision making (Daw & Doya, 2006; Lee, 2006; Dayan & Niv, 2008), especially when the environment is dynamic and incompletely known (Rushworth & Behrens, 2008). Previous studies have also found that when monkeys play a competitive game against a computer opponent, they tend to approximate the optimal

strategy using a reinforcement learning algorithm (Lee, Conroy, McGreevy, & Barraclough, 2004; Lee, McGreevy, & Barraclough, 2005). In addition, the activity in the prefrontal cortex encodes the signals related to the animal's previous choices and their outcomes (Barraclough, Conroy, & Lee, 2004; Seo & Lee, 2007, 2008). Therefore, the prefrontal cortex might contribute to improving the animal's decision-making strategies by providing the information necessary to adjust the value functions (Lee, Rushworth, Walton, Watanabe, & Sakagami, 2007). In this paper, we briefly review these behavioral findings and also summarize the key features of prefrontal activity observed during the matching pennies game. Combined with the results from the experiment on inter-temporal choice, these results suggest that the prefrontal cortex plays an important role not only in evaluating the possible outcomes from various actions, but also in providing accurate information for such evaluations.

2. Prefrontal cortex and inter-temporal choice

2.1. Mathematical forms of temporal discounting

The subjective value of a delayed reward tends to decrease with its delay. This has been accounted for by assuming that the decision makers choose among various rewards by computing for each reward a quantity referred to as temporally discounted value (DV). The DV for a delayed reward can be given by the product of its magnitude, R , and the discount function, $F(D)$, which describes the change in the value of the delayed reward as a function of its delay, D . In other words,

$$DV(R,D)=R F(D).$$

A large number of studies have explored various mathematical functions that might best describe this so-called temporal discounting or delay discounting. For simplicity, most of these studies have examined the discount function for a single type of reward, although there is evidence that discounting might be steeper for consumable goods than for money (McClure et al., 2004, 2007; Estle, Green, Myerson, & Holt, 2007). Denoting the discount function and its time derivative as $F(D)$ and $F'(D)$, respectively, discount rate r can be defined as $r=-F'(D)/F(D)$. If the discount rate is constant, then this leads to an exponential discount function, namely,

$$F(D)=\exp(-k_E D),$$

where k_E refers to the discount rate and determines how steeply the discount function decreases with delay. Despite its mathematical simplicity, empirical studies in both humans and animals have almost universally found that the exponential discount function does not account for the empirical data accurately (Frederick et al., 2002; Green & Myerson, 2004; Kalenscher & Pennartz, 2008). In particular, if the preference for a delayed reward is given by an exponential discount function, then the relative preference for two different reward magnitudes (R_1 and R_2) expected after waiting periods of D_1 and D_2 , respectively, should not change when the delays for both rewards are changed by the same amount, Δ . This can be illustrated as follows. For exponential discounting, the temporally discounted value of a reward with the magnitude R_1 and delay $D_1+\Delta$ is given by

$$DV(R_1,D_1+\Delta)=R_1 F(D_1+\Delta)=R_1 \exp\{-k_E(D_1+\Delta)\}=R_1 \exp(-k_E D_1) \exp(-k_E \Delta)=F_\Delta DV(R_1,D_1),$$

where $F_{\Delta} = \exp(-k_E \Delta)$. Similarly, the temporally discounted value of a reward with the magnitude R_2 and delay $D_2 + \Delta$ is given by

$$DV(R_2, D_2 + \Delta) = R_2 F(D_2 + \Delta) = R_2 \exp\{-k_E (D_2 + \Delta)\} = R_2 \exp(-k_E D_2) \exp(-k_E \Delta) = F_{\Delta} DV(R_2, D_2).$$

Therefore, $DV(R_1, D_1) > DV(R_2, D_2)$, if and only if $DV(R_1, D_1 + \Delta) > DV(R_2, D_2 + \Delta)$, since $F_{\Delta} > 0$. This property is referred to as time consistency. The violation of time consistency is often referred to as preference reversal, and has been empirically demonstrated in both humans and animals (Rachlin & Green, 1972; Ainslie & Herrnstein, 1981; Green et al., 1981, 1994). In contrast to exponential discount function, a hyperbolic discount function allows the discount rate to change with the reward delay, and is given by the following

$$F(D) = 1 / (1 + k_H D),$$

in which the value of the parameter k_H controls the steepness of temporal discounting. The discount rate for a hyperbolic discount function is $r_H = k_H / (1 + k_H D)$. Since $r_H = k_H$ when $D = 0$, k_H corresponds to the discount rate for a small delay ($D \approx 0$). A large number of studies have shown that a hyperbolic discount function accounts for the behavioral data during inter-temporal choice tasks in both humans and animals (Mazur, 1987; Rachlin, Raineri, & Cross, 1991; Simpson & Vuchinich, 2000; Myerson & Green, 1995; Kirby & Maraković, 1995; Kirby, 1997; Murphy, Vuchinich, & Simpson, 2001; Madden, Begotka, Raiff, & Kastern, 2003; Woolverton, Myerson, & Green, 2007; Kalenscher & Pennartz, 2008; Kim et al., 2008).

2.2. Behavioral studies on temporal discounting

Several studies have investigated how monkeys discount the values of delayed rewards (Tobin, Logue, Chelonis, & Ackerman, 1996; Stevens, Hallinan, & Hauser, 2005). Some investigators have also found that the steepness of temporal discount is comparable for chimpanzees and humans when they are tested with primary rewards (Rosati, Stevens, Hare, & Hauser, 2007). In addition, it has been shown that rhesus monkeys discount the value of self-administered cocaine hyperbolically according to its delay (Woolverton et al., 2007). However, it has not been tested whether monkeys show hyperbolic discounting for natural rewards. In addition, most previous studies of intertemporal choice in animals have gradually varied the amount or delay of reward to identify the point in which the animals become indifferent between immediate and delayed rewards. Under such conditions, the animal's choices tend to be correlated in successive trials (Cardinal, Daw, Robbins, & Everitt, 2002). This is not ideal for neurophysiological studies on temporal discounting, because neurons involved in decision making might change their activity according to the animal's previous choices and their outcomes (Barraclough et al., 2004; Seo, Barraclough, & Lee, 2007), making it difficult to distinguish it from the activity related to the decision variables in the current trial. Therefore, we have developed a novel inter-temporal choice task in which the amount and delay of reward resulting from a particular choice are indicated by visual stimuli and hence can be manipulated independently across trials (Kim et al., 2008).

In our study, rhesus monkeys were trained to indicate their choices between two different peripheral targets by making an eye movement towards one of them (Fig 1). At the beginning of each trial, the animal was required to fixate a white square presented at the center of a computer screen. After a 1-s fore-period, two peripheral targets were presented along the horizontal meridian. One of the targets (TS) was green and delivered a small amount of juice reward (0.27 ml) when it was chosen by the animal, whereas the other target (TL) was red and

delivered a larger amount of reward (0.4 ml). When the central target was extinguished after a 1-s cue period, the animal was required to shift its gaze towards one of these two peripheral targets. Each peripheral target could be accompanied by a number of yellow disks, which indicated the interval between the onset of fixation on a peripheral target and the time of reward delivery. Namely, when the chosen target included N yellow disks in its “clock”, the reward was delivered N seconds after the animal fixated the target. The behavioral data described below (Fig 2) were obtained while the reward delay ranged from 0 to 8 s in steps of 1 s, corresponding to 0 to 8 yellow disks. All possible delay combinations for TS and TL were included in this experiment as long as the delay for the TL was equal to or longer than the delay for the TS. This produced 45 different combinations of reward delays. In addition, the positions of the TS and TL were counter-balanced across trials, resulting in 90 trials in each block. Two male rhesus monkeys (D and J) were tested in this experiment. Each animal performed 10 blocks in a daily session, and was tested for 5 days.

When both green and red targets were presented without yellow disks, indicating that any chosen reward would be delivered immediately, the animal almost always chose the red, hence large-reward, target (TL; Fig 2). In contrast, as the delay for the TL increased, the probability of choosing the TS increased gradually. Similarly, the animals were more likely to choose the TL as the delay for TS increased (Fig 2). In the following, the set of parameters corresponding to the magnitudes (R_S and R_L) and delays (D_S and D_L) for small and large rewards is denoted by the symbol Ω . To test whether these results are accounted for better by an exponential or hyperbolic discount function, we assumed that the animal chose the TS with the probability given by the softmax transformation of the temporally discounted values (DV) of the small and large rewards. Namely, the probability that the animal would choose the small-reward target (TS) given Ω , $P(TS|\Omega)$, is given by the following.

$$P(TS|\Omega) = 1 / [1 + \exp - \beta \{DV(R_S, D_S) - DV(R_L, D_L)\}],$$

where β denotes the inverse temperature parameter controlling the degree of randomness in the animal's choices. The probability that the animal would choose the large reward target (TL) was given by $P(TL|\Omega) = 1 - P(TS|\Omega)$. The likelihood of the animal's choices in the entire dataset was then given by the following.

$$L = \prod_t p(c_t|\Omega_t) = p(c_1|\Omega_1) p(c_2|\Omega_2) \dots p(c_N|\Omega_N),$$

where c_t (TS or TL) indicates the animal's choice in trial t , Ω_t the magnitudes and delays for the rewards in trial t , and N the number of trials. The parameters (k_E or k_H , and β) that maximize the likelihood function (Pawitan, 2001) were estimated separately for exponential and hyperbolic discount functions.

The results of this analysis showed that for both monkeys, the hyperbolic discount function fit the behavioral data substantially better than the exponential discount function (Fig 2). A close examination of the data illustrated in Figure 2 shows that compared to the predictions of the hyperbolic discount function, the exponential discount function tends to underestimate the tendency for the animal to choose the small reward that is available immediately without any delays. The log likelihood for the exponential discount function computed for the entire data set was $-1,907.2$ and $-2,097.9$ for monkeys D and J, respectively, whereas corresponding values for the hyperbolic discount function were $-1,742.9$ and -1927.8 . The corresponding values of log likelihood ratio were therefore 164.3 and 170.1, indicating that the hyperbolic discount function was strongly preferred. The maximum likelihood estimate for the parameter

k_H in the hyperbolic discount function was 0.25 and 0.21, for monkeys D and J, respectively. For hyperbolic discount function, $F(D)=0.5$, when $D=1/k_H$. Therefore, these results indicate that for the monkeys tested in this study, the subjective value of a delayed reward was halved when it was delayed by approximately 4 to 4.8 s.

2.3. Neural coding of temporally discounted values in the prefrontal cortex

To investigate whether and how the signals related to the magnitudes and delays of the expected rewards are encoded in the prefrontal cortex, we recorded the activity of individual neurons in the dorsolateral prefrontal during the inter-temporal choice task described above (Kim et al., 2008). During this recording experiment, the delay for the small reward was 0 or 2 s, whereas the delay for the larger reward was 0, 2, 4, 6, or 8 s. All possible combinations of delays were tested as long as the delay for the large reward was no less than the delay for the small reward. With the positions of the large-reward and small-reward targets counter-balanced across trials, this resulted in 18 trials/block. In this study, we also included a control task to test whether neural activity seemingly related to the temporally discounted values might result from the visual stimuli used to indicate the magnitude and delay of reward. During the control task, the color of the fixation target (green or red) indicated which target the animal was required to choose, and the animal always received a small, fixed amount of juice reward (0.27 ml) immediately after it fixated the correct target. Otherwise, the stimuli and procedure used in the control task were the same as in the inter-temporal choice. The inter-temporal choice task and the control task were given in alternating blocks.

The activity of 164 neurons recorded in the dorsolateral prefrontal cortex during this task (116 daily sessions) was analyzed. All of these neurons were tested in at least 6 blocks for each condition (216 trials), and most of them (135 neurons) were tested in 10 blocks (360 trials). Consistent with the results described above, the behavioral data obtained during the inter-temporal choice task in this recording experiment were fit better by hyperbolic discount functions than by exponential discount functions. For example, for the behavioral data illustrated in Figure 3, the log likelihood for the exponential discount function was -41.9 and -48.7 for monkeys D and J, respectively, whereas the corresponding values for the hyperbolic discount function were -36.9 and 38.1. Overall, the behavioral data were fit better by the hyperbolic discount function than by the exponential discount function in 81.0% of the sessions. The average log likelihood ratio between these two models was 1.81 ± 0.38 and 6.35 ± 0.66 for monkeys D and J (mean $\pm$ SEM), and this difference was statistically significant (paired t-test, $p < 10^{-5}$) for both animals. The median value of the parameter k_H in the hyperbolic discount function was 0.12 and 0.35 s^{-1} for monkeys D and J, respectively.

To test whether the activity of individual neurons in the DLPFC encoded signals that influence the animal's decisions during this task, we applied the following regression model to the spike count of each neuron during the 1-s cue period.

$$S = a_0 + a_1 C + a_2 (DV_L - DV_R) + a_3 DV_{cho},$$

where S denotes the spike count, C the animal's choice (0 and 1 for right-ward and left-ward choice), and DV_L , DV_R , and DV_{cho} denote the temporally discounted values for the rewards expected from the left-ward and right-ward targets, and from the target chosen by the animal, respectively. The results from the behavioral analysis show that the animals tend to prefer the target with the greater temporally discounted value. Therefore, neurons modulating their activity according to the difference in the temporally discounted values for the two targets, $(DV_L - DV_R)$, might provide the signals necessary to determine the animal's choice. The results from this analysis showed that some neurons in the DLPFC changed their activity

significantly according to the difference in the temporally discounted values for the two alternative targets. For example, the DLPFC neuron shown in Figure 4a increased its activity as the temporally discounted value for left-ward target increased relative to the temporally discounted value for the right-ward target (Fig 4a left; t-test, $p < 0.005$). The activity of the same neuron was not related to the temporally discounted value for the target chosen by the animal (Fig 4a, bottom left). To test whether the activity seemingly related to the temporally discounted values could reflect the response to the visual features used to indicate the magnitudes and delays of different rewards, we analyzed the activity of the same neuron during the control task. It should be noted that the temporally discounted values for all the targets in the control task were constant. Therefore, to test whether the colors of the targets and the number of disks in their clocks similarly influenced the activity of DLPFC neurons during the inter-temporal choice and control tasks, we computed fictitious temporally discounted values (FDV) for each target in the control task as if the magnitudes and delays of their rewards changed as in the inter-temporal choice task. If the neural activity seemingly related to the temporally discounted values during the inter-temporal choice were indeed due to the visual features used to indicate the magnitudes and delays of different rewards, then it should be also modulated similarly by the FDV. The activity of the neuron shown in Figure 4a, however, was not significantly related to the FDV in the control task (Fig 4a, right). To further test whether this difference is statistically significant or not, we applied the following regression model with a series of interaction terms.

$$S = a_0 + a_1 C + a_2 (DV_L - DV_R) + a_3 DV_{cho} + a_4 T + a_5 T \times C + a_6 T \times (DV_L - DV_R) + a_7 T \times DV_{cho},$$

where T refers to the task (0 and 1 for inter-temporal choice and control tasks, respectively). The neuron shown in Figure 4a showed a significant interaction between the task and the difference in the temporally discounted values, suggesting that its activity was modulated by the temporally discounted values only during the inter-temporal choice task. Overall, the difference in the temporally discounted values significantly modulated the activity for 16.5% of the neurons in the DLPFC (Fig 5a), and approximately half of them (48.2%) showed significant interactions between the task and the difference in the temporally discounted values.

In addition to the difference in the temporally discounted values for the two alternative targets, neurons in the DLPFC also encoded the temporally discounted value of the reward expected from the target chosen by the animal in a given trial. For example, the activity of the neuron illustrated in Figure 4b increased its activity as the temporally discounted value of the left-ward target during the cue period of the trials in which the animal chose the left-ward target, whereas its activity increased with the temporally discounted value of the right-ward target when the animal chose the right-ward target. Indeed, the regression analysis showed that the effect of the temporally discounted value for the chosen target was statistically significant (t-test, $p < 0.0005$), whereas the effect of the difference in the temporally discounted values was not significant ($p = 0.91$). In addition, similar to the neuron shown in Figure 4a, this neuron changed its activity according to the temporally discounted value of the chosen target only during the inter-temporal choice task, but not during the control task ($p = 0.23$). The regression analysis also showed that this neuron showed a significant interaction between the task and the temporally discounted value of the chosen target (t-test, $p < 0.05$). Overall, 30.5% of the neurons in the DLPFC showed significant modulation in their activity related to the temporally discounted value of the chosen target (Fig 5a). Among them, 28% of the neurons also showed significant interactions between the task and the temporally discounted value of the chosen target.

We have also examined the time course of neural signals in the DLPFC related to the temporally discounted values and the animal's choice by performing the same regression analysis described above using a 200-ms sliding window with a 50-ms step size. For each time step, we determined the fraction of neurons that showed statistically significant modulations in their activity according to each variable, as well as the average amount of variance accounted for by the same variable determined by the coefficient of partial determination (CPD; Kim et al., 2008). The results show that the signals related to the difference in the temporally discounted values of the two alternative rewards and the temporally discounted value of the reward chosen by the animal emerged immediately after the cue onset and almost simultaneously in the activity of neurons in the DLPFC (Fig 5b and 5c). In contrast, the signals related to the animal's choice increased gradually during the cue period. This suggests that the signals related to the animal's preference for delayed rewards were gradually converted to the animal's overt behavioral response in the DLPFC.

3. Prefrontal cortex and reinforcement learning

3.1. Learning mixed strategies in competitive games

During the inter-temporal choice task discussed above, the animals were given complete information about the outcomes of their choices. However, for many decision-making problems in real life, there is often uncertainty about the outcomes of alternative actions. This is especially true in social contexts, in which the outcome of a decision maker's choice is determined by the interactions among multiple decision makers or players. Game theory refers to the mathematical studies on such decision making in social contexts (von Neumann & Morgenstern, 1944). In game theory, a game is defined by a payoff matrix which describes the payoff to each player as a function of the choices of all the players. In addition, Nash equilibrium refers to a set of strategies defined for all players such that no players can increase their payoffs by deviating from such strategies individually. Although many human and animal behaviors can be at least approximately accounted for by the concept of Nash equilibrium, the predictions of Nash equilibrium can be sometimes violated. This might result because human decision makers are sometimes altruistic and therefore choose actions that can increase the well-beings of others at the expense of their own (Fehr & Fischbacher, 2003; Fehr & Camerer, 2007). In addition, gradual convergence of decision-making strategies towards the Nash equilibrium strategies might indicate that decision makers might approximate the optimal strategies through experience, for example, by using a reinforcement learning algorithm (Sutton & Barto, 1998; Camerer, 2003; Lee, 2008; Dayan, 2009; Zhang, 2009).

To investigate whether and how the primate prefrontal cortex contributes to the process of approximating an optimal strategy during socially interactive decision making, we trained rhesus monkeys in a computerized matching pennies game (Barraclough et al., 2004; Lee et al., 2004; Seo & Lee, 2007, 2008). During this experiment, the animal started each trial by fixating a small square at the center of a computer screen (Fig 6A). After a 0.5-s fore-period, two green peripheral targets were presented along the horizontal meridian. When the central fixation target was extinguished after a 0.5-s delay period, the animal was required to shift its gaze towards one of the peripheral targets and maintain its fixation on the chosen target for 0.5 s. At the end of this fixation period, the computer displayed a red ring around the peripheral target it had selected, and the animal was rewarded with a small juice reward only when it chose the same target as the computer opponent. The computer was programmed to simulate a rational decision maker in the matching pennies game and made its choices so as to minimize the expected payoff to the animal. This was accomplished by estimating a series of conditional probabilities for the animal's choices and using this information to determine the target that the animal was more likely to choose (see Lee et al., 2004). For example, if the animal had a tendency to choose the same target that was rewarded in the previous trial, then the computer opponent increased the probability of choosing the right-ward target after the animal was

rewarded for choosing the left-ward target. Therefore, the optimal (Nash-equilibrium) strategy for the animal during this matching pennies game was to choose the two targets with the same probabilities and to do so independently from its previous choices and their outcomes.

Behavioral data were obtained from 6 rhesus monkeys for a total of 319 daily sessions and 200,228 trials. All animals chose the two targets almost equally frequently, and the probability of choosing each target averaged across all the sessions did not deviate more than 1% from 50% in any of the animals. In contrast, all animals displayed the tendency to choose the same target that was rewarded in the previous trial and switch to the other target otherwise. Averaged across all the animals, the probability that the animal would choose its target according to this so-called win-stay-lose-switch strategy was 0.532. Across different sessions, there was a significant negative correlation between the probability of win-stay-lose-switch strategy and the animal's reward rate ($r=0.212$, $p<0.001$). In other words, as the animal chose its targets more frequently according to the win-stay-lose-switch strategy, this was exploited by the computer opponent and decreased the animal's reward rate. To quantify how the animal's choice in a given trial and its outcome influenced the animal's choices in subsequent trials, the following logistic regression model was applied.

$$\begin{aligned} \text{logit } P_t(R) &\equiv \log P_t(R)/\{1-P_t(R)\} \\ &= a_0 + \mathbf{A}_c [c_{t-1} \ c_{t-2} \ \dots \ c_{t-10}] + \mathbf{A}_o [o_{t-1} \ o_{t-2} \ \dots \ o_{t-10}], \end{aligned}$$

where c_t and o_t indicate the choice of the animal and that of the computer opponent in trial t , respectively (1 and -1 for right-ward and left-ward choices, respectively). Individual elements of $\mathbf{A}_c$ indicate whether the animal has the tendency to choose the same target as in previous trials, whereas those of $\mathbf{A}_o$ indicate whether the animal tends to choose the same target that was chosen by the computer and therefore was or would have been rewarded in previous trials. The results from this logistic regression analysis showed that there was a little or no consistent tendency for the animals to choose the same target repeatedly (Fig 6B, left). In contrast, all animals showed strong tendencies to choose the same target chosen by the computer opponent in any of the last several trials (Fig 6B, right), suggesting that they might have used a reinforcement learning algorithm to improve their decision making strategies (Sutton & Barto, 1998; Lee et al., 2004; Seo & Lee, 2007).

3.2. Neural basis of reinforcement learning in the prefrontal cortex

In the reinforcement learning theory (Sutton & Barto, 1998), the probability of taking a given action is given by a set of value functions, which are the estimates of long-term rewards. For the binary choices during the matching pennies task, the probability of choosing the right-ward target can be given by the soft-max transformation of the value functions for the left-ward and right-ward targets. Namely,

$$P_t(R) = 1 / [1 + \exp - \beta \{Q_t(R) - Q_t(L)\}],$$

where $Q_t(x)$ denotes the value function for target x in trial t , and β the inverse temperature. The value function for the target chosen in trial t was updated according to the reward prediction error, as follows.

$$Q_{t+1}(x) = Q_t(x) + \alpha \{r_t - Q_t(x)\},$$

where r_t denotes the reward in trial t (0 and 1 for unrewarded and rewarded trials, respectively), and α the learning rate. The value function for the target not chosen by the animal was not updated. The two parameters of this model, α and β , were estimated for the entire behavioral data obtained from each animal using a maximum likelihood procedure (Pawitan, 2001). The value functions estimated from this reinforcement learning model were then used in the following regression model to test whether they are encoded in the activity of each neuron in the DLPFC.

$$S_t = b_0 + b_c c_t + b_d \{Q_t(R) - Q_t(L)\} + b_s \{Q_t(R) + Q_t(L)\},$$

where S_t denotes the spike rate during the 0.5-s delay period in trial t , and c_t the animal's choice. In this model, how the neural activity was influenced by the difference in the value functions for the two alternative choices and their sum was modeled separately. The difference in the value functions might be used by the animal to determine its choice, whereas their sum might be related to the state value function (Seo & Lee, 2008; Belova, Paton, & Salzman, 2008). This analysis was applied to 322 neurons recorded in the DLPFC of 5 monkeys (Seo et al., 2007), and the statistical significance of each regression coefficient was determined using a permutation test (Seo & Lee, 2008, 2009). The results of this analysis showed that, during the delay period, the activity of DLPFC neurons was sometimes modulated by the difference in the value functions for the two targets. For example, the neuron illustrated in Figure 7 increased its activity significantly as the value function for the right-ward target increased relative to that for the left-ward target. Overall, 13.4% of the neurons in the DLPFC significantly modulated their activity according to the difference in the value functions. Therefore, many neurons in the DLPFC encoded the signals that can be potentially used to determine the animal's choice. In addition, 28.2% and 14.0% of the DLPFC neurons showed significant modulations in their activity according to the animal's choice and the sum of the value functions for the two targets (Seo & Lee, 2008).

The results described above suggest that the neurons in the DLPFC might play an important role in integrating the signals related to the animal's previous choices and their outcomes to update the value functions. To test this directly, we applied the following regression model that includes the previous choices of the animal and computer opponent as well as the animal's choice outcomes.

$$S_t = \mathbf{B} [1 \ u_t \ u_{t-1} \ u_{t-1} \ u_{t-1}]',$$

where u_t is a row vector consisting of 3 dummy variables corresponding to the animal's choice (0 and 1 for the left-ward and right-ward choices, respectively), the computer's choice (coded in the same way as the animal's choice), and the reward (1 for rewarded trials and 0 otherwise) in trial t , and $\mathbf{B}$ is a vector of 13 regression coefficients. This analysis was performed separately for the spike rates measured with a series of non-overlapping 0.5-s bins defined relative to the time of target onset or feedback onset. The results showed that many DLPFC neurons encoded the signals related to the animal's choice and its outcome as well as the computer's choice, whereas the magnitudes and precise time courses of such signals varied substantially across different neurons (Seo et al., 2007; Seo & Lee, 2008). For example, the DLPFC neuron shown in Figure 8 displayed significantly higher activity during the peripheral fixation period and the feedback period when the animal chose the right-ward target compared to when it chose the left-ward target (Fig 8, top, Trial Lag=0). The activity of the same neuron was also significantly higher during the delay period when the animal had chosen the right-ward target in the previous trial than when its previous choice was the left-ward target (Fig 8, top, Trial Lag=1). In addition,

the same neuron showed significantly higher activity from the feedback period through the delay period in the following trial, when the computer opponent chose the right-ward target, compared to when the choice of the computer opponent was the left-ward target (Fig 8, middle, Trial Lag=0 and 1). Finally, the activity of this neuron was also modulated by the outcome of the animal's choice (Fig 8, bottom). Across the population, a large fraction of the neurons in the DLPFC showed significant modulations in their activity according to these three variables (Fig 9).

4. Discussion

Almost all theories and models of decision making postulate a set of variables that quantify the desirability of outcomes expected from alternative actions. In economic theories, they correspond to utility functions, and summarize crisply the preference of the decision makers for all possible actions. In reinforcement learning, value functions represent the empirical estimates of possible outcomes expected from various actions, and are continually updated according to the experience of the decision maker. These two theoretical frameworks can account for a wide range of choice behaviors observed in humans and animals. Therefore, the neural structures, such as the prefrontal cortex, in which the signals related to utilities and value functions are identified, are likely to play a key role in decision making. For example, inter-temporal choice behaviors in many different animal species can be concisely accounted for by hyperbolic discount functions. Previously, we have shown that the activity of individual neurons in the DLPFC encodes signals related to the difference in the temporally discounted values, suggesting that the DLPFC provides the neural signals necessary for evaluating the magnitude and delay of each reward during inter-temporal choice (Kim et al., 2008). In the present study, we also found that the activity of some neurons in the DLPFC might encode the temporally discounted value of the target chosen by the animal. Similar findings have been obtained in the orbitofrontal cortex during a behavioral task in which monkeys chose between two different types of juices that varied in their magnitudes (Padoa-Schioppa & Assad, 2006). Therefore, the encoding of subjective value for an option or action selected by the animal might be a common feature in the primate prefrontal cortex.

In most studies on inter-temporal choice, the information about the magnitude and delay of each reward is presented unambiguously in order to examine the effect of reward delay independently of the effect of reward uncertainty. In contrast, in many real-life problems of decision making, the outcomes of various actions and their timing can be highly uncertain, and the decision makers are often required to adjust their decision-making strategies based on the actual outcomes of their previous choices. This adaptive process can be formally accounted for by the reinforcement learning theory (Sutton & Barto, 1998). Similar to the findings from previous studies in human subjects (Mookherjee & Sopher, 1994; Erev & Roth, 1998; Camerer, 2003), monkeys tend to approximate the optimal decision-making strategy during a computer-simulated competitive game using a simple reinforcement learning algorithm (Lee et al., 2004). In addition, the activity of individual neurons in the DLPFC encoded several different types of signals required to implement such an algorithm (Barraclough et al., 2004; Seo et al., 2007). For example, many neurons in the DLPFC modulate their activity according to the difference in the value functions for the two alternative targets, which in turn determines the probability of choosing each of the two targets (Seo & Lee, 2008). During the matching pennies game, the probability of obtaining reward by choosing a particular target is determined by the probability that the same target would be chosen by the opponent. Similarly, some neurons in the DLPFC changed their activity according to the choice of the computer opponent, suggesting that such neurons might be specifically involved in updating the value functions for alternative targets. Neurons in the DLPFC also encoded signals related to the animal's previous choices, which might be used to associate particular choices and their subsequent outcomes. Such memory signals related to the animal's previous choices are referred to as eligibility trace

(Sutton & Barto, 1998). Therefore, the primate prefrontal cortex might provide the source of eligibility trace necessary for reinforcement learning, which has also been observed in the striatum (Kim et al., 2007; Lau & Glimcher, 2007). Finally, many neurons in the DLPFC show modulations in their activity according to the outcomes of the animal's previous choices, namely the animal's reward history. These signals might be used to compute the overall reward rate, which could provide valuable information about possible changes in the animal's environment.

Although the present study focused on the role of the DLPFC in decision making, many of the computational steps described above are likely to involve a network of cortical and subcortical areas connected to the DLPFC (Lee, 2006; Levine, 2009). For example, the signals related to the magnitude and probability of expected reward have been identified in a number of brain areas, including the posterior parietal cortex (Platt & Glimcher, 1999; Sugrue, Corrado, & Newsome, 2004; Dorris & Glimcher, 2004) and basal ganglia (Hollerman, Tremblay, & Schultz, 1998; Kawagoe, Takikawa, & Hikosaka, 1998; Samejima, Ueda, Doya, & Kimura, 2005). In addition, signals related to the outcome of the animal's choice, such as the reward prediction error, have been found in the midbrain dopamine neurons (Schultz, 1998) as well as in the anterior cingulate cortex (Matsumoto, Matsumoto, Abe, & Tanaka, 2007; Seo & Lee, 2007). Therefore, it remains an important question how the signals related to the outcomes of previous choices can be incorporated into the value functions represented in multiple brain areas. For example, signals related to the value functions, previous choices of the animals, and their outcomes might converge in the basal ganglia (Kim et al., 2007; Lau & Glimcher, 2007, 2008) in addition to the DLPFC. This suggests that the prefrontal cortex and basal ganglia might play a particularly important role in updating the value functions. However, this hypothesis needs to be tested more rigorously in future studies.

Neurobiological studies on decision making that adopted computational framework of economics and reinforcement learning theories have provided many important insights. Nevertheless, in their present forms, neither framework is likely to provide a complete account of the cognitive capabilities of human and other animal decision makers. For example, reinforcement learning algorithms that do not rely on explicit models of the animal's environment are often referred to simple or model-free (Sutton and Barto, 1998). In addition, most reinforcement learning algorithms used in previous studies update their value functions only according to the actual outcomes of the animal's previous choices. However, human decision makers and possibly other non-human primates might be capable of improving their decision-making strategies more efficiently by exploiting the structural knowledge of their environment (Hampton, Bossaerts, & O'Doherty, 2006; Pan, Sawa, Tsuda, Tsukada, & Sakagami, 2008; Dayan & Niv, 2008; Dayan, 2009) as well as by observing the hypothetical payoffs that could have been obtained from the actions not chosen by the decision maker (Heyman, Mellers, Rishcenko, & Schwartz, 2004; Lee et al., 2005; Lohrenz, McCabe, Camerer, & Montague, 2007; Lee, 2008). Although the neural basis of these more efficient learning algorithms might overlap with the brain areas involved in simple reinforcement learning, how these different learning algorithms coexist and interact in the brain is still not well understood and remains as an important topic for future research.

Acknowledgments

We thank Dominic Barraclough for his contribution to some of the studies described in this manuscript. This study was supported by the grants from the National Institute of Health (MH 073246 and DA 024855).

References

- Ainslie G, Herrnstein RJ. Preference reversal and delayed reinforcement. *Animal Learning & Behavior* 1981;9:476–482.

- Ballard K, Knutson B. Dissociable neural representations of future reward magnitude and delay during temporal discounting. *Neuroimage* 2009;45:143–150. [PubMed: 19071223]
- Barracough DJ, Conroy ML, Lee D. Prefrontal cortex and decision making in a mixed-strategy game. *Nature Neuroscience* 2004;7:404–410.
- Belova MA, Paton JJ, Salzman CD. Moment-to-moment tracking of state value in the amygdala. *Journal of Neuroscience* 2008;28:10023–10030. [PubMed: 18829960]
- Berns GS, Laibson D, Loewenstein G. Intertemporal choice - towards an integrative framework. *Trends in Cognitive Science* 2007;11:482–488.
- Camerer, CF. *Behavioral game Theory: experiments in strategic interaction*. Princeton, NJ: Princeton University Press; 2003.
- Cardinal RN, Daw N, Robbins TW, Everitt BJ. Local analysis of behavior in the adjusting-delay task for assessing choice of delayed reinforcement. *Neural Networks* 2002;15:617–634. [PubMed: 12371516]
- Daw ND, Doya K. The computational neurobiology of learning and reward. *Current Opinion in Neurobiology* 2006;16:199–204. [PubMed: 16563737]
- Dayan P. Goal-directed control and its antipodes. *Neural Networks*. 2009This issue
- Dayan P, Niv Y. Reinforcement learning: the good, the bad and the ugly. *Current Opinion in Neurobiology* 2008;18:185–196. [PubMed: 18708140]
- Dorris MC, Glimcher PW. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 2004;44:365–378. [PubMed: 15473973]
- Erev I, Roth AE. Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 1998;88:848–881.
- Estle SJ, Green L, Myerson J, Holt DD. Discounting of monetary and directly consumable rewards. *Psychological Science* 2007;18:58–63. [PubMed: 17362379]
- Fehr E, Camerer CF. Social neuroeconomics: the neural circuitry of social preferences. *Trends in Cognitive Science* 2007;11:419–427.
- Fehr E, Fischbacher U. The nature of human altruism. *Nature* 2003;425:785–791. [PubMed: 14574401]
- Frederick S, Loewenstein G, O'Donoghue T. Time discounting and time preference: a critical review. *Journal of Economic Literature* 2002;40:351–401.
- Green L, Fisher EB Jr, Perlow S, Sherman L. Preference reversal and self control: choice as a function of reward amount and delay. *Behaviour Analysis Letters* 1981;1:43–51.
- Green L, Fry AF, Myerson J. Discounting of delayed rewards - a life-span comparison. *Psychological Sciences* 1994;5:33–36.
- Green L, Myerson J. A discounting framework for choice with delayed and probabilistic reward. *Psychological Bulletin* 2004;130:769–792. [PubMed: 15367080]
- Gregorios-Pippas L, Tobler PN, Schultz W. Short-term temporal discounting of reward value in human ventral striatum. *Journal of Neurophysiology* 2009;101:1507–1523. [PubMed: 19164109]
- Hampton AN, Bossaerts P, O'Doherty JP. The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *Journal of Neuroscience* 2006;26:8360–8367. [PubMed: 16899731]
- Heyman J, Mellers B, Tishchenko S, Schwartz A. I was pleased a moment ago: how pleasure varies with background and foreground reference points. *Motivation and Emotion* 2004;28:65–83.
- Hollerman JR, Tremblay L, Schultz W. Influence of reward expectation on behavior-related neuronal activity in primate striatum. *Journal of Neurophysiology* 1998;80:947–963. [PubMed: 9705481]
- Kable JW, Glimcher PW. The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience* 2007;10:1625–1633.
- Kahneman D, Tversky A. Prospect theory: an analysis of decision under risk. *Econometrica* 1979;47:263–291.
- Kalenscher T, Pennartz CMA. Is a bird in the hand worth two in the future? The neuroeconomics of intertemporal decision-making. *Progress in Neurobiology* 2008;84:284–315. [PubMed: 18207301]
- Kawagoe R, Takikawa Y, Hikosaka O. Expectation of reward modulates cognitive signals in the basal ganglia. *Nature Neuroscience* 1998;1:411–416.
- Kim S, Hwang J, Lee D. Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron* 2008;59:161–172. [PubMed: 18614037]

- Kim YB, Huh N, Lee H, Baeg EH, Lee D, Jung MW. Encoding of action history in the rat ventral striatum. *Journal of Neurophysiology* 2007;98:3548–3556. [PubMed: 17942629]
- Kirby KN. Bidding on the future: evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General* 1997;126:54–70.
- Kirby KN, Maraković NN. Modeling myopic decisions: evidence for hyperbolic delay-discounting within subjects and amounts. *Organizational Behavior and Human Decision Processes* 1995;64:22–30.
- Lau B, Glimcher PW. Action and outcome encoding in the primate caudate nucleus. *Journal of Neuroscience* 2007;27:14502–14514. [PubMed: 18160658]
- Lau B, Glimcher PW. Value representations in the primate striatum during matching behavior. *Neuron* 2008;58:451–463. [PubMed: 18466754]
- Lee D. Neural basis of quasi-rational decision making. *Current Opinion in Neurobiology* 2006;16:191–198. [PubMed: 16531040]
- Lee D. Game theory and neural basis of social decision making. *Nature Neuroscience* 2008;11:404–409.
- Lee D, Conroy ML, McGreevy BP, Barraclough DJ. Reinforcement learning and decision making in monkeys during a competitive game. *Cognitive Brain Research* 2004;22:45–58. [PubMed: 15561500]
- Lee D, McGreevy BP, Barraclough DJ. Learning and decision making in monkeys during a rock-paper-scissors game. *Cognitive Brain Research* 2005;25:416–430. [PubMed: 16095886]
- Lee D, Rushworth MFS, Walton ME, Watanabe M, Sakagami M. Functional specialization of the primate frontal cortex during decision making. *Journal of Neuroscience* 2007;27:8170–8173. [PubMed: 17670961]
- Leon MI, Shadlen MN. Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron* 1999;24:415–425. [PubMed: 10571234]
- Levine DS. Brain pathways for cognitive-emotional decision making in the human animal. *Neural Networks*. 2009This issue
- Lohrenz T, McCabe K, Camerer CF, Montague PR. Neural signature of fictive learning signals in a sequential investment task. *Proceedings of the National Academy of Science of the United States of America* 2007;104:9493–9498.
- Luhmann C, Chun MM, Yi D-J, Lee D, Wang X-J. Neural dissociation of delay and uncertainty in intertemporal choice. *Journal of Neuroscience*. 2009In press
- Madden GJ, Begotka AM, Raiff BR, Kastern LL. Delay discounting of real and hypothetical rewards. *Experimental and Clinical Psychopharmacology* 2003;11:139–145. [PubMed: 12755458]
- Matsumoto M, Matsumoto K, Abe H, Tanaka K. Medial prefrontal cell activity signaling prediction errors of action values. *Nature Neuroscience* 2007;10:647–656.
- Mazur, JE. An adjusting procedure for studying delayed reinforcement. In: Commons, ML.; Mazur, JE.; Nevin, JA.; Rachlin, H., editors. *Quantitative analyses of behaviors: the effect of delay and of intervening events on reinforcement value*. Hillsdale, NJ: Erlbaum; 1987. p. 55-73.
- McClure SM, Ericson KM, Laibson DI, Loewenstein G, Cohen JD. Time discounting for primary rewards. *Journal of Neuroscience* 2007;27:5792–5804.
- McClure SM, Laibson DI, Loewenstein G, Cohen JD. Separate neural systems value immediate and delayed monetary rewards. *Science* 2004;306:503–507. [PubMed: 15486304]
- Mookherjee D, Sopher B. Learning behavior in an experimental matching pennies game. *Games and Economic Behavior* 1994;7:62–91.
- Murphy JG, Vuchinich RE, Simpson CA. Delayed reward and cost discounting. *Psychological Record* 2001;51:571–588.
- Myerson J, Green L. Discounting of delayed rewards: models of individual choice. *Journal of the Experimental Analysis of Behavior* 1995;64:263–276. [PubMed: 16812772]
- Padoa-Schioppa C, Assad JA. Neurons in the orbitofrontal cortex encode economic value. *Nature* 2006;441:223–226. [PubMed: 16633341]
- Pan X, Sawa K, Tsuda I, Tsukada M, Sakagami M. Reward prediction based on stimulus categorization in primate lateral prefrontal cortex. *Nature Neuroscience* 2008;11:703–712.
- Pawitan, Y. *In all likelihood: statistical modelling and inference using likelihood*. New York: Oxford University Press; 2001.

- Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. *Nature* 1999;400:233–238. [PubMed: 10421364]
- Rachlin H, Green L. Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior* 1972;17:15–22. [PubMed: 16811561]
- Rachlin H, Raineri A, Cross D. Subjective probability and delay. *Journal of the Experimental Analysis of Behavior* 1991;55:233–244. [PubMed: 2037827]
- Rick S, Loewenstein G. Intangibility in intertemporal choice. *Philosophical Transactions of the Royal Society B* 2008;363:3813–3824.
- Roesch MR, Olson CR. Neuronal activity dependent on anticipated and elapsed delay in macaque prefrontal cortex, frontal and supplementary eye fields, and premotor cortex. *Journal of Neurophysiology* 2005;94:1469–1497. [PubMed: 15817652]
- Rosati AG, Stevens JR, Hare B, Hauser MD. The evolutionary origins of human patience: temporal preferences in chimpanzees, bonobos, and human adults. *Current Biology* 2007;17:1663–1668. [PubMed: 17900899]
- Rushworth MFS, Behrens TEJ. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience* 2008;11:389–397.
- Samejima K, Ueda Y, Doya K, Kimura M. Representation of action-specific reward values in the striatum. *Science* 2005;310:1337–1340. [PubMed: 16311337]
- Schultz W. Predictive reward signal of dopamine neurons. *Journal of Neurophysiology* 1998;80:1–27. [PubMed: 9658025]
- Seo H, Barraclough DJ, Lee D. Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. *Cerebral Cortex* 2007;17:i110–i117. [PubMed: 17548802]
- Seo H, Lee D. Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *Journal of Neuroscience* 2007;27:8366–8377. [PubMed: 17670983]
- Seo H, Lee D. Cortical mechanisms for reinforcement learning in competitive games. *Philosophical Transactions of the Royal Society B* 2008;363:3845–3857.
- Seo H, Lee D. Behavioral and neural changes after gains and losses of conditioned reinforcers. *Journal of Neuroscience*. 2009in press
- Simpson CA, Vuchinich RE. Reliability of a measure of temporal discounting. *Psychological Research* 2000;50:3–16.
- Starmer C. Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* 2000;38:332–382.
- Stevens JR, Hallinan EV, Hauser MD. The ecology and evolution of patience in two new world monkeys. *Biology Letters* 2005;1:223–226. [PubMed: 17148172]
- Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. *Science* 2004;304:1782–1787. [PubMed: 15205529]
- Sutton, R.; Barto, AG. Reinforcement learning: an introduction. Cambridge, MA: MIT Press; 1998.
- Tanaka SC, Doya K, Okada G, Ueda K, Okamoto Y, Yamawaki S. Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nature Neuroscience* 2004;7:887–893.
- Tobin H, Logue AW, Chelonis JJ, Ackerman KT. Self-control in the monkey *Macaca fascicularis*. *Animal Learning & Behavior* 1996;24:168–174.
- Tsujimoto S, Sawaguchi T. Neuronal activity representing temporal prediction of reward in the primate prefrontal cortex. *Journal of Neurophysiology* 2005;93:3687–3692. [PubMed: 15634707]
- von Neumann, J.; Morgenstern, O. Theory of games and economic behavior. Princeton, NJ: Princeton University Press; 1944.
- Weber BJ, Huettel SA. The neural substrates of probabilistic and intertemporal decision making. *Brain Research* 2008;1234:104–115. [PubMed: 18710652]
- Woolverton WL, Myerson J, Green L. Delay discounting of cocaine by rhesus monkeys. *Experimental and Clinical Psychopharmacology* 2007;15:238–244. [PubMed: 17563210]
- Zhang J. Adaptive learning via selectionism and Bayesianism. Part I: connection between the two. *Neural Networks*. 2009This issue

Figure 1. Spatio-temporal sequence of inter-temporal choice task. In this example, the delays for the rewards resulting from the small-reward (green) and large-reward (red) targets were 2 and 6 s, respectively, and the animal chose the large-reward target.

Figure 2.

Probability of choosing the small-reward target, $P(TS)$, during the inter-temporal choice task plotted for various combinations of delays for small (TS) and large (TL) reward targets. These results were obtained in separate behavioral experiments conducted prior to the neurophysiological experiments. Dots correspond to the data, whereas lines are the predictions from the best-fitting hyperbolic discount function.

Figure 3.

Probability of choosing the small-reward target, $P(TS)$, for an example daily session during neurophysiological experiments. Solid and dotted lines correspond to the predictions from the best-fitting hyperbolic and exponential discount functions, respectively. TS (TL) refers to the small-reward (large-reward) target.

Figure 4.

Activity of two example neurons in the DLPFC (a and b) during the inter-temporal choice (left) and control (right) tasks. Top two rows show the average spike density functions for the trials in which the animal chose the right-ward and left-ward targets, respectively. Activity was averaged separately according to the difference in the temporally discounted values for the two targets, and the magnitudes and delays of the rewards for the two targets are indicated by the legend on the right. For example, S(0):L(0) indicates that the left-ward target delivered a small reward, and that the delays for both rewards was 0s, whereas L(4):S(2) indicate that the left-ward target gave a large reward, and that the delays for the small and large rewards were 2 and 4 s, respectively. Third row shows the average spike rates during the cue period plotted according to the difference in the temporally discounted values (DV) and fictitious discounted values (FDV) for the inter-temporal choice and control task, respectively. Finally, the bottom row shows the average spike rates during the cue period plotted as a function of the DV or FDV for the target chosen by the animal in the inter-temporal choice or control task. Filled and empty symbols indicate the results for the trials in which the animals chose the right-ward and left-ward targets, respectively.

Figure 5.

Population summary of neural activity in the DLPFC during the inter-temporal choice task. a. Fraction of neurons that significantly modulated their activity during the cue period according to the animal's choice (Cho), the difference in the temporally discounted values ($\Delta DV = DV_L - DV_R$), and the temporally discounted value of the chosen target (DV_{cho}). b. Fraction of neurons in the DLPFC that show significant modulations in their activity related to the same variables estimated from a sliding-window regression analysis. c. Coefficient of partial determination (CPD) for the same variables estimated by the sliding-window regression analysis. Shaded areas correspond to the mean $\pm$ SEM. Gray background in b and c corresponds to the cue period.

Figure 6.

Sequence of visual stimuli (a) and performance (b) during the matching pennies task. The line plots in (b) show the average regression coefficients determined by a logistic regression model related to the animal's previous choices (left) or the choices of the computer opponent (right). Small dots indicate the results from individual animals, whereas the large black disks indicate that the overall average values were significantly different from zero (t-test, $p < 0.05$). The gray histograms in (b) show the proportion of the daily sessions in which the corresponding regression coefficient was significantly positive (upper half) or negative (lower half).

Figure 7.

Activity of an example neuron in the DLPFC showing significant modulation in its activity according to the difference in the value functions during the matching pennies (left). The activity of the same neuron was not significantly related to the sum of the value functions (right). The empty and filled circles indicate the average activity in each decile of trials sorted according to the difference in the value functions (left) or sum of the value functions (right) for the trials in which the animal chose the left-ward and right-ward targets, respectively. Error bars, SEM. Histograms show the distribution of trials along the same variables for the trials in which the animal chose the left-ward (gray) or right-ward (black) trials.

Figure 8.

Activity of an example DLPFC neuron (same neuron shown in Fig 7) during the matching pennies task. Each pair of small panels displays the spike density functions estimated relative to the time of target onset (left panels) or feedback onset (right panels), separately according to the animal's choice (top), the choice of the computer opponent (middle), or the animal's choice outcome (bottom) in the current (Trial Lag=0) or previous (Trial Lag = 1 to 3) trials. Gray (black) lines correspond to the activity associated with right-ward (left-ward) choices in the top two panels or rewarded (unrewarded) trials in the bottom panels. Circles show the regression coefficients from a multiple linear regression model, which was performed separately for a series of 0.5-s windows. Large circles indicate the coefficients significantly different from zero (t-test, $p < 0.05$). Gray background indicates the delay (left panels) or feedback (right panels) periods.

Figure 9.

Time course of DLPFC activity related to the animal's choice (top), the choice of the computer opponent (middle), and the animal's choice outcome (reward, bottom). Each symbol indicates the fraction of the neurons that displayed significant modulations in their activity according to the corresponding variable in the same regression analysis used in Figure 8. Large circles indicate that the percentage of neurons was significantly higher (binomial test, $p < 0.05$) than the significance level used in the regression analysis (0.05). Gray background indicates the delay (left panels) or feedback (right panels) periods.