Behavioral/Systems/Cognitive

Temporal Discounting of Reward and the Cost of Time in Motor Control

Reza Shadmehr, Jean Jacques Orban de Xivry, Minnan Xu-Wilson, and Ting-Yu Shih

Department of Biomedical Engineering, The Johns Hopkins University School of Medicine, Baltimore, Maryland 21205

Why do movements take a characteristic amount of time, and why do diseases that affect the reward system alter control of movements? Suppose that the purpose of any movement is to position our body in a more rewarding state. People and other animals discount future reward as a hyperbolic function of time. Here, we show that across populations of people and monkeys there is a correlation between discounting of reward and control of movements. We consider saccadic eye movements and hypothesize that duration of a movement is equivalent to a delay of reward. The hyperbolic cost of this delay not only accounts for kinematics of saccades in adults, it also accounts for the faster saccades of children, who temporally discount reward more steeply. Our theory explains why saccade velocities increase when reward is elevated, and why disorders in the encoding of reward, for example in Parkinson's disease and schizophrenia, produce changes in saccade. We show that delay of reward elevates the cost of saccades, reducing velocities. Finally, we consider coordinated movements that include motion of eyes and head and find that their kinematics is also consistent with a hyperbolic, reward-dependent cost of time. Therefore, each voluntary movement carries a cost because its duration delays acquisition of reward. The cost depends on the value that the brain assigns to stimuli, and the rate at which it discounts this value in time. The motor commands that move our eyes reflect this cost of time.

Introduction

Passage of time discounts the value of reward. For example, college students would rather receive $400 now than wait for 5 years to receive $1000 (Myerson and Green, 1995). This implies that for young people the value of $1000 drops to less than $400 in 5 years. For older people, this value drops more slowly, and for children, the value drops more quickly (Green et al., 1999). Psychologists have characterized this behavior via a hyperbolic reward discount function. If α represents the value of something at present, and β is the rate at which we discount this value in time, then the value at some time t in the future is as follows:

$$V(t) = \frac{\alpha}{1 + \beta t}. \tag{1}$$

Response of dopamine neurons to stimuli that promise future reward also follows this hyperbolic form. When monkeys view a stimulus that predicts when they will receive a drop of juice, in response to the stimulus that predicts the shortest time delay midbrain dopamine neurons discharge strongly, whereas for the

stimulus that predicts a longer delay the discharge declines hyperbolically (Kobayashi and Schultz, 2008).

Here, we suggest that there is a connection between how the brain discounts reward in time and how it controls movements. We begin with the assumption that the purpose of any voluntary movement is to change the state of our body to one that is more valuable. Because the passage of time discounts this value (i.e., we would rather receive the reward now than later), the duration of movement carries a specific penalty (Eq. 1). We will ask whether this penalty can explain why movements of people and other animals take a characteristic amount of time.

Our focus will be on control of saccades, as this movement has been measured in numerous populations and conditions. Eye kinematics during a saccade exhibits curious properties. For example, people produce higher velocity saccades when they view a face (Xu-Wilson et al., 2009). Aging of the brain alters saccade velocities: velocities are highest in children and lowest in the elderly (Fioravanti et al., 1995; Munoz et al., 2003). Patients with Parkinson's disease (PD) have reduced saccade velocities (Nakamura et al., 1991), whereas schizophrenic patients have increased velocities (Mahlberg et al., 2001). Saccades of some species of monkeys are nearly twice as fast as those of humans (Straube et al., 1997; Chen-Harris et al., 2008). We will suggest that, in all these cases, the specific velocities and durations of saccades arise from a desire to maximize reward in a setting in which reward loses value hyperbolically as a function of movement duration. Finally, we will consider the fact that natural eye movements in people typically accompany head movements (Guitton and Volle, 1987) (i.e., voluntary movements rarely involve a single body part). We will show that the timing and velocities of these coordinated movements, as well as some of their variability at-

tributable to task conditions (Epelboim et al., 1997), are also consistent with our theory. We suggest that our brain views duration of movements as an implicit cost because passage of time discounts the value of future reward.

Materials and Methods

Theory. Let us assume that to make a movement the brain solves the following problem: generate motor commands to acquire as much reward as possible while expending as little effort as possible. Suppose that at time t, the state of our eye is described by vector $\mathbf{x}(t)$ (representing position, velocity, etc.), our motor commands are $u(t)$, and our target is a stimulus at position g (with respect to the fovea). Furthermore, suppose that the brain assigns some reward value α to the target. For example, faces may be more valuable than inanimate objects. The reward is acquired when the image of the target is on the fovea, which will require a movement that will take time. The key assumption is that the motor system will incur a cost for delaying the acquisition of reward because of the time p that it takes to place the valuable image on the fovea as follows:

$$J_p = \alpha\left(1 - \frac{1}{1 + \beta p}\right). \quad (2)$$

Therefore, the longer it takes to get the target on the fovea, the larger the loss of reward value (Eq. 2).

To move the eyes, we will have to spend some effort in terms of motor commands. There is little information about how the brain represents effort. Two recent results suggest that this cost is approximately a quadratic function of force (Fagg et al., 2002; O'Sullivan et al., 2009) as follows:

$$J_u = \lambda\int_0^p u^2(t)dt. \quad (3)$$

When our movement ends at time $t = p$, our eye position $x(p)$ should coincide with where reward is (i.e., target position g). This constitutes an accuracy cost, and it is convenient to represent it also as a quadratic function as follows:

$$J_x = E[(x(p) - g)^2]. \quad (4)$$

In Equation 4, $E[]$ is the expected value operator. In summary, we assume that in performing a movement the brain attempts to produce motor commands that minimize a cost that depends on accuracy, effort, and temporal discounting of reward as follows:

$$J = J_x + J_u + J_p. \quad (5)$$

The idea of a cost associated with endpoint accuracy (e.g., variance) was introduced by Harris and Wolpert (1998). The idea of a cost associated with effort was introduced by Todorov and Jordan (2002). These two costs by themselves are insufficient to explain movements because, without a cost for time, all movements are unnaturally slow. Recently, Harris and Wolpert (2006) suggested a cost for time that increased linearly as a function of movement duration. Here, we will show that, if we assume that the cost of time is related to the temporal discounting of reward (i.e., a hyperbolic cost of time), we will not only account for saccade kinematics better than any previous model but also explain why there are changes in movements when there are changes in reward processing in the brain.

Our first objective is to ask whether a hyperbolic temporal cost can account for the kinematics (i.e., duration, velocity, etc.) of saccades. Our second objective is to ask whether this temporal cost is related to reward processing in the brain. These objectives require solving an optimal control problem in which Equation 5 serves as a cost function. The crucial prediction of the theory is that there should be specific changes in saccade durations and velocities because of changes in the reward discounting function (Eq. 2), for example, because of changes in stimulus reward value α or temporal discounting rate β.

Control of saccades. We modeled the dynamics of the human eye as a discrete linear system with signal-dependent noise as follows:

$$\begin{aligned} \mathbf{x}^{(k+1)} &= A\mathbf{x}^{(k)} + \mathbf{b}(u^{(k)} + \varepsilon^{(k)}) \\ \varepsilon^{(k)} &\sim N(0, \kappa^2(u^{(k)})^2). \\ \mathbf{y}^{(k)} &= C\mathbf{x}^{(k)} \end{aligned} \quad (6)$$

The superscripts on the above equations refer to time steps. The term ε is signal-dependent noise (i.e., a random variable with a normal distribution of mean zero and SD that linearly scales with the motor commands). Our objective was to find motor commands $\mathbf{u}_h = [u^{(0)}, u^{(1)}, \ldots, u^{(p-1)}]^T$ that minimized the cost as follows:

$$J = E[(\mathbf{y}^{(p)} - \mathbf{r})^T T(\mathbf{y}^{(p)} - \mathbf{r})] + \mathbf{u}_h^T L\mathbf{u}_h + \alpha\left(1 - \frac{1}{1 + \beta p}\right), \quad (7)$$

where

$$T \equiv \begin{pmatrix} \nu_1 & 0 & 0 \\ 0 & \nu_2 & 0 \\ 0 & 0 & \nu_3 \end{pmatrix} \quad \mathbf{r} \equiv \begin{bmatrix} g \\ 0 \\ 0 \end{bmatrix}$$

$$L \equiv \begin{pmatrix} \lambda^{(0)} & 0 & 0 & 0 \\ 0 & \lambda^{(1)} & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & \cdots & 0 & \lambda^{(p-1)} \end{pmatrix}. \quad (8)$$

The first term in Equation 7 enforces our desire to have endpoint accuracy. It penalizes the expected squared difference between the state of the eye at movement end and the goal state (i.e., the sum of bias and variance of the movement). The second term penalizes effort. The third term is a cost of time, as passage of time discounts reward. We define the following:

$$\boldsymbol{\varepsilon}_h = \begin{bmatrix} \varepsilon^{(0)} \\ \varepsilon^{(1)} \\ \vdots \\ \varepsilon^{(p-1)} \end{bmatrix} \quad U = \begin{bmatrix} u^{(0)} & 0 & \cdots & 0 \\ 0 & u^{(1)} & 0 & 0 \\ \vdots & 0 & \ddots & \vdots \\ 0 & 0 & 0 & u^{(p-1)} \end{bmatrix}. \quad (9)$$

The mean and variance of our noise vector are as follows:

$$E[\boldsymbol{\varepsilon}_h] = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} \quad \text{var}[\boldsymbol{\varepsilon}_h] = \kappa^2 UU. \quad (10)$$

The state at end of the movement is as follows:

$$\mathbf{x}^{(p)} = A^p\mathbf{x}^{(0)} + F\Gamma(\mathbf{u}_h + \boldsymbol{\varepsilon}_h), \quad (11)$$

where

$$F \equiv [A^{p-1} \ A^{p-2} \ A^{p-3} \cdots I] \quad \Gamma \equiv \begin{bmatrix} \mathbf{b} & 0 & \cdots & 0 \\ 0 & \mathbf{b} & 0 & 0 \\ \vdots & 0 & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{b} \end{bmatrix}. \quad (12)$$

The expected value and variance of our state at the end of the movement are as follows:

$$\begin{aligned} E[\mathbf{x}^{(p)}] &= A^p\mathbf{x}^{(0)} + F\Gamma\mathbf{u}_h \\ \text{var}[\mathbf{x}^{(p)}] &= \kappa^2 F\Gamma UU\Gamma^T F^T \end{aligned} \quad (13)$$

Therefore, we have:

$$E[J] = E[\mathbf{x}^{(p)}]^T C^T TC E[\mathbf{x}^{(p)}] + tr[C^T TC\text{var}[\mathbf{x}^{(p)}]]$$

$$- 2E[\mathbf{x}^{(p)}]^T C^T T\mathbf{r} + \mathbf{r}^T T\mathbf{r} + \mathbf{u}_h^T L\mathbf{u}_h + \alpha\left(1 - \frac{1}{1 + \beta p}\right). \quad (14)$$

In the above equation, $tr[]$ is the trace operator. We can simplify the trace operator as follows:

$$
\begin{aligned}
tr[C^T TC\ var[\mathbf{x}^{(p)}]] &= \kappa^2 tr[C^T\ TCF\Gamma UU\Gamma^T\ F^T] \\
&= \kappa^2 tr[U\Gamma^T\ F^T\ C^T\ TCF\Gamma U] \\
&= \kappa^2 \mathbf{u}_h^T\ diag\ [\Gamma^T\ F^T\ C^T\ TCF\Gamma]\mathbf{u}_h
\end{aligned}
\quad . \quad (15)
$$

The term $diag[X]$ in Equation 15 is the diagonal operator that generates a matrix with only the diagonal elements of the square matrix X. Setting the derivative of Equation 15 with respect to $\mathbf{u}_h$ to zero and solving for $\mathbf{u}_h$ gives us the optimal sequence of motor commands for a given duration of movement:

$$
S = diag[\Gamma^T\ F^T\ C^T\ TCF\Gamma]
$$

$$
\mathbf{u}_h^\star(p) = (L + \Gamma^T\ F^T\ C^T\ TCF\Gamma + \kappa^2 S)^{-1} \qquad (16)
$$

$$
\times\ \Gamma^T\ F^T\ C^T T(\mathbf{r}_h - CA^p\ \mathbf{x}^{(0)}).
$$

However, what is the optimum duration of our movement? To arrive at this, we divide the problem into two parts: first, we select an arbitrary duration p and find the optimal set of motor commands $\mathbf{u}_h^\star(p)$ via Equation 16, and then compute the cost of this movement via Equation 7. Next, we search the space of p for the one movement duration that provides the minimum cost J, as illustrated in Figure 1. In our simulations, all saccades started from $\mathbf{x}^{(0)} = \mathbf{0}$.

Parameter values. Our eye plant model in continuous form is as follows:

$$
\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} =
\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ \dfrac{-c_4}{c_1} & \dfrac{-c_3}{c_1} & \dfrac{-c_2}{c_1} \end{bmatrix}
\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} +
\begin{bmatrix} 0 \\ 0 \\ \dfrac{1}{c_1} \end{bmatrix} (\boldsymbol{u} + \boldsymbol{\varepsilon}).
$$

$$(17)$$

For the human eye, we used time constants of 224, 13, and 4 ms (Keller, 1973; Robinson et al., 1986). For the eyes of the rhesus monkey, we used time constants of 260, 12, and 1 ms (Fuchs et al., 1988). The constants in Equation 17 are related to these time constants as follows: $c_4 = 1, c_3 = \tau_1 + \tau_2 + \tau_3, c_2 = \tau_1\tau_2 + \tau_2\tau_3 + \tau_1\tau_3$, and $c_1 = \tau_1\tau_2\tau_3$. For example, for the human eye, $\tau_1 = 0.224, \tau_2 = 0.013$, and $\tau_3 = 0.004$. The continuous equations were transformed to discrete time using matrix exponentials with time interval of 1 ms. The goal of the movement is to position the eyes at the target with zero velocity and acceleration, $\mathbf{r} = [g\ 0\ 0]^T$, while minimizing effort and reward costs. For head-fixed saccades, the matrix C is the identity matrix and the motor costs $\lambda^{(i)} = 1$. The only unknown parameters are the accuracy cost ν and noise κ. To find these parameters, we considered a 50° saccade, which has a peak velocity of ~450°/s. The parameters that reproduced such a movement are $\nu = [5 \times 10^9\ 1 \times 10^6\ 50]$, and $\kappa = 0.0075$. These parameters were then kept constant for all simulations here. Given these parameters, we searched for reward costs that reproduced the durations of saccades. This search involved the two parameters of the reward cost function: α and β in Equation 2.

To simulate saccades of children and other special populations, we varied α, which in turn altered both the value of the stimulus and the rate of reward discounting as a function of movement duration. In other words, the parameter α is the only variable that we manipulated to generate saccades for various populations and conditions in this paper.

Eye–head coordination. People and some species of monkeys (e.g., rhesus and macaques) rarely move their eyes in isolation. Rather, to view a stimulus, we typically move both our eyes and our head. In this head-free setting, a stimulus at position g does not produce a displacement of the eyes by amount g. Rather, both the eyes and the head contribute to the movement. Therefore, in response to a given stimulus, the eyes move differently in the head-free versus the head-fixed conditions. To test the strength of our model, we asked whether our cost function could account for saccade kinematics in both the head-fixed and head-free conditions.

To model movements of the eyes and the head, we augmented the state vector $\mathbf{x}$ to include three new states associated with the third-order dynamics of the head. That is, $\mathbf{x} = [\mathbf{x}_1\ \mathbf{x}_2]^T$, where $\mathbf{x}_1$ is state of the eye and $\mathbf{x}_2$ is state of the head. The time constants for the head were 270, 150, and

10 ms, which we extrapolated from monkey data (Bizzi, 1974). In contrast to head-fixed condition in which the objective was to position the eye at the target, now our objective was to position gaze at the target (i.e., the sum of eye and head positions). This means that

$$
C = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}. \qquad (18)
$$

We kept the eye plant parameters unchanged from the head-fixed simulations. The addition of the head model required addition of one new parameter, the motor cost associated with the head. We assumed that the motor cost term λ_2 was larger for the head than for the eye λ_1 (4 for the head, 1 for the eye). As before, we varied the parameter α to investigate the relationship between stimulus value and movement kinematics.

Experimental methods. An interesting prediction of our theory is that delay of reward should alter saccade kinematics. Specifically, our theory predicts that if at saccade completion, the stimulus is not present until after some time delay, the time delay should act as a reward prediction error, discounting the value of the stimulus, reducing saccade velocities. To check this prediction, we recruited healthy volunteers ($n = 8$; mean age, 26; range, 18–39; two females) and asked them to make saccades to targets that appeared on the horizontal meridian at displacements of 30°. Our procedures were approved by the Johns Hopkins Institutional Review Board. We measured eye movements using a high-speed infrared camera (EyeLink 1000; SR Research), which sampled eye position at 1000 Hz. An experimental session consisted of 12 sets, with each set composed of 40 targets. Subjects were tested on two sessions, performed on different days. Stimuli were presented on a 19 inch CRT monitor (frame rate, 120 Hz) and viewed from a distance of 37 cm. All fixation and target points were red dots (0.3° in diameter) presented against a black background. In each session, a set began with a fixation point, as illustrated in Figure 4A. Targets were displayed at 15° to the left and right of center resulting in 30° saccades symmetric about center. After saccade onset, the target was either maintained on the screen (first 10 and last 10 targets of the set as well all trials in the 0 ms delay sets) or extinguished (middle 20 targets of the set). If the target was extinguished, on saccade completion it was redisplayed after a time delay Δ. This delay was constant within each set, and then reset to a randomly selected value for the next set. Five nonzero delays were explored on day 1, and five were explored on day 2, along with two sets of the zero-delay condition during each session. To analyze the data, we measured the within-subject change in saccade parameters between the trials in which there was a delay and the trials for which delay was zero. It is important to note that, in our task, there were no explicit rewards associated with the saccades. Indeed, no score or feedback of any kind was provided to the volunteers regarding their performance. They were simply instructed to look at the target.

Data analysis. Abnormal saccades were excluded from analysis using global criteria that were applied to all subjects: (1) saccade amplitude, <20° (67% of target displacement) and >35°; (2) saccade duration, <60 and >300 ms; (3) saccade reaction time, <100 or >400 ms. For each subject, outliers for amplitude and peak velocity, which are those outside of two times the interquartile range, were also removed. Overall, ~9% of saccades were excluded from analysis.

Results

There are two basic ideas that we want to test: (1) movement durations carry a hyperbolic cost for the brain, and (2) this cost arises because the duration of a movement is equivalent to a delay in acquisition of reward. To test these ideas, we will first compare a hyperbolic cost of time with other kinds of cost functions to ask how well it can account for movement kinematics. Next, we will link the hyperbolic cost of time to discounting of reward by showing that variations in how the brain represents reward appear to produce variations in kinematics of movements.

Hyperbolic costs versus other costs of time

Figure 1A, right panel, plots the cost for a 20° saccade under a hyperbolic cost of time. Short duration saccades have a large cost

because the penalties associated with inaccuracy and effort increase as saccade duration decreases. With increasing saccade duration, the cost of delaying the reward increases. According to our hypothesis, the optimum movement duration is one that balances the need to be accurate versus the need to maximize reward (i.e., minimize the devaluation associated with delaying the reward). To test our hypothesis, let us consider kinematics of saccades that result under a hyperbolic discounting function, and compare it with saccades that result from other functions that penalize time.

For example, consider a quadratic cost of time $J_p = \alpha p^2$, as shown in the left panel of Figure 1A. We see that, for both hyperbolic and quadratic costs, there are parameter values so that a 20° saccade will have its minimum total cost at ~85 ms (this is the duration of a typical 20° saccade). Therefore, there is nothing special about a hyperbolic cost, as any increasing function of time can account for the observed kinematics of a 20° saccade. However, if we consider a family of movements (i.e., all amplitudes), then the implications for the choice of cost becomes clear. A quadratic cost of time implies that the cost as a function of movement duration increases rapidly. Therefore, with a quadratic cost, there is little increased penalty when we compare movements of 50 and 100 ms in duration (Fig. 1A, red dotted lines), but much greater increased penalty when we compare movements of 350 and 400 ms in duration. In contrast, for a hyperbolic function, there is greater increase in cost for short-duration saccades than for long-duration saccades. That is, for a hyperbolic cost, as movement durations increase the sensitivity to passage of time decreases.

Indeed, with a hyperbolic cost of time, we can account for an important property of saccades: on average, the duration of a saccade as a function of amplitude grows faster than linearly (Collewijn et al., 1988). For a quadratic cost of time, increasing movement amplitudes produce smaller and smaller changes in saccade durations, as shown by the tick marks in Figure 1B. In contrast, for a hyperbolic cost, the increasing movement amplitudes accompany a faster than linear increase in saccade durations. Figure 2 summarizes this idea for three kinds of temporal costs: quadratic, linear, and hyperbolic. This figure includes data from Collewijn et al. (1988), as well as a line of best fit that Collewijn et al. (1988) computed for saccades of small amplitude. A quadratic temporal cost produces reasonable estimates of saccade parameters for small amplitudes but fails for larger amplitudes. The reason is that, with a quadratic cost, the rate of increase in the penalty increases with time. If we consider a linear cost of time, an approach that was used by Harris and Wolpert (2006), the rate of increase in the penalty is constant, and we can produce reasonable trajectories for small-amplitude saccades. However, as Figure 2 illustrates, a linear cost of time underestimates saccade durations for large-amplitude movements. Therefore, with a hyperbolic cost of time we can account for durations of both small as well as large-amplitude saccades, but not with linear or qua-
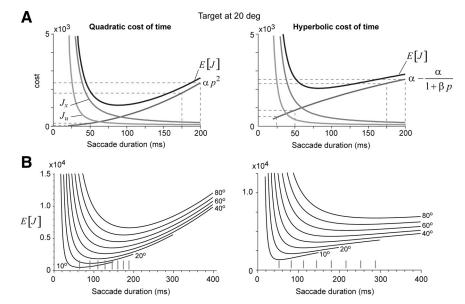

Figure 1. The cost of a saccade. Here, two forms of temporal discounting are considered: quadratic and hyperbolic. **A**, For a 20° saccade, the cost in Equation 5 is plotted as a function of movement duration p. Both quadratic and hyperbolic costs of time can produce a total cost that has a minimum at ~85 ms. **B**, Expected value of the cost as a function of movement duration. Each curve is the cost for a movement of constant amplitude. The curves are drawn for saccade amplitudes in the range 10–80°, by intervals of 10°. The tick marks near the x-axis are the optimal durations (i.e., the movement durations that produce minimum cost). For quadratic cost of time, movement durations get closer to each other as movement amplitude increases. For a hyperbolic cost, the durations get farther apart as amplitudes increase. Quadratic discount parameter: $\alpha = 5.75 \times 10^4$. Hyperbolic discount parameters: $\alpha = 0.8 \times 10^4$ and $\beta = 2.5$.
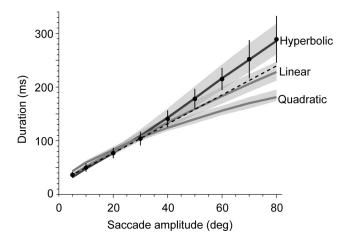

Figure 2. Effect of cost of time on movement durations. The data points are from Collewijn et al. (1988). The dashed line, also from Collewijn et al. (1988), is a good predictor of saccade durations in the range of 5–30° but underestimates durations for larger amplitudes. Quadratic $J_p = \alpha p^2$ or linear $J_p = \alpha p$ costs cannot account for the fact that saccade durations increase faster than linearly as a function of saccade amplitudes. The shaded areas along each curve represent the effect of changing stimulus value α by ±20%. The hyperbolic discounting not only accounts for the faster than linear increase in durations but also for the variability in this relationship: as stimulus value α changes, it has little effect on saccade durations for short amplitudes, but a greater effect for large amplitudes. Quadratic: $\alpha = 5.75 \times 10^4$. Linear: $\alpha = 1.2 \times 10^4$. Hyperbolic: $\alpha = 0.8 \times 10^4$ and $\beta = 2.5$. The red error bars are SD.

dratic costs of time. The fact that saccade durations increase faster than linearly is consistent with a hyperbolic cost of time.

Cost of time and temporal discounting of reward
Why should the brain impose a hyperbolic cost on duration of movements? The answer, in our opinion, is that this cost expresses how the brain temporally discounts reward. That is, the
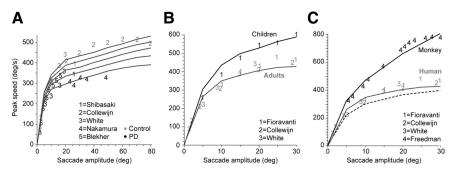
Figure 3. Change in the reward discount function predicts change in saccade velocities. The lines are simulation results, and the numbers refer to data from previous publications. For each line, the stimulus value α was kept constant. **A**, Saccade velocities in Parkinson's disease and healthy controls from data in the studies by Shibasaki et al. (1979), Collewijn et al. (1988), White et al. (1983), Blekher et al. (2000), and Nakamura et al. (1991). Reducing the stimulus value decreases saccade speeds. The changes in saccade speeds are bigger for large-amplitude saccades than small amplitudes. Parameter values are as follows: $\alpha = 0.52 \times 10^4$ to 1.08×10^4 and $\beta = 2.5$. **B**, Saccade velocities in children and young adults. Increasing the rate of discounting of reward (α in Eq. 6) by a factor of 2 produces saccade velocities that are similar to those seen in children. The data are from the studies by Fioravanti et al. (1995), Collewijn et al. (1988), and White et al. (1983). Parameter values are as follows: children, $\alpha = 2.16 \times 10^4$; adults, $\alpha = 1.08 \times 10^4$, $\beta = 2.5$. **C**, Saccade velocities in adult humans and rhesus monkeys. The dashed line represents simulations for which a rhesus monkey eye plant was combined with a human temporal discount function. The black line represents simulations for which a monkey eye plant was combined with a monkey temporal discount function ($\alpha = 6.5 \times 10^4$). For the human simulations, $\alpha = 1.08 \times 10^4$. The data on monkey saccades are from the study by Freedman (2008).

brain penalizes movement durations because passage of time delays the acquisition of reward. If this hypothesis is true, then it follows that movement kinematics should vary as a function of the amount of reward. For example, if we make a movement in response to a stimulus that promises little reward, α in Equation 1 is small and the motor and accuracy costs become relatively more important. As a consequence, when our brain assigns a low value to the stimulus, our movement toward that movement should be slow. To explore this idea, let us consider what happens to saccades when we alter the value of the stimulus α. Movement durations depend on the rate at which reward value is discounted in time. That is, movement duration depends on the derivate of cost J_p. This derivative is as follows:

$$\frac{dJ_p}{d_p} = \frac{\alpha\beta}{(1 + \beta p)^2}. \quad (19)$$

As α increases, so does the derivative of the reward discount function. Therefore, the cost of time rises faster when the stimulus has a larger value. As a consequence, movements in response to stimuli that have larger value will have shorter durations, exhibiting higher velocities. For example, the opportunity to look at a face is a valued commodity, and physical attractiveness is a dimension along which value rises (Hayden et al., 2007). As α increases, durations of simulated saccades decrease (as shown by the lower bound of the "error bars" in Fig. 2), resulting in higher velocities. This potentially explains why people make faster saccades to look at faces (Xu-Wilson et al., 2009).

A hyperbolic function is a good fit to discharge of dopamine cells in the brain of monkeys that have been trained to associate visual stimuli with delayed reward (Kobayashi and Schultz, 2008). That is, the response of these cells to stimuli is a good predictor of the temporally discounted value of these stimuli. In PD, many of the dopaminergic cells die. Let us hypothesize that this is reflected in a devaluation of the stimulus (i.e., a smaller than normal α). In Figure 3A, we have plotted velocity–amplitude data from a number of studies that have examined saccades of people with moderate to severe PD. The saccades of PD patients exhibit an intriguing property: the peak speeds are normal for small amplitudes but become much slower than normal for

large amplitudes. If we simply reduce stimulus value α, the model reproduces velocity–amplitude characteristics of PD patients (Fig. 3A).

Consider another curious fact regarding saccades: as we age, the kinematics of our saccades changes: children produce faster saccades than young adults (Fioravanti et al., 1995; Munoz et al., 2003). According to our theory, the differences in saccade kinematics should be a consequence of the way the child's brain temporally discounts reward. Green et al. (1999) measured the temporal discount rate of reward in both young children and adults and found that the initial slope of the discount function was two to three times larger in children than adults. That is, children discount reward more steeply than adults. They would rather take a single cookie now than wait for a brief period to receive two cookies. Figure 3B shows that, if we increase the slope of our temporal cost function (Eq. 19) by a factor of 2 (via parameter α), the resulting saccades share the velocity–amplitude relationship found in children's saccades.

As we age, saccade kinematics changes continuously so that by the time we reach our sixties, velocities are significantly lower than when we were in our twenties (Irving et al., 2006). Our theory accounts for this by noting that, as we age, the slope of the temporal discount function declines (Green et al., 1999).

In Table 1, we have summarized some of the data available on the rate of discounting of reward in various populations. We find a remarkable pattern: changes in saccade kinematics are generally consistent with the change in the rate of discounting of reward. For example, people with melancholic depression exhibit a steeper than normal temporal discounting of reward (Takahashi et al., 2008). Saccades in this population exhibit higher than normal velocities (Winograd-Gurvich et al., 2006). In schizophrenia, there is increased rate of temporal discounting (Klöppel et al., 2008), and this patient population also exhibits higher than normal saccade velocities (Mahlberg et al., 2001). In people who suffer from substance abuse, or people with gambling tendencies, there is increased impulsivity in tests that measure the rate of temporal discounting of reward (in conditions in which the subjects are not under influence of the substance). In all these cases, our theory predicts that saccade velocities will be higher than normal.

Let us now consider the fact that saccade velocities differ across species. For example, rhesus monkeys exhibit velocities that are approximately twice as fast as humans (Straube et al., 1997; Chen-Harris et al., 2008). One possibility is that this is attributable to interspecies differences in the eye plant. To check for this, we simulated saccades while taking into account eye dynamics of rhesus monkeys with a temporal discount function found in humans (Fig. 3C, dashed line). We found that the simulated monkey saccades were somewhat slower than in humans. Therefore, the differences in the eye plant did not appear to account for the differences in saccades. According to our theory, the differences in saccades should be related to interspecies differences in valuation of stimuli and temporal discounting of reward. Indeed, rhesus monkeys exhibit a much greater temporal discount rate: when making a choice between stimuli that promise

Table 1. Summary of data on the rate of discounting of reward in various populations

Condition	Discounting of reward (rate)	Saccades (peak velocity)
Schizophrenia	Increased (Heerey et al., 2007; Gold et al., 2008)	Increased (Mahlberg et al., 2001)
Melancholic depression	Increased (Takahashi et al., 2008)	Increased (Winograd-Gurvich et al., 2006)
Parkinson		Decreased (Nakamura et al., 1991)
DA-medicated Parkinson's disease patients	Increased (Voon et al., 2010)	Increased (Nakamura et al., 1991)
Huntington		Decreased (Lasker and Zee, 1997)
Low serotonin	Increased (Schweighofer et al., 2008)	Increased (Long et al., 2009)
Testosterone	Increased (Takahashi et al., 2006)	
Premenstrual syndrome		Decreased (Sundström and Bäckström, 1998)
Progesterone		Decreased (van Broekhoven et al., 2006)
Nicotine	Increased (Bickel et al., 1999; Reynolds et al., 2003; Ohmura et al., 2005)	
Cocaine, heroine, or amphetamine	Increased (Coffey et al., 2003; Kirby and Petry, 2004; Hoffman et al., 2006)	
Ecstasy	Increased (Morgan et al., 2006)	
Alcoholism	Increased (Mitchell et al., 2005)	
Pathological gamblers	Increased (Petry, 2001; Alessi and Petry, 2003; Dixon et al., 2003, 2006)	

reward (juice) over a range of tens of seconds, thirsty adult rhesus monkeys (Kobayashi and Schultz, 2008; Hwang et al., 2009) exhibit discount rates that are many times that of thirsty undergraduate students (Jimura et al., 2009). When we took into account this much faster temporal discount rate, our simulated monkey saccades had velocities (Fig. 3C) that were fairly consistent with the velocities that have been recorded from this species (Freedman, 2008).

It is noteworthy that among various species (Luhmann, 2009), pigeons exhibit some of the highest temporal discount rates (Green et al., 2004). Our theory suggests that their very fast, almost robotic-like movements are a reflection of this impulsivity.

Effect of delaying the reward

There are at least two shortcomings in the approach that we have taken in testing our theory: first, in experiments that are performed on humans, one is generally not explicitly rewarded for a saccade (i.e., one is not paid, given juice, etc.). The reader may doubt the idea that, in a darkened room, the brain would assign a value to a point of light that serves as the goal of the movement. A second problem is that we have used our theory to fit existing data, but we have not made predictions and tested the theory on new data. We designed an experiment to address both shortcomings.

Volunteers were asked to make a saccade to a visual stimulus (a point of light on a video monitor), as illustrated in Figure 4A. No explicit reward or performance measures were provided. Rather, the only manipulation was that on some blocks of trials the stimulus disappeared after saccade onset and then reappeared at a delay Δ after saccade end. Therefore, the saccade completed but the stimulus was not present. Based on our hypothesis, at trial onset the brain assigned a value to the target stimulus, and at saccade end this value had declined as specified by Equation 1. Because the saccade ended without the expected "reward," each trial induces a reward prediction error: if the movement completed at time p but the stimulus appeared at time $p + \Delta$, the reward prediction error is as follows:

$$V(p + \Delta) - V(p) = -\frac{\alpha\beta\Delta}{(1 + \beta p)(1 + \beta(p + \Delta))}.$$

(20)

We have plotted Equation 20 in Figure 4B. We see that the introduction of a delay will always produce a negative reward prediction error. More importantly, with increasing delay the reward prediction error tends to saturate.
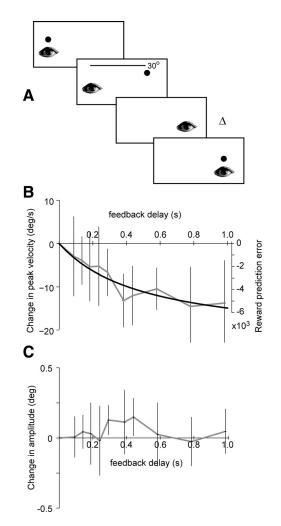

Figure 4. Delaying the stimulus discounts stimulus value. *A*, Experimental paradigm. Volunteers were asked to look at a stimulus, but after saccade initiation, the stimulus was removed. The stimulus was redisplayed at time Δ after saccade end. *B*, The black line is the theoretical estimate of reward prediction error (Eq. 17). Parameter values are as follows: $\alpha = 1.08 \times 10^4$, $\beta = 2.5$. Saccade duration is $p = 110$ ms. The data points are experimental results, showing within-subject change in peak saccade velocity with respect to the no-delay condition. The changes in saccade velocity are proportional to reward prediction error. *C*, Within-subject change in saccade amplitudes were uncorrelated with feedback delay. The horizontal and vertical error bars are SEM.

In response to a reward prediction error, the brain should update the value it assigns to the stimulus (i.e., it should devalue it because it did not receive the reward that it was expecting). Our value function is linear in α (Eq. 1). Therefore, an effective ap-

proach to minimize the reward prediction error is to update α by an amount proportional to the error as follows:

$$\alpha^{(n+1)} = \alpha^{(n)} + \frac{\eta}{1 + \beta p} (V(p + \Delta) - V(p)). \quad (21)$$

In Equation 21, the superscript refers to trial number and η is a learning rate (i.e., sensitivity to reward prediction error, which is unknown to us). Equation 21 predicts that the change in stimulus value should be proportional to the change in reward prediction error. Previously, we showed that as stimulus value decreased, so did saccade velocity. Importantly, for small changes in stimulus value the changes in velocity are proportional to changes in value. Therefore, our model makes two concrete predictions: (1) a delay in the availability of the stimulus with respect to movement completion will act as a reward prediction error, resulting in stimulus devaluation and reduced saccade velocities, and (2) the change in saccade velocities as a function of stimulus delay will be proportional to reward prediction error (i.e., Eq. 20).

Figure 4B plots the changes that we recorded in saccade kinematics of our volunteers as a consequence of delay Δ. We found that delaying the stimulus resulted in reduced saccade velocities (test for linear trend, $p < 0.05$), without producing consistent changes in saccade amplitudes (Fig. 4C) (no significant linear trend, $p = 0.56$). As the theory had predicted, the changes in saccade velocities were proportional to the function specified in Equation 20. That is, the changes in saccade velocities were proportional to the hypothetical reward prediction error.

Cost of time during eye–head movements

Does our theory generalize to other, more complicated movements? The movements that we have considered thus far are unusual in the sense that the head is kept fixed during the eye movement. In the natural setting, the brain responds to a stimulus at position g by moving both the eyes and the head. These head-free movements exhibit interesting characteristics: eye displacements grow slower than linearly as a function of g (Goossens and Van Opstal, 1997), whereas head displacements grow faster than linearly as a function of g (Guitton and Volle, 1987). Furthermore, duration of movement grows faster than linearly as a function of g (Epelboim et al., 1997). We wondered whether the hyperbolic cost of time could account for these coordinated movements.

To simulate head-free movements, we replaced the accuracy cost associated with eye position with an accuracy cost associated with gaze, where gaze is the sum of eye and head positions. We did not alter the parameters associated with this cost (i.e., kept T as before). Mathematically, our control problem is identical with one in which two arms cooperate to move a single cursor (Todorov and Jordan, 2002; Diedrichsen, 2007). Here, the eyes and head cooperate to move the fovea to the stimulus. However, unlike the two-arm situation, because of the substantially larger motor commands required to move the head than the eyes, the motor and accuracy costs ensure that the eyes lead the head, as shown for a simulated gaze change to a target at 45° in Figure 5A. Note that gaze is brought to the target through a combination of eye and head movements. Specifically, the eye contributions grow slower than linearly as a function of target displacement g, as shown in Figure 5B. These results are consequences of motor and accuracy costs and are generally unaffected by how we penalize time during a movement.

The cost of time, however, is strongly reflected in the relationships between gaze amplitude, duration, and velocity. To com-
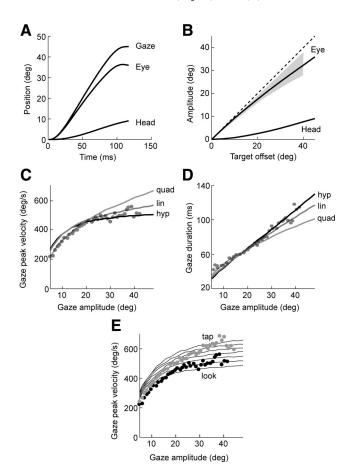

Figure 5. Kinematic characteristics of gaze appear consistent with a hyperbolic temporal discounting of reward. **A**, Simulation results for a gaze change to a target at 45°. Both the eyes and the head contribute to the gaze change, with the eyes leading the movement. **B**, Displacement of eye and head as a function of gaze amplitude. The gray region represents data from the study by Goossens and Van Opstal (1997). **C**, Peak gaze velocity as a function of gaze amplitude for three forms of temporal discounting. The data points are from the study by Epelboim et al. (1997). Parameter values are as follows: hyperbolic, $\alpha = 1.35 \times 10^4$ and $\beta = 2.5$; linear, $\alpha = 2.6 \times 10^4$; quadratic, $\alpha = 2.0 \times 10^5$. **D**, Gaze duration as a function of gaze amplitude. Parameters are same as in **C**. **E**, Effect of context on gaze velocities. The data points are from the study by Epelboim et al. (1997). The gray data points correspond to the tap task in which volunteers looked at the target that they were reaching for, and the black data points correspond to the look task, in which they only looked at the target. Simulation results of the hyperbolic model are shown by the lines. The stimulus value α was increased from $\alpha = 1.25 \times 10^4$ to $\alpha = 2.45 \times 10^4$.

pare our model with other costs of time, we once again considered two other functions that penalized time: linear and quadratic. For small-amplitude gaze changes, the three costs were indistinguishable in that they produced velocity–amplitude and duration–amplitude relationships that matched the available data (for example, 10–30° amplitudes) (Fig. 5C,D). However, as the movement amplitude increased, linear and quadratic costs tended to underestimate gaze duration and overestimate gaze velocity. This is a direct result of the fact that, with a hyperbolic cost of time, the incremental cost associated with increased movement duration becomes smaller as durations increase (i.e., the derivative of cost of time is decreasing). As a result, a hyperbolic cost once again reproduced velocity–amplitude–duration relationships for both short- and long-amplitude movements.

If the duration and speed of our actions are dictated by how we value the stimulus, then changing the context in which we view the stimulus might change the value we attribute to it. Contextual effects have been reported in control of gaze: when people are

asked to look and tap objects, they produce faster gaze changes than when they are asked to simply look at the objects (Epelboim et al., 1997). If we assume that, in the tap condition, the brain assigns a greater value α to the stimulus than in the look condition, then the model produces faster gaze velocities in the tap condition, as shown in Figure 5*E*. Our results suggest that the hyperbolic cost might be as relevant for eye–head movements as for eye movements alone.

State-dependent value of a stimulus

Finally, let us consider the curious fact that the kinematics of saccades to target of a reaching movement is affected by the load that one might impose on the arm. For example, the peak speed of a saccade is higher when there is a load that resists the reach, and lower when the load assists the reach (van Donkelaar et al., 2004). Why should varying the effort required to perform a reach to a target affect saccade velocities to that target?

Animals do not assign a value to a stimulus based on its inherent properties, but based on their own state when the stimulus was encountered. For example, birds that are initially trained to obtain equal rewards after either large or small effort, and are then offered a choice without the effort, generally choose the reward previously associated with the greater effort (Clement et al., 2000). The choice indicates a greater utility (i.e., relative usefulness, rather than absolute value) for the reward that was attained after a more effortful action. This phenomenon is called state-dependent valuation learning and is present in a wide variety of species from mammals to invertebrates (for review, see Pompilio and Kacelnik, 2010). In this framework, a reaching movement that is resisted by a load arrives at the target after a larger effort than one that is assisted. The more effortful state in which the reward is encountered favors assignment of a greater utility for that stimulus. This greater utility in our framework produces a faster saccade.

Discussion

Let us assume that our brain assigns a value to every part of the visible space, and each saccade is a voluntary movement with which the brain directs the fovea to a region where, currently, the value is highest. This framework naturally applies to the process with which the brain selects an action. However, the puzzling fact has been that the landscape of the value map also affects the motor commands that move the eyes. For example, saccades to faces are faster (Xu-Wilson et al., 2009), as are saccades to objects that are subject of a reach (Epelboim et al., 1997; Snyder et al., 2002; van Donkelaar et al., 2004). In monkeys, stimuli that promise greater reward result in saccades that have higher velocities (Takikawa et al., 2002). What is the link between the value that the brain assigns to a stimulus and the motor commands that it programs to acquire that stimulus?

We imagined that the objective of any voluntary movement is to place the body at a more valuable state. The value of the goal state is not static but is discounted in time. This forms an implicit cost of time (i.e., a penalty for the duration of our movement). To formulate this cost, we relied on experiments in which subjects were asked to choose between two amounts of reward: one that would be given to them now versus one that would be given later. These experiments measured time in years (Myerson and Green, 1995), or seconds (Jimura et al., 2009), and found that people's choices fit a hyperbolic function of time. Based on these results, we imagined that discounting of reward might remain hyperbolic even in the scale of milliseconds in which movements such as

saccades take place. Therefore, we imposed a cost on movement durations as a hyperbolic function of time.

Previous research had suggested that there are other costs associated with voluntary movements: a cost for accuracy (Harris and Wolpert, 1998) and a cost for effort (Todorov and Jordan, 2002; Izawa et al., 2008; O'Sullivan et al., 2009). To improve accuracy and minimize effort, one must slow the movement and increase its duration. However, if we also impose a cost of time based on temporal discounting of reward, then a natural balance arises between the desire to get as much reward as possible but be as lazy as possible. When we applied this idea to control of saccades, we found that the hyperbolic shape of temporal costs was essential to reproduce the velocity–duration–amplitude relationship found in saccades of healthy people.

A principal neuronal system involved in the encoding of reward is the dopamine system. Dopamine cells have a phasic discharge that varies hyperbolically with respect to stimuli that promise future reward (Kobayashi and Schultz, 2008). If a movement is required to obtain this reward, our theory indicates that the current value of this future reward should discount the motor costs. Indeed, a smaller phasic discharge of dopamine neurons precedes a slow reaching movement toward a food reward, whereas a larger discharge precedes a fast reaching movement toward the same reward [Ljungberg et al. (1992), their Tables 1, 2]. However, movement speeds are affected not only by the value of the reward predicted by the stimulus but also by the subject's global motivational state. Niv et al. (2007) suggested that the tonic discharge of dopamine neurons may encode the long-term average rate of reward per unit of time, discounting the effort needed to perform all actions. This model suggests that duration of a movement carries a cost because of missed opportunities to perform other actions. For example, it can account for the fact that hungry animals are more active, as well as more vigorous in each action. It is possible that tonic dopamine sets a baseline for reward per unit of time as applied for all actions, whereas phasic dopamine sets the reward per unit of time for the specific stimulus that affords the upcoming movement.

In Parkinson's disease, dopamine cells tend to die. Our inference that movements in PD are slow because of abnormally low temporal costs is in close agreement with results obtained in reaching movements of PD patients (Mazzoni et al., 2007). In that study, Mazzoni and colleagues demonstrated that PD patients do not move slowly because they are incapable of making fast and accurate movements: fast movements in PD are no more inaccurate than in healthy people. They speculated that slowness was related to a problem in how the PD brain evaluated effort, which is equivalent to an abnormally large L in Equation 7. In our formulation, slowness in PD arises because of an abnormally low stimulus value α. Mathematically, these two mechanisms produce fairly similar saccades in the small-amplitude range for which data are available.

If an abnormally small stimulus value can produce slow saccades, then an abnormally large value should produce fast saccades. In schizophrenia, saccade velocities are faster than in healthy controls (Mahlberg et al., 2001). Schizophrenia is a complex disease that likely involves dysfunction of generation and uptake of many neurotransmitters including dopamine, glutamate, and GABA. Stone et al. (2007) suggested that, in the striatum of schizophrenic patients, there is greater than normal dopamine synthesis. Kapur (2003) noted schizophrenics assign an unusually high salience to stimuli so that "every stimulus becomes loaded with significance and meaning." Indeed, all currently available antipsychotic medications have one common

feature: they block dopamine D_2 receptors. The reward temporal discount function in schizophrenia has a higher slope with respect to controls (Heerey et al., 2007; Klöppel et al., 2008), implying a greater discount rate. In our theory, this produces a faster rise in the cost of time, increasing saccade speeds.

Our inference is that the processes with which the brain temporally discounts reward are reflected in the kinematics of movements. Psychologists have quantified discounting of reward in diverse groups of patients and conditions, and physiologists have measured saccade kinematics of many of these same groups. Our theory suggests that there is a link between these two large bodies of science (Table 1).

The hyperbolic form of the reward discount function is favored by psychologists, whereas the exponential form is favored by economists and other theorist. Here, we chose the hyperbolic form because, empirically, it is a better fit to choices that animals make (Myerson and Green, 1995). However, in simulating saccades, the timescales are too short to allow us to dissociate between hyperbolic and exponential temporal discount functions. Reaching movements may provide a better way to test this dissociation.

Previous efforts in modeling voluntary movements such as head-fixed (Harris and Wolpert, 1998) and head-free saccades (Kardamakis and Moschovakis, 2009) had assumed a "best duration" for each movement amplitude. These models could not explain why movements are a particular duration. More recent efforts have suggested that duration of movements are linked to a desired level of endpoint accuracy (Tanaka et al., 2006), implying that faster movements that accompany more rewarding stimuli are attributable to a reduced accuracy cost. It is hard to see why eye movements should become less accurate when one is reaching for the stimulus versus when one is simply looking at the stimulus. Our proposed link between a cost of movement duration and temporal discounting of reward potentially resolves this issue.

Although there are physiological data that link our cost of time with temporal discounting of reward in the dopamine system (Kobayashi and Schultz, 2008), the dorsolateral prefrontal cortex (Kim et al., 2008), and the posterior parietal cortex (Louie and Glimcher, 2010), there is comparatively little known regarding the costs associated with effort and accuracy. Accuracy is a form of spatial cost, referring to a measure of distance between state of the eye and the rewarding state. As we move away from the fovea on the retina, the neuronal density drops exponentially, and as a result visual acuity drops exponentially. It is likely that, for saccades, spatial accuracy costs are not quadratic as we have assumed, but exponential. The implications of this idea remain to be explored. Furthermore, we assumed that cost of time interacts additively with other costs. An alternative, however, is to have cost of time multiplicatively interact with accuracy costs. This formulation does not produce satisfactory results with the current accuracy costs, suggesting the need for additional theoretical work.

In our theory, the cost of time during a movement depended on two parameters: the value α that the brain assigned the stimulus and the rate β that discounted this value in time. Our simulations here only varied α because this variation altered the rate of change in the cost of time, affecting velocities of small-amplitude saccades for which data are available in various populations. Importantly, for small-amplitude movements, it is difficult to dissociate the effect of α versus β. However, for large amplitudes, α alters the asymptotic velocities, whereas β has no effect on the asymptote. If we could develop robust techniques to measure α

and β in the reward function of individuals, it would be possible to test for within-subject correlations between the reward function and movement kinematics.

A prediction of our theory is that some of the interspecies differences that exist in movement kinematics may be attributable to differences in the cost of time arising from processing of reward. It will be useful to test our theory on different kinds of movements across various species and inquire about the evolutionary basis of temporal discount rates and its link to changes in motor control.

References

Alessi SM, Petry NM (2003) Pathological gambling severity is associated with impulsivity in a delay discounting procedure. Behav Processes 64:345–354.

Bickel WK, Odum AL, Madden GJ (1999) Impulsivity and cigarette smoking: delay discounting in current, never, and ex-smokers. Psychopharmacology (Berl) 146:447–454.

Bizzi E (1974) The coordination of eye-head movements. Sci Am 231:100–106.

Blekher T, Siemers E, Abel LA, Yee RD (2000) Eye movements in Parkinson's disease: before and after pallidotomy. Invest Ophthalmol Vis Sci 41:2177–2183.

Chen-Harris H, Joiner WM, Ethier V, Zee DS, Shadmehr R (2008) Adaptive control of saccades via internal feedback. J Neurosci 28:2804–2813.

Clement TS, Feltus JR, Kaiser DH, Zentall TR (2000) Work ethic in pigeons: reward value is directly related to the effort or time required to obtain the reward. Psychon Bull Rev 7:100–106.

Coffey SF, Gudleski GD, Saladin ME, Brady KT (2003) Impulsivity and rapid discounting of delayed hypothetical rewards in cocaine-dependent individuals. Exp Clin Psychopharmacol 11:18–25.

Collewijn H, Erkelens CJ, Steinman RM (1988) Binocular co-ordination of human horizontal saccadic eye movements. J Physiol 404:157–182.

Diedrichsen J (2007) Optimal task-dependent changes of bimanual feedback control and adaptation. Curr Biol 17:1675–1679.

Dixon MR, Marley J, Jacobs EA (2003) Delay discounting by pathological gamblers. J Appl Behav Anal 36:449–458.

Dixon MR, Jacobs EA, Sanders S (2006) Contextual control of delay discounting by pathological gamblers. J Appl Behav Anal 39:413–422.

Epelboim J, Steinman RM, Kowler E, Pizlo Z, Erkelens CJ, Collewijn H (1997) Gaze-shift dynamics in two kinds of sequential looking tasks. Vision Res 37:2597–2607.

Fagg AH, Shah A, Barto AG (2002) A computational model of muscle recruitment for wrist movements. J Neurophysiol 88:3348–3358.

Fioravanti F, Inchingolo P, Pensiero S, Spanio M (1995) Saccadic eye movement conjugation in children. Vision Res 35:3217–3228.

Freedman EG (2008) Coordination of the eyes and head during visual orienting. Exp Brain Res 190:369–387.

Fuchs AF, Scudder CA, Kaneko CR (1988) Discharge patterns and recruitment order of identified motoneurons and internuclear neurons in the monkey abducens nucleus. J Neurophysiol 60:1874–1895.

Gold JM, Waltz JA, Prentice KJ, Morris SE, Heerey EA (2008) Reward processing in schizophrenia: a deficit in the representation of value. Schizophr Bull 34:835–847.

Goossens HH, Van Opstal AJ (1997) Human eye-head coordination in two dimensions under different sensorimotor conditions. Exp Brain Res 114:542–560.

Green L, Myerson J, Ostaszewski P (1999) Discounting of delayed rewards across the life span: age differences in individual discounting functions. Behav Processes 46:89–96.

Green L, Myerson J, Holt DD, Slevin JR, Estle SJ (2004) Discounting of delayed food rewards in pigeons and rats: is there a magnitude effect? J Exp Anal Behav 81:39–50.

Guitton D, Volle M (1987) Gaze control in humans: eye-head coordination during orienting movements to targets within and beyond the oculomotor range. J Neurophysiol 58:427–459.

Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. Nature 394:780–784.

Harris CM, Wolpert DM (2006) The main sequence of saccades optimizes speed-accuracy trade-off. Biol Cybern 95:21–29.

Hayden BY, Parikh PC, Deaner RO, Platt ML (2007) Economic principles motivating social attention in humans. Proc Biol Sci 274:1751–1756.

Heerey EA, Robinson BM, McMahon RP, Gold JM (2007) Delay discounting in schizophrenia. Cogn Neuropsychiatry 12:213–221.

Hoffman WF, Moore M, Templin R, McFarland B, Hitzemann RJ, Mitchell SH (2006) Neuropsychological function and delay discounting in methamphetamine-dependent individuals. Psychopharmacology (Berl) 188:162–170.

Hwang J, Kim S, Lee D (2009) Temporal discounting and inter-temporal choice in rhesus monkeys. Front Behav Neurosci 3:9.

Irving EL, Steinbach MJ, Lillakas L, Babu RJ, Hutchings N (2006) Horizontal saccade dynamics across the human life span. Invest Ophthalmol Vis Sci 47:2478–2484.

Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. J Neurosci 28:2883–2891.

Jimura K, Myerson J, Hilgard J, Braver TS, Green L (2009) Are people really more patient than other animals? Evidence from human discounting of real liquid rewards. Psychon Bull Rev 16:1071–1075.

Kapur S (2003) Psychosis as a state of aberrant salience: a framework linking biology, phenomenology, and pharmacology in schizophrenia. Am J Psychiatry 160:13–23.

Kardamakis AA, Moschovakis AK (2009) Optimal control of gaze shifts. J Neurosci 29:7723–7730.

Keller EL (1973) Accommodative vergence in the alert monkey. Motor unit analysis. Vision Res 13:1565–1575.

Kim S, Hwang J, Lee D (2008) Prefrontal coding of temporally discounted values during intertemporal choice. Neuron 59:161–172.

Kirby KN, Petry NM (2004) Heroin and cocaine abusers have higher discount rates for delayed rewards than alcoholics or non-drug-using controls. Addiction 99:461–471.

Klöppel S, Draganski B, Golding CV, Chu C, Nagy Z, Cook PA, Hicks SL, Kennard C, Alexander DC, Parker GJ, Tabrizi SJ, Frackowiak RS (2008) White matter connections reflect changes in voluntary-guided saccades in pre-symptomatic Huntington's disease. Brain 131:196–204.

Kobayashi S, Schultz W (2008) Influence of reward delays on responses of dopamine neurons. J Neurosci 28:7837–7846.

Lasker AG, Zee DS (1997) Ocular motor abnormalities in Huntington's disease. Vision Res 37:3639–3645.

Ljungberg T, Apicella P, Schultz W (1992) Responses of monkey dopamine neurons during learning of behavioral reactions. J Neurophysiol 67:145–163.

Long AB, Kuhn CM, Platt ML (2009) Serotonin shapes risky decision making in monkeys. Soc Cogn Affect Neurosci 4:346–356.

Louie K, Glimcher PW (2010) Separating value from choice: delay discounting activity in the lateral intraparietal area. J Neurosci 30:5498–5507.

Luhmann CC (2009) Temporal decision-making: insights from cognitive neuroscience. Front Behav Neurosci 3:39.

Mahlberg R, Steinacher B, Mackert A, Flechtner KM (2001) Basic parameters of saccadic eye movements—differences between unmedicated schizophrenia and affective disorder patients. Eur Arch Psychiatry Clin Neurosci 251:205–210.

Mazzoni P, Hristova A, Krakauer JW (2007) Why don't we move faster? Parkinson's disease, movement vigor, and implicit motivation. J Neurosci 27:7105–7116.

Mitchell JM, Fields HL, D'Esposito M, Boettiger CA (2005) Impulsive responding in alcoholics. Alcohol Clin Exp Res 29:2158–2169.

Morgan MJ, Impallomeni LC, Pirona A, Rogers RD (2006) Elevated impulsivity and impaired decision-making in abstinent Ecstasy (MDMA) users compared to polydrug and drug-naive controls. Neuropsychopharmacology 31:1562–1573.

Munoz DP, Armstrong IT, Hampton KA, Moore KD (2003) Altered control of visual fixation and saccadic eye movements in attention-deficit hyperactivity disorder. J Neurophysiol 90:503–514.

Myerson J, Green L (1995) Discounting of delayed rewards: models of individual choice. J Exp Anal Behav 64:263–276.

Nakamura T, Kanayama R, Sano R, Ohki M, Kimura Y, Aoyagi M, Koike Y (1991) Quantitative analysis of ocular movements in Parkinson's disease. Acta Otolaryngol Suppl 481:559–562.

Niv Y, Daw ND, Joel D, Dayan P (2007) Tonic dopamine: opportunity costs and the control of response vigor. Psychopharmacology (Berl) 191:507–520.

Ohmura Y, Takahashi T, Kitamura N (2005) Discounting delayed and probabilistic monetary gains and losses by smokers of cigarettes. Psychopharmacology (Berl) 182:508–515.

O'Sullivan I, Burdet E, Diedrichsen J (2009) Dissociating variability and effort as determinants of coordination. PLoS Comput Biol 5:e1000345.

Petry NM (2001) Pathological gamblers, with and without substance use disorders, discount delayed rewards at high rates. J Abnorm Psychol 110:482–487.

Pompilio L, Kacelnik A (2010) Context-dependent utility overrides absolute memory as a determinant of choice. Proc Natl Acad Sci U S A 107:508–512.

Reynolds B, Karraker K, Horn K, Richards JB (2003) Delay and probability discounting as related to different stages of adolescent smoking and non-smoking. Behav Processes 64:333–344.

Robinson DA, Gordon JL, Gordon SE (1986) A model of the smooth pursuit eye movement system. Biol Cybern 55:43–57.

Schweighofer N, Bertin M, Shishida K, Okamoto Y, Tanaka SC, Yamawaki S, Doya K (2008) Low-serotonin levels increase delayed reward discounting in humans. J Neurosci 28:4528–4532.

Shibasaki H, Tsuji S, Kuroiwa Y (1979) Oculomotor abnormalities in Parkinson's disease. Arch Neurol 36:360–364.

Snyder LH, Calton JL, Dickinson AR, Lawrence BM (2002) Eye-hand coordination: saccades are faster when accompanied by a coordinated arm movement. J Neurophysiol 87:2279–2286.

Stone JM, Morrison PD, Pilowsky LS (2007) Glutamate and dopamine dysregulation in schizophrenia—a synthesis and selective review. J Psychopharmacol 21:440–452.

Straube A, Fuchs AF, Usher S, Robinson FR (1997) Characteristics of saccadic gain adaptation in rhesus macaques. J Neurophysiol 77:874–895.

Sundström I, Bäckström T (1998) Patients with premenstrual syndrome have decreased saccadic eye velocity compared to control subjects. Biol Psychiatry 44:755–764.

Takahashi T, Sakaguchi K, Oki M, Homma S, Hasegawa T (2006) Testosterone levels and discounting delayed monetary gains and losses in male humans. Neuro Endocrinol Lett 27:439–444.

Takahashi T, Oono H, Inoue T, Boku S, Kako Y, Kitaichi Y, Kusumi I, Masui T, Nakagawa S, Suzuki K, Tanaka T, Koyama T, Radford MH (2008) Depressive patients are more impulsive and inconsistent in intertemporal choice behavior for monetary gain and loss than healthy subjects–an analysis based on Tsallis' statistics. Neuro Endocrinol Lett 29:351–358.

Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. Exp Brain Res 142:284–291.

Tanaka H, Krakauer JW, Qian N (2006) An optimization principle for determining movement duration. J Neurophysiol 95:3875–3886.

Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. Nat Neurosci 5:1226–1235.

van Broekhoven F, Bäckström T, Verkes RJ (2006) Oral progesterone decreases saccadic eye velocity and increases sedation in women. Psychoneuroendocrinology 31:1190–1199.

van Donkelaar P, Siu KC, Walterschied J (2004) Saccadic output is influenced by limb kinetics during eye-hand coordination. J Mot Behav 36:245–252.

Voon V, Reynolds B, Brezing C, Gallea C, Skaljic M, Ekanayake V, Fernandez H, Potenza MN, Dolan RJ, Hallett M (2010) Impulsive choice and response in dopamine agonist-related impulse control behaviors. Psychopharmacology (Berl) 207:645–659.

White OB, Saint-Cyr JA, Tomlinson RD, Sharpe JA (1983) Ocular motor deficits in Parkinson's disease. II. Control of the saccadic and smooth pursuit systems. Brain 106:571–587.

Winograd-Gurvich C, Georgiou-Karistianis N, Fitzgerald PB, Millist L, White OB (2006) Ocular motor differences between melancholic and non-melancholic depression. J Affect Disord 93:193–203.

Xu-Wilson M, Zee DS, Shadmehr R (2009) The intrinsic value of visual information affects saccade velocities. Exp Brain Res 196:475–481.