

Vigor of Movements and the Cost of Time in Decision Making

Jennie E. S. Choi,^{1*} Pavan A. Vaswani,^{2*} and Reza Shadmehr¹

Departments of ¹Biomedical Engineering and ²Neuroscience, Laboratory for Computational Motor Control, Johns Hopkins University School of Medicine, Baltimore, Maryland 21205

If we assume that the purpose of a movement is to acquire a rewarding state, the duration of the movement carries a cost because it delays acquisition of reward. For some people, passage of time carries a greater cost, as evidenced by how long they are willing to wait for a rewarding outcome. These steep discounters are considered impulsive. Is there a relationship between cost of time in decision making and cost of time in control of movements? Our theory predicts that people who are more impulsive should in general move faster than subjects who are less impulsive. To test our idea, we considered elementary voluntary movements: saccades of the eye. We found that in humans, saccadic vigor, assessed using velocity as a function of amplitude, was as much as 50% greater in one subject than another; that is, some people consistently moved their eyes with high vigor. We measured the cost of time in a decision-making task in which the same subjects were given a choice between smaller odds of success immediately and better odds if they waited. We measured how long they were willing to wait to obtain the better odds and how much they increased their wait period after they failed. We found that people that exhibited greater vigor in their movements tended to have a steep temporal discount function, as evidenced by their waiting patterns in the decision-making task. The cost of time may be shared between decision making and motor control.

Key words: impulsivity; motor control; reward; saccade; temporal discounting; vigor

Introduction

Among healthy people, there are similarities in how we walk, reach, or move our eyes. To explain these regularities, theories have suggested that the nervous system produces motor commands to minimize metabolic costs (Hoyt and Taylor, 1981; Willis et al., 2005) or kinematic variability (Harris and Wolpert, 1998). Yet these theories cannot explain the fact that people (Xu-Wilson et al., 2009) and other primates (Kawagoe et al., 1998; Takikawa et al., 2002; Opris et al., 2011) move sooner or faster when there is an opportunity to acquire a greater amount of reward. For example, people produce saccades that have higher velocities in environments that offer greater rate of reward (Haith et al., 2012), and walk faster in cities that have a larger population (Bornstein and Bornstein, 1976). These observations suggest that in addition to efficiency and variability, the reward landscape affects the speed with which we move (Niv et al., 2007).

In principle, why should reward affect speed of movements? If we assume that the purpose of any movement is to arrive at a more rewarding state, then movement duration carries a cost, because passage of time discounts reward; that is, it is better to receive the reward sooner rather than later. Therefore, a move-

ment that takes longer to complete produces a greater devaluation of reward. Motor commands that guide a movement may be a balance between a desire to reduce inaccuracy (move slowly and improve precision) and a desire to maximize reward (move quickly and get reward sooner; Shadmehr et al., 2010).

Suppose that we have two subjects who have similar biomechanics, but who temporally discount reward differently. The theory predicts that the subject who discounts reward steeply should generally move faster (Shadmehr and Mussa-Ivaldi, 2012). Indeed, in populations in which development or disease affects temporal discounting, there are between-population differences in saccade velocity (Shadmehr et al., 2010). The critical question, however, is whether the between-subject differences in movement are related to between-subject differences in discounting of reward.

A temporal discount function can be measured in scenarios in which subjects compare a rewarding state that can be attained soon, with a more rewarding state that can be attained later (Millar and Navarick, 1984; Myerson and Green, 1995). For example, suppose that you purchased a new device and are offered a choice: you may have your device now, or you may wait a day and get the device engraved with your name. Which one would you prefer? The person with a steeper discount function would forgo the engraving and take the device home today.

Here, we measured saccadic eye movements and observed that some people moved their eyes with a peak velocity that was 50% faster than others. This difference was consistent in repeated measurements, appearing to be a trait. We then estimated temporal discounting in a decision-making task in which people decided how long to wait to improve their odds of success. In our

Received July 1, 2013; revised Oct. 28, 2013; accepted Nov. 19, 2013.

Author contributions: J.E.S.C. performed research; J.E.S.C., P.A.V., and R.S. analyzed data; P.A.V. and R.S. wrote the paper.

This work was supported by NIH Grant NS078311 and the Human Frontiers Science Program.

*J.E.S.C. and P.A.V. contributed equally to this work.

Correspondence should be addressed to Reza Shadmehr, 419 Traylor Building, 720 Rutland Avenue, Baltimore, MD 21205. E-mail: shadmehr@jhu.edu.

DOI:10.1523/JNEUROSCI.2798-13.2014

Copyright © 2014 the authors 0270-6474/14/341212-12\$15.00/0

task, every choice resulted in a real and immediate consequence, reinforcing the choice and affecting subsequent choices. We found that people with faster movements, as evidenced by saccade velocities, also tended to have a steep temporal discount function, as evidenced by the shorter periods of time they chose to wait to obtain additional reward.

Materials and Methods

Each subject sat in a darkened room in front of a CRT monitor (36.5 × 27.5 cm, 1024 × 768 pixel, light gray background, 120 Hz frame rate) with their head restrained using a dental bite bar. Visual targets (black; diameter, 1°) were presented on a CRT monitor with Matlab 7.4 (MathWorks) using Psychophysics Toolbox 3. The screen was placed at a distance of 31 cm from the subject's face, and an EyeLink 1000 (SR Research) infrared camera recording system (sampling rate, 1000 Hz) was used to record movement of the right eye. The experiments were approved by the Johns Hopkins Institutional Review Board. Volunteers were healthy with no known neurological disorders.

We wished to answer two questions: (1) How much did movement vigor, as measured by peak saccade velocity as a function of amplitude, vary across healthy individuals? And (2) was an individual's temporal discounting of reward as measured in a decision-making task a predictor of that individual's movement vigor? Twenty-three volunteers (14 females, 26.9 ± 6.8 years old, mean ± SD) participated in our two-part study, which was conducted on two separate days.

Part 1: movement vigor. In this part of the experiment, we wished to determine the range of movement speeds across our population of healthy individuals. We measured the kinematics of saccadic eye movements and determined the within-subject reproducibility of these movements and between-subject differences. Targets that were 5, 10, 15, 20, 25, 30, 35, or 40° apart on the horizontal axis were presented on a CRT monitor, centered on the midline of the right eye. Target amplitudes were ordered pseudorandomly in a blockwise fashion. Each target was presented 30 times in a row, resulting in 29 saccades. (We discarded the first saccade because this saccade was from a midline location to the first target and therefore was half the target amplitude.)

A trial began with display of a fixation spot. Our instructions were as follows: "A sequence of targets will appear on the screen. Please look at each target and maintain fixation until you see the next target." Each target was displayed for 1 s plus a random time distributed uniformly over −100 to 100 ms (Fig. 1A). Appearance of the target acted as a go cue. We did not enforce any gaze precision requirements; subjects received only the verbal instruction to look at the targets. The subjects received a short break after completion of two target amplitude blocks.

To assess the reproducibility of our results, five of the subjects were examined repeatedly on this task on 4 separate days. To ensure that time of day was not a factor, we selected four test times during the day (ranging from early morning to late afternoon) and tested each of these five subjects once on each test time.

To define vigor, we considered the peak velocity of saccades as a function of saccade amplitude. We measured amplitude via end point displacement of the eye, with positive displacement indicating temporal saccades and negative displacement indicating nasal saccades. Saccade peak velocity tends to increase as amplitude increases and saturates around 30°. As we will show, the between-subject differences in the velocity–amplitude function is accurately summarized by a scaling factor. Let us label the across-subject mean of the velocity–amplitude function as $g(x)$, where x is end point displacement and $g(x)$ describes the relationship between displacement and average velocity across the population:

$$g(x) = E[v(x)]. \quad (1)$$

In Equation 1, $E[\]$ is the expected value operator, computing the across-subject mean of the velocity–amplitude relationship. We will show that each subject's velocity–amplitude relationship is a scaled version of this function; that is, for subject i , peak velocity at displacement x is described by the following:

$$v_i(x) = \alpha_i g(x). \quad (2)$$

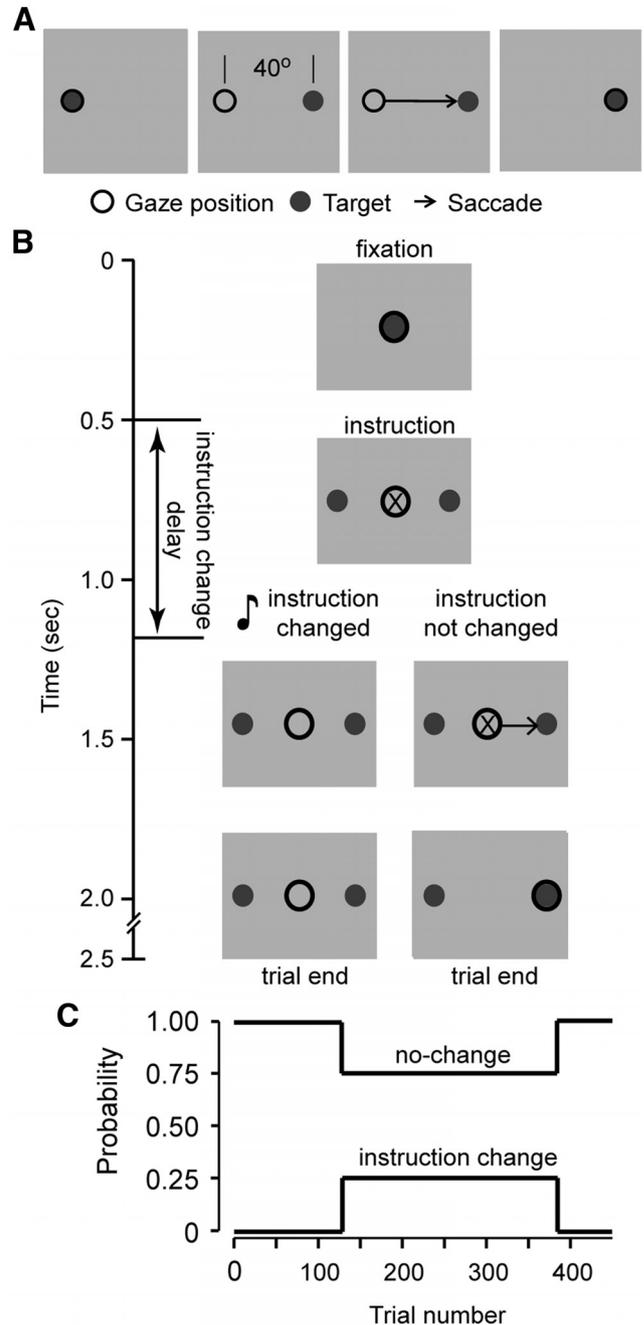


Figure 1. Experimental protocol. **A**, Part 1 of the experiment: measurement of saccade vigor. The trial began with a fixation spot of 0.5°, and then presentation of a target spot of 0.5° at a given displacement along the horizontal axis. Targets were presented for 1 s plus or minus a random time period. The targets were centered about the midline of the right eye. **B**, Part 2 of the experiment: measurement of temporal discount function. The trial began with a central fixation spot. Two targets were presented at 20° from fixation along with an instruction at the fixation spot indicating which target was the direction of the correct saccade. In Blocks 3–6 s, there was a 25% probability that following a variable delay period a second instruction would be given, indicating that the previously instructed saccade should be canceled. The delay period was adaptively adjusted to the success and failure of the subject on previous trials: success made the delay period 30 ms longer. The experiment attempted to measure the length of time the subject was willing to wait to improve their probability of success. **C**, Schedule of instruction probabilities in Part 2 of the experiment.

The scale factor α_i is our proxy for vigor of saccades for subject i . Saccade beginning and end were marked using a 30°/s velocity threshold (held for at least 4 ms). We used the following criteria to accept a saccade: no blinking during the saccade, displacement of <100°, and peak velocity of <1500°/s.

Part 2: temporal discount function. There are two classes of experiments that are used to measure temporal discounting in humans (Navarick, 2004). In one class, subjects are presented with potentially rewarding outcomes, the resulting choice is measured, and the consequences are immediately applied. The key element of this “operant class” of experiments is that the choices have real and immediate consequences that are experienced before any other choices are made. These consequences act as reinforcements or punishments, which then affect the next choice. All experiments in animals and some experiments in humans (Jimura et al., 2009) are of this class.

In the “nonoperant” class of experiments, rewarding states are presented (often small amounts of money soon vs larger amounts later) and a choice is measured, but the consequences of that choice are not experienced before the next choice is made. This is because the delay associated with the two rewarding states is typically days or weeks (rather than seconds, as in the operant experiments). Furthermore, nearly all rewards are hypothetical. However, participants are sometimes instructed that a couple of their choices will be selected for real payment after the session. Importantly, because all decisions are made before the money is received, the reward or punishment is not a reinforcement that affects the subsequent choices that are made in the experiment.

Here, we designed an experiment to measure temporal discounting that relied on the operant procedure. Every choice produced immediate and real consequences that in case of success could positively reinforce the choice, and in case of failure could negatively reinforce the choice. Using a model (described later), we predicted how the consequence of each choice would affect the subsequent choice, and how this trial-to-trial effect would be a proxy for the steepness of the temporal discount function.

Let us explain our task first intuitively, and then in a mathematical framework (in Results). Imagine joining a line where one has to wait to experience an event. We join the line with a prior belief regarding how long we have to wait. In our task, we control this prior belief by manipulating the history of when that event takes place. As we wait in line, with the passage of time we update our expectation of how much longer we have to wait. At some point, we may decide that the waiting is not worthwhile and leave the line. According to our model, the time at which we abandon the line is the time at which the temporally discounted value of reward has reached and passed a local maximum. The time when we abandon the line is a measure that will act as a proxy for the steepness of the temporal discount function.

Our task is shown in Figure 1B. Subjects were instructed to look at the central fixation spot (0.5°) presented for 500 ms. Subjects were instructed as follows: “If the central fixation spot turns into an X, move your eyes to look at the target on the right. If the central fixation point turns into an O, move your eyes to look at the target on the left.” Next, we presented two visual targets of size 0.5° at $\pm 20^\circ$, along with an instruction at the fixation spot indicating which target the subject should saccade to: an “X” instructed a saccade to the right target, and an “O” instructed a saccade to the left target.

The experiment consisted of seven blocks of 64 trials. In the first two blocks, the subjects were told to respond to the center instruction by making a saccade to the appropriate target. In the first block, while subjects were learning the instruction, if a saccade was made in the wrong direction, the computer played a distinct tone to indicate an error had been made. After the first block, subjects were not given feedback regarding movement direction; however, they made saccades in the wrong direction on only $1.0 \pm 0.2\%$ (mean \pm SEM across subjects) of trials after the first block. Visual observation of the target and the error tone were the main sources of feedback in this task.

Before the start of the third block, the subjects were given new instructions: “For some of the trials, the first instruction may be followed, after a delay, by a tone [the second instruction]. Occurrence of this tone means that the first instruction has been canceled and replaced. In this case, you should continue fixation.” In Blocks 3–6, on 25% of the trials after a variable delay period, the instruction changed, signaled by a distinct sound. This instruction-change cue was different from the error tone. Success or failure on these trials was determined only by whether subjects responded to the instruction-change cue, and was independent of the

saccade direction. Therefore, if the subject followed the first instruction and made a saccade, and the instruction did not change, that trial was a success. If the subject followed the first instruction but the instruction changed, then the trial was a failure, and the error tone, the same as that from Block 1, was played. If the subject waited, maintaining fixation despite the first instruction, and subsequently the instruction changed, then the trial was a success. The only feedback was the success or failure of the current trial determined only by whether or not the subject made a saccade, indicated by the error tone. Making a saccade in the incorrect direction was not penalized, though this happened rarely. We did not provide scores regarding the number of successful trials or any other cumulative feedback. In the final block, the instruction did not change, but the subjects were not provided verbal information regarding this fact.

If one was to react only to the first instruction, then one was successful with 75% probability. Waiting for the second instruction improves the probability of success by 25%. How long would an individual be willing to wait to improve their odds? The variable of interest was the delay period that could be sustained by each individual. The instruction-change delay period started at 200 ms for all subjects. If on the instruction-change trial the subject was successful (i.e., the subject had waited), the instruction-change delay increased by 30 ms, requiring them to wait longer in the future. If on the instruction-change trial the subject failed, the instruction-change delay decreased by 30 ms. Therefore, with this adaptive algorithm we attempted to find how long the subject was willing to wait to acquire the greater odds of success. A formal analysis of this task is provided in Results.

Each trial was 2.5 s in duration. This duration was fixed regardless of events that occurred in that trial. In this way, both the subject that waited a brief period of time for the second instruction and the subject that waited a long period experienced the same total experiment time and the same overall rate of movement.

After completion of the task, subjects filled out two questionnaires that are commonly used to measure impulsivity as a psychological profile. These questionnaires are the Barratt Impulsiveness Scale (BIS-11; Patton et al., 1995) and the I7 Impulsiveness Questionnaire (Eysenck et al., 1985). For the I7 questionnaire, we did not use the components in the empathy category.

Modeling. We considered a model to describe the process of decision making in this task. This model is described in Equations 5–9 in Results. As a trial began, the model decided at what time it would move using its temporal discount function and expected probability of success, given the expected arrival time of the second instruction. As time progressed (in 1 ms increments), the model truncated its expectation of the probability of the time of the second instruction, updated its desired movement time, and, if the desired movement time was at the current time or sooner, responded to the first instruction, i.e., stopped waiting. Otherwise, the model waited. If the model moved before the arrival of the second instruction (a failed trial), the delay of the second instruction was reduced by 30 ms, as in our experiment. Otherwise (a successful trial), the delay was increased by 30 ms. Our model changed its estimate of $\hat{\Delta}$, the expected arrival time of the second instruction, only in trials in which there was a second instruction. Therefore, we simulated our model with 64 trials in which the instruction changed.

Results

In Part 1 of our experiment we asked whether there were consistencies in the saccade velocities of healthy individuals across several amplitudes. Using these velocities, we defined a measure of movement vigor for each subject. In Part 2 we asked whether an individual’s temporal discounting of reward, as measured in a decision-making task, was a predictor of that individual’s movement vigor.

Between-subject differences in movement velocities

Figure 2A shows the eye velocity trajectories of two representative subjects during saccades of various amplitudes. Saccade peak velocity and duration increased with amplitude in both subjects,

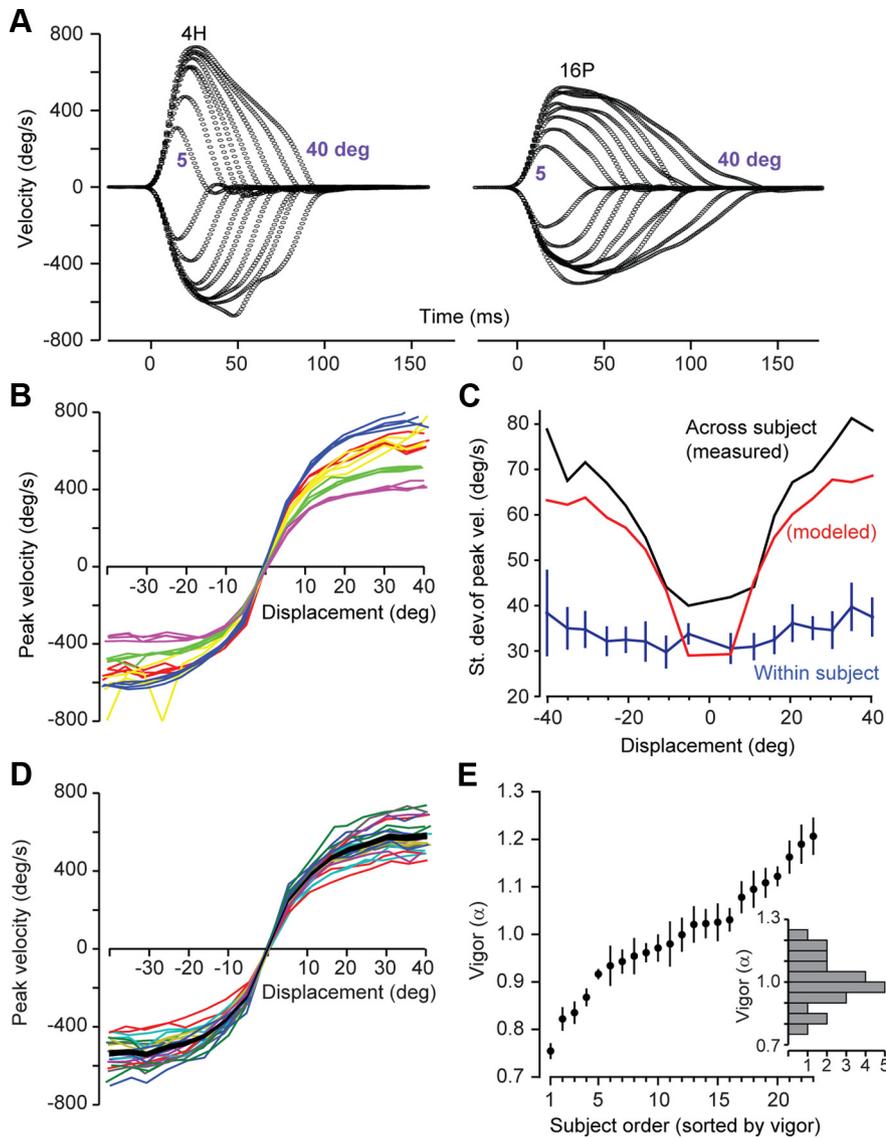


Figure 2. Vigor of saccades. **A**, Average eye velocity traces for horizontal saccades of various amplitudes for two representative subjects. Saccades were averaged for each amplitude in 5° increments, centered at 5 to 40°. Recordings are from the right eye. Negative velocities refer to nasal saccades. **B**, Peak velocity versus end point displacement of the eye during saccades for five subjects that were examined on four separate days. Each line represents data from one subject on 1 d, and each color is one subject. **C**, Within-subject variability of peak velocity and across-subject variability of peak velocity as functions of displacement. The error bars represent 1 SEM. The black line is the across-subject measured variability, and the red line is the variability accounted for using the model of Equation 2. Peak velocity was much less variable within a subject than across subjects. **D**, Peak velocity–displacement relationship for all subjects. Each line represents the data for a single subject. The thick line is the across-subject mean, and the black region is ± 1 SEM. The velocity–displacement relationship for each subject appears to be a multiplicative scaling of the mean function. **E**, The distribution of vigor of saccades, as defined in Equation 2, across the population of subjects. A vigor of 1 represents the mean of the population. Error bars are 95% confidence intervals in estimating vigor for each subject.

but for any given amplitude, Subject 4H had peak velocities that were higher than those of Subject 16P. One way to summarize these data is to consider peak velocity as a function of end point displacement for each subject. Figure 2B provides these data for five representative subjects, measured over 4 d. In this figure, each line represents data from one subject on 1 d.

We first asked whether there were significant between-subject differences in the saccade peak velocity–amplitude relationship. To determine whether the between-subject differences were statistically robust, we performed a one-way repeated measure ANOVA on peak velocity measurements where

displacement on each day was the within-subject factor and subject identity was the between-subject factor. We found significant effect of subject identity ($F_{(4,15)} = 22.5, p < 10^{-5}$), and a subject-by-displacement interaction ($F_{(60,225)} = 45.5, p < 10^{-9}$). This indicates that there were highly significant between-subject differences in the amplitude–velocity relationship of saccades: day after day, some subjects moved their eyes with reliably higher velocities than others.

To quantify the consistency of velocities within a subject, we asked whether across-subject peak velocities are more variable within or across subjects. We measured the SD of peak velocity at a given displacement for each subject. The resulting within-subject distribution is shown in Figure 2C (blue line). In comparison, consider the across-subject distribution of peak velocities (Fig. 2C, black line). The across-subject SD is about twice that of the within-subject SD. At all displacements, a *t* test produced a significant difference in the comparison of the within- and between-subject measures (in all cases, $p < 0.0015$). A Bonferroni-Holm correction of the *p* values for the $m = 16$ family of multiple comparisons demonstrated that the differences remained significant after this correction. Therefore, peak velocity was much less variable within a subject than across subjects. This implies that individuals have a characteristic, trait-like velocity with which they move their eyes.

Movement vigor

The data in Figure 2B suggest that the between-subject differences in velocity may be summarized by a scaling factor. To see this, consider the velocity–displacement relationship for all subjects, as shown in Figure 2D. The heavy black line is the across-subject mean of the data. Let us label this across-subject mean as $g(x)$, where x is end point displacement, and $g(x)$ describes the relationship between displacement and average velocity across the population. In Equation 2, we hypothesized that each subject’s velocity–dis-

placement relationship is a scaled version of this function, where the scaling factor α_i is our proxy for vigor of saccades for subject i . To test whether Equation 1 is an accurate representation of the data in Figure 2D, we found the parameter α_i that in a least-squares sense best fitted the amplitude–velocity data for subject i . Next, we used Equation 2 to predict the between-subject SD of saccade velocities as a function of displacement:

$$SD[v(x)] = SD[|\alpha|g(x)]. \tag{3}$$

In Equation 3, $||$ indicates absolute value, and $SD[]$ is the SD operator. We then compared the predicted SD–displacement re-

relationship of Equation 3 (modeled across subjects, red line; Fig. 2C) with the actual relationship (measured across subjects, black line; Fig. 2C). We found that the model and data correlated at $r = 0.94$, $F_{(1,14)} = 116$, and $p < 10^{-8}$. The goodness of fit of our model of vigor for each subject is shown by the confidence intervals in Figure 2E, and the resulting distribution of saccade vigor, i.e., parameter α , is shown in the inset of Figure 2E. Therefore, it appears that Equation 2 is a reasonable representation of the velocity–displacement data of each subject and the parameter α can be used as a measure of the vigor of each subjects’ eye movements.

We wondered whether differences in vigor are related to differences in end point variability; that is, do people who have higher vigor move their eyes with less accuracy? We compared mean SD of saccade end points across all target distances with vigor and found that increased vigor did not correspond to more variability ($F_{(1,21)} < 1$, $p = 0.99$). Therefore, accuracy is not a cost that can readily account for between-subject differences in vigor.

If there is a general cost of time for control of movements, then subjects that exhibit a greater vigor (and therefore a greater cost of time) may also exhibit a faster reaction time (RT). We observed a trend in this direction, but the trend was not statistically significant ($r = -0.28$, $F_{(1,21)} = 1.76$, $p = 0.20$); that is, people who had higher vigor did not react faster to a stimulus.

Estimating the temporal discount function

It is possible that between-subject differences in movement vigor are related to between-subject differences in the reward system of the brain (Shadmehr et al., 2010): populations that show increased saccade velocity may also exhibit increased rates of temporal discounting in decision-making tasks. For example, rhesus monkeys have saccade velocities that are about twice as fast as those of humans (Straube et al., 1997; Chen-Harris et al., 2008). Monkeys have eye biomechanics that are somewhat different from those of humans (Fuchs et al., 1988), but once these differences are accounted for, there remains persistent differences in movement vigor (Shadmehr et al., 2010). Intriguingly, monkeys exhibit a greater temporal discount rate: when making a choice between stimuli that promise juice over a range of tens of seconds, thirsty monkeys (Kobayashi and Schultz, 2008; Hwang et al., 2009) exhibit discounts rates that are higher than those of thirsty humans (Jimura et al., 2009).

Let us define temporal discounting as follows:

$$V(r, t_0 + t) = V(r, t_0) F(t). \tag{4}$$

The value of reward at current time t_0 , written as $V(r, t_0)$, is discounted by a function $F(t)$ to produce value at time $t_0 + t$, with $F(0) = 1$. Suppose Subject 1 is given a choice between a small amount of reward now (r, t_0) and a large amount of reward later ($r + R, t_0 + t$), and this subject picks the smaller reward. In comparison, Subject 2 is given the same choice, but picks the larger reward. In this choice, Subject 1 is more impulsive, preferring the sooner but smaller reward. Therefore, for Subject 1, $V_1(r, t_0) < V_1(r + R, t_0)F_1(t)$, whereas for Subject 2, $V_2(r, t_0) < V_2(r + R, t_0)F_2(t)$. If we assume that the two subjects value a given reward equally at the current time, i.e., $V_1(r, t_0) = V_2(r, t_0)$, then we infer that the temporal discount function of Subject 1 devalues reward more than Subject 2, $F_1(t) < F_2(t)$, which implies that Subject 1 is a steeper discounter. According to our hypothesis, Subject 1 should generally move with greater velocity than Subject 2.

To test this prediction, we designed a task to measure temporal discounting (Fig. 1B). A critical component of our task was that each choice produced an immediate and real consequence

(success or failure), which was experienced before the next choice. As we will see, the consequence of the choice affects the next choice, and this trial-to-trial change in behavior is related to the individual’s temporal discount function.

On each trial subjects were given instructions to make a movement. However, on some fraction of trials ϵ after some time delay Δ , there was a second instruction. On trials with only a single instruction, one was successful by following that instruction. On trials with a second instruction, one was successful only after waiting for that instruction. The subjects did not know whether a trial had one or two instructions. Only by waiting the subjects discovered the nature of the trial. The result of each trial, success or failure, reinforced the choice that was made.

The probability of success in a trial increased with waiting. The probability of success, given that the second instruction came at Δ seconds, was described by a logistic function:

$$\Pr(\text{success} | \Delta) = (1 - \epsilon) + \frac{\epsilon}{1 + \exp(-b(t - \tau - \Delta))}. \tag{5}$$

In Equation (5), ϵ is the fraction of trials in which there is a second instruction, τ is the amount of time it takes to respond to the second instruction (i.e., reaction time), t is the time at which the movement takes place, and b reflects the variance in the ability of a subject to reproduce the predicted timing of the second instruction. In this experiment, $\epsilon = 0.25$. In our simulations $\tau = 0.1$ was used, reflecting the approximate response time (~ 110 ms, see Reaction to the instruction-change cue, below) to the second instruction. We also used $b = 100$ Hz, which corresponds to an SD in the estimate of time of ~ 17 ms. This is similar to estimates of the SD of the ability of subjects to produce a time interval in prior work (17 ms for 400 ms intervals; Ivry and Hazeltine, 1995). We have plotted Equation 5 via a blue curve in Figure 3A. The longer one waited before making a movement, the higher the chances of success.

Suppose that from the history of previous trials, the subject estimates the time Δ that the second instruction will come. For example, in Figure 3A (top row) the subject expects that the second instruction will come at a time as shown by the red distribution, labeled as $p(\Delta|t)$. $\hat{\Delta}$ is the median of this distribution, and this is the best guess, at current time t , regarding when the second instruction will come. If the subject’s objective is to maximize the probability of success, then the subject would wait indefinitely on each trial. However, suppose time discounts reward such that we have the following:

$$F(t) = \frac{1}{1 + \beta t} \tag{6}$$

Equation 6 is a temporal discount function, representing hyperbolic discounting of reward. This function is shown by the green line in Figure 3A. Consider two hypothetical subjects: one who has a steep discount function (large $\beta = 0.58$), as shown in the left column of Figure 3A, and one who has a shallow discount function, as shown in the right column of Figure 3A (small $\beta = 0.2$). These two values of β were selected to illustrate the differences in the behavior of subjects with steep and shallow discount functions. Using probability of success and their personal temporal discount function, the two subjects choose the amount of time they are willing to wait so that they maximize the discounted value of success:

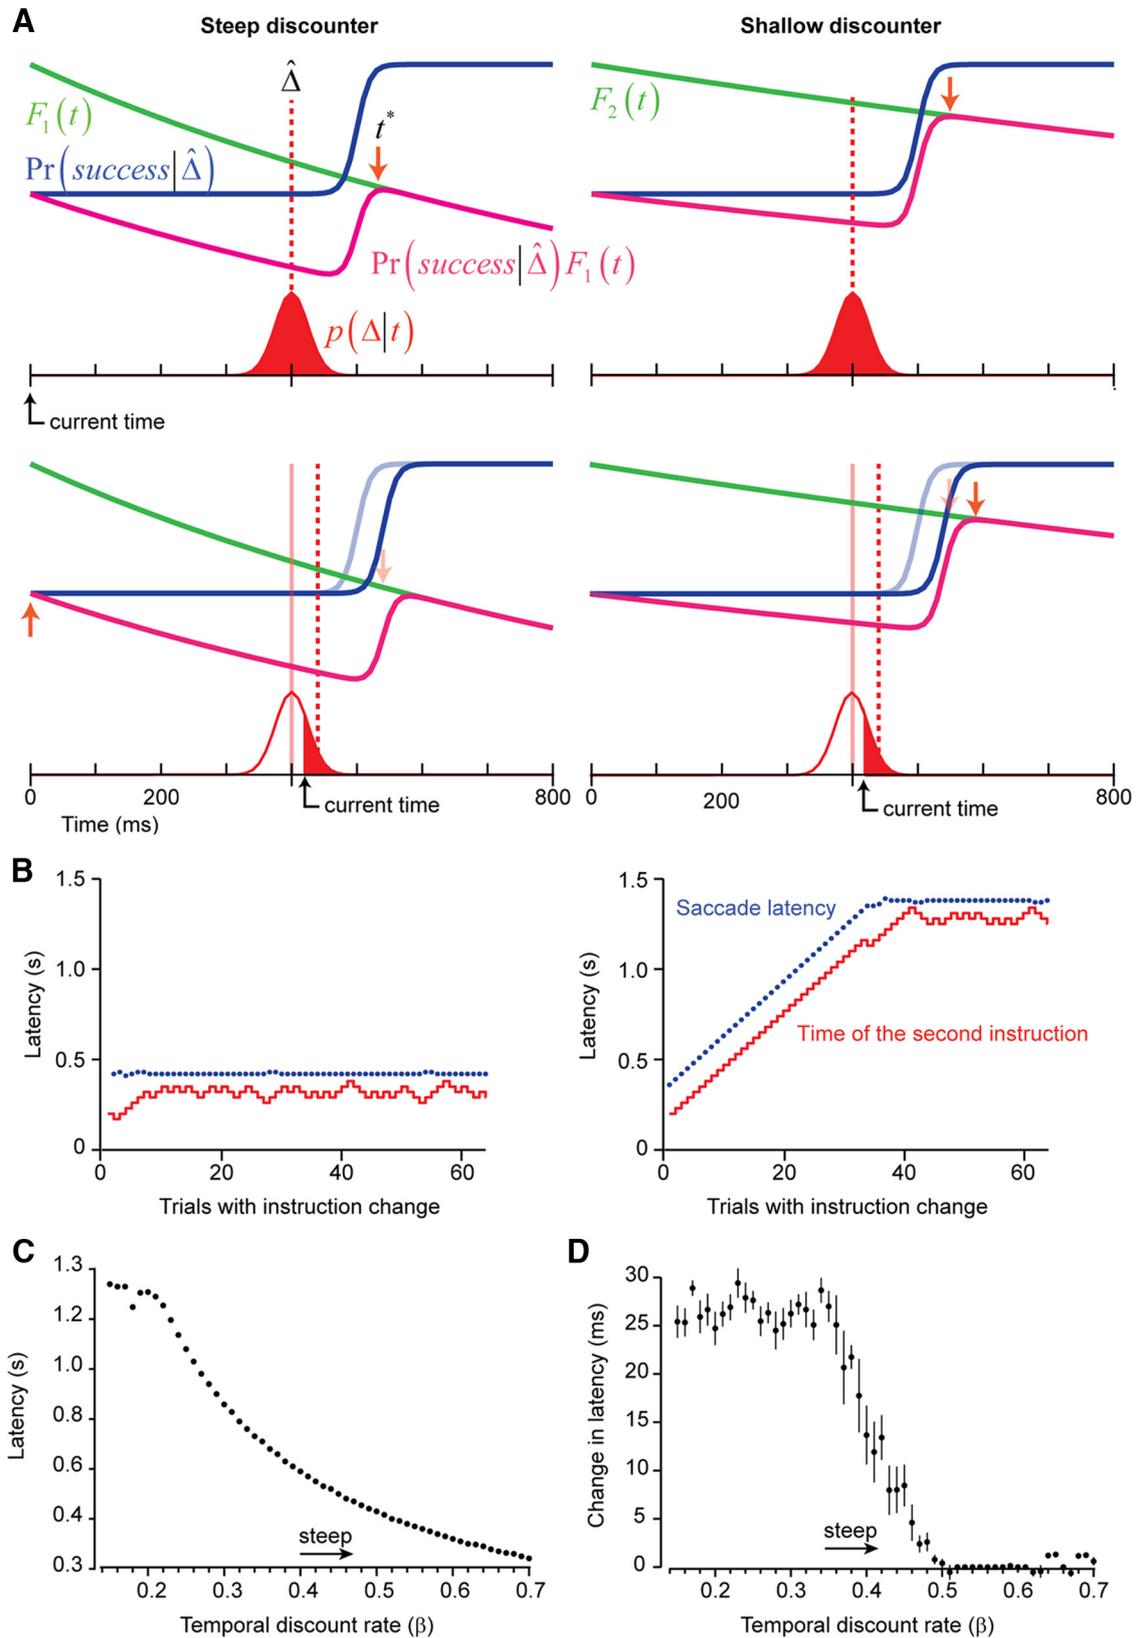


Figure 3. Performance of two simulated subjects (one a steep discounter and one a shallow discounter) in the decision-making task. **A**, Based on the history of previous trials and the current time t with in the trial, the subject estimates the time $\hat{\Delta}$ the second instruction will come. This is represented by the red distribution and is labeled as $p(\Delta|t)$. $\hat{\Delta}$ is the median of this distribution (dotted red line). The probability of success (blue line) is 0.75, and rises to 1.0 after $\hat{\Delta}$. The temporal discount function is represented by a hyperbola (green line) for each subject and discounts the probability of success to produce the value represented by the pink line. For these two subjects, $F_1(t) < F_2(t)$. As the trial starts, $t = 0$ (top row), both subjects estimate that the time of the second instruction will be $\hat{\Delta} = 400$ ms. For both subjects it is worthwhile to wait because the peak of the discounted probability $\Pr(\text{success}|\hat{\Delta})F_1(t)$ is in the future. The time of this peak is labeled t^* . As they wait and time passes, the probability distribution $p(\Delta|t)$ becomes truncated because the time of the second instruction cannot be in the past (second row). As a result, $\hat{\Delta}$ shifts to the right (dotted red line). This change causes a change in the probability of success (blue curve), as well as the discounted probability of success (red curve). After 400 ms of waiting, the steep (Figure legend continues.)

$$t^* = \arg \max \{ \Pr(\text{success} | \hat{\Delta}) F(t) \}. \tag{7}$$

The discounted value of success is plotted via the pink curve in the top row of Figure 3A, and the optimum wait time t^* is labeled with an arrow.

Let us illustrate how these two hypothetical subjects would behave on a given trial. As trial n starts, i.e., $t = 0$, suppose both subjects estimate that the second instruction will come at $\hat{\Delta} = 400$ ms; the median of the probability distribution $p(\Delta | t)$. The distribution $p(\Delta | t)$ was simulated as a Gaussian with a mean of 400 ms and SD of 25 ms. This SD is similar to the SD of subjects' perceptions of 400 ms time intervals in prior work [20.3 ms in the study by Ivry and Hazeltine (1995); ~32 ms in the study by Westheimer (1999)].

For this $\hat{\Delta}$, at $t = 0$ the optimum amount of time to wait is in the future (Eq. 7): the discounted value of success $\Pr(\text{success} | \hat{\Delta}) F_1(t)$ for both subjects has a maximum that lies in the future. Therefore, both subjects wait. As they wait and time passes, the probability distribution $p(\Delta | t)$ becomes truncated because the time of the second instruction cannot be in the past (Fig. 3A, second row). This means that as time passes in the trial and the subject waits, $\hat{\Delta}$ is not constant, but becomes larger, reflecting the median of the now truncated $p(\Delta | t)$:

$$\hat{\Delta}(t) = \text{Median} [p(\Delta | t)]. \tag{8}$$

This change in $\hat{\Delta}$ (dotted red line) causes a change in the probability of success (blue curve), which in turn produces a change in the discounted probability of success (pink curve). The second row of Figure 3A shows discounted value of success at $t = 410$ ms. At this time (i.e., 410 ms into the trial), for the impatient subject (steep discounter) the peak discounted value is no longer in the future, but is in the past (the red arrow is now at $t = 0$). The impatient subject stops waiting and initiates their movement, responding to the first instruction. The time of the movement represents the saccade latency of this subject. In contrast, for the person with the shallow discount function (patient discounter), at time $t = 410$ ms the discounted value of success has a maximum that is still in the future. This person will continue to wait.

In our task, the time of the second instruction, represented by Δ , was adjusted so that it tracked the amount of time each subject was willing to wait. If the second instruction occurred and the subject had waited for it (successful trial), Δ was increased by 30 ms. If the second instruction occurred and the subject had not waited for it (failed trial), Δ was decreased by 30 ms. Using the temporal discount functions shown in Figure 3A, we simulated the behavior of the two hypothetical subjects, as shown in Figure

3B. After a trial in which the second instruction occurred, regardless of success or failure, the model updated its expectation of the time of this event as follows, where, in our simulations, $\eta = 0.7$:

$$\hat{\Delta}^{(n+1)} = \hat{\Delta}^{(n)} + \eta (\hat{\Delta}^{(n)} - \hat{\Delta}^{(n)}). \tag{9}$$

In the simulated steep discounter, Δ reached a maximum of ~400 ms, whereas in the simulated shallow discounter, Δ reached a maximum of ~1400 ms. In Figure 3C, we have plotted the asymptotic value of latency for various temporal discount rates. The y -axis of this figure indicates the final saccade latency in the simulated experiment. The model suggests that at the end of the experiment, a subject that has a shallow temporal discount function will have a longer saccade latency, i.e., will wait longer to respond to the first instruction, than a subject that has a steep temporal discount function.

Whereas the simulations in Figure 3C describe the asymptotic behavior of our simulated subjects, the model also makes an interesting prediction with regard to trial-to-trial change in behavior, in particular near the start of the experiment. Suppose that on trial n , both of our hypothetical subjects predict that the second instruction will come at time $\hat{\Delta}$. Further suppose that in fact the second instruction comes at a time Δ , later than expected. The two subjects have the same prediction error, and both shift their estimate $\hat{\Delta}$ for the next trial, $n + 1$ (Eq. 9). Because of the shape of the discount functions, this change in $\hat{\Delta}$ produces a small trial-to-trial change in the latency for the steep discounter, but a larger change in latency for the shallow discounter. To illustrate this idea, we ran our model for various discount functions and focused on the latencies in the first block of trials with a second instruction. We computed how much the simulated subjects changed their latency in response to a trial in which the second instruction occurred but they did not wait for it; that is, we computed the change in behavior in response to a failed trial. The results are shown in Figure 3D. The model predicted that subjects with shallow temporal discount functions should respond to a failed trial with relatively large change in latency, whereas steep discounters should show a small change in latency.

Our simulations also illustrate that the change in latency in response to a failed trial is a steeper function of temporal discounting than asymptotic latency. For example, the ability of the model to distinguish a 0.35 discounter from a 0.45 discounter is about four times better with the change in latency measure compared to the asymptotic latency measure. This implies that for two people who are near the mean of the population, small differences in temporal discount rates will be more easily observed in terms of change in latency compared to asymptotic latency.

In summary, the results of the decision-making task provide two proxies for the rate of temporal discounting: trial-to-trial change in latency following a failed trial as expressed early in the experiment, and asymptotic latency as expressed late in the experiment.

Relationship between vigor and willingness to wait

To verify the validity of our vigor model, we first asked to what extent the vigor estimate for a subject in Experiment 1 was a predictor of their saccade peak velocities in Experiment 2. (The experiments were conducted on separate days.) We computed the mean saccade velocities in the first two blocks of Experiment 2 (i.e., baseline blocks) and found that vigor in Experiment 1 was strongly correlated with velocities recorded in Experiment 2 ($r = 0.89$, $F_{(1,21)} = 81.9$, $p < 10^{-7}$).

←

(Figure legend continued.) discounter has a discounted success function with a peak that is now in the past. This subject will stop waiting and respond to the first instruction. The shallow discounter, however, still has a discounted probability function with a peak that is in the future. For this subject, the optimum action is still in the future. This subject will continue to wait. **B**, Simulated decision making of the two subjects. Saccade latency (blue line) is the time that the simulated subject decided to move. On every trial with a second instruction, the delay Δ was adjusted (increased by 30 ms if the subject's decision resulted in success, decreased by 30 ms if the subject's decision resulted in failure). The steep discounter reached a maximum delay time that was much less than that of the shallow discounter. **C**, We simulated various discount rates and computed the final latency achieved for each discount function. Steeper discounters are expected to have smaller asymptotic latencies, i.e., wait shorter periods of time. **D**, Average trial-to-trial change in latency following a failed trial (a trial in which the second instruction arrived, but the simulated subject had chosen not to wait) for the first block (16 trials) for 100 simulated subjects. After a failed trial, shallow discounters increase their latency by a larger amount than steep discounters.

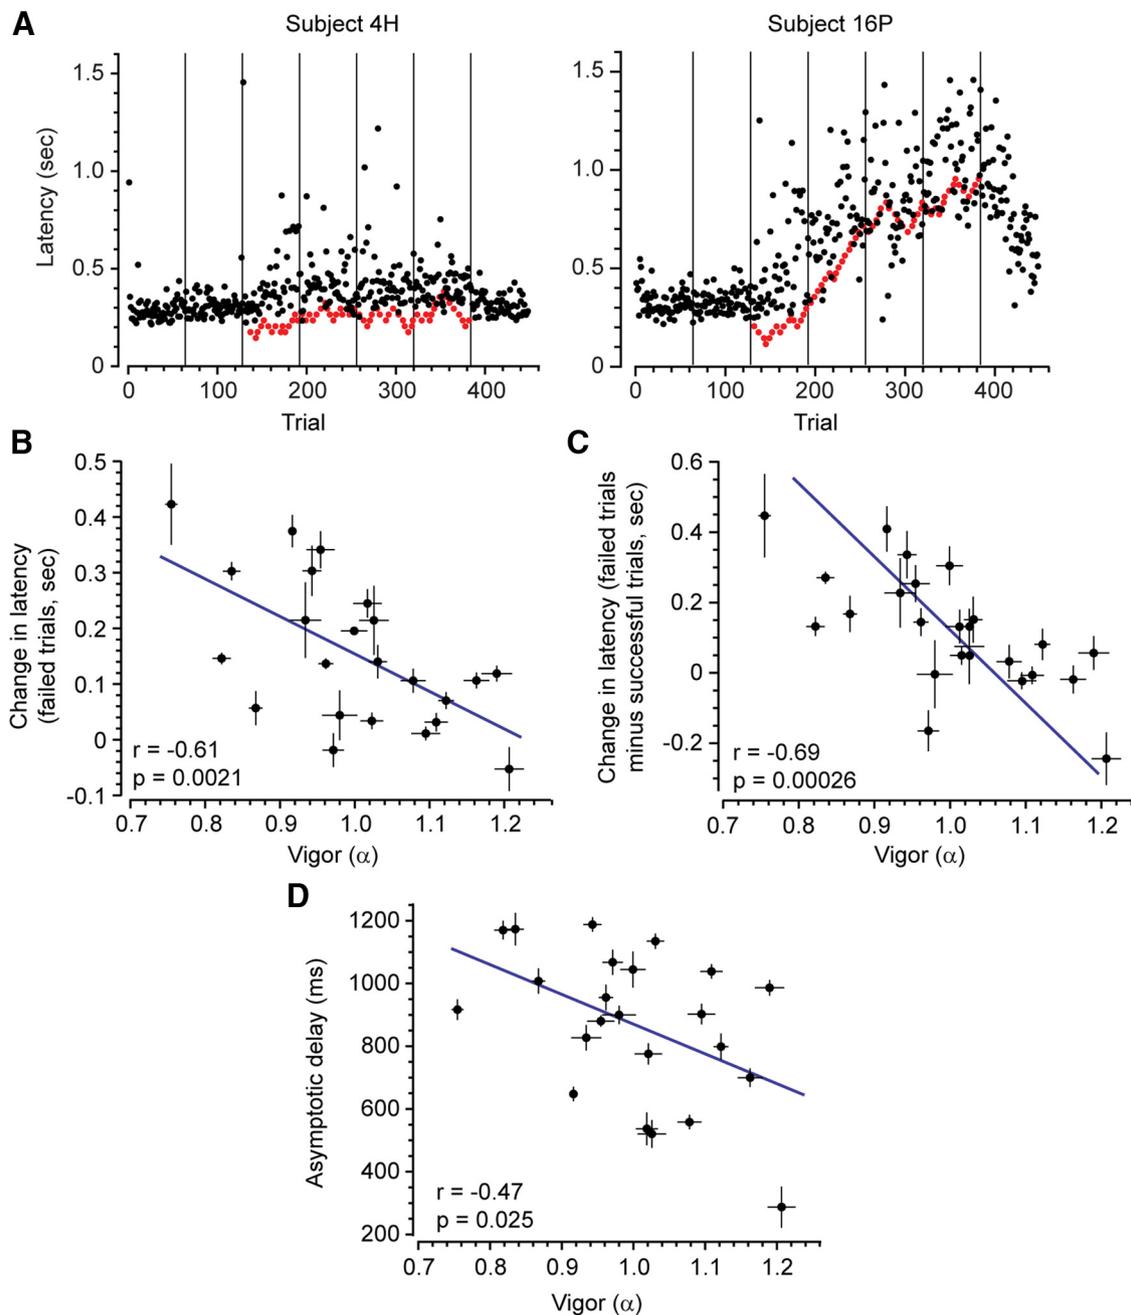


Figure 4. Relationship between willingness to wait and movement vigor. **A**, Latency of saccades and delay of the second instruction for two subjects. (Saccade velocities of these subjects, as measured in Part 1 of the experiment, are shown in Fig. 2A.) The vertical lines denote breaks between blocks. Saccade latencies are shown with black dots, and the latency of the instruction-change cue is shown with red dots. **B**, Trial-to-trial change in saccade latency in response to an error trial, i.e., a trial in which a second instruction occurred but the subject chose not to wait. The latencies were measured in the third block of the experiment (first block in which the second instruction occurred). Following an error trial, people who exhibited lower vigor tended to increase their wait time by a larger amount. **C**, Trial-to-trial change in saccade latency in response to an error trial, normalized with respect to change in latency in trials in which the second instruction occurred and the subjects had waited for it; that is, this measure quantifies behavioral change in response to an error trial, minus the change with respect to a successful trial. The latencies were measured in the third block of the experiment (first block in which the second instruction occurred). **D**, Relationship between asymptotic delay of the second instruction and saccade vigor. The delays were measured during the last half of the final block of trials in which there was a second instruction (Block 6). Data are mean \pm SEM for latency and delay and mean \pm SD for vigor. Blue lines show the results of linear regression.

Saccade latencies of two subjects in Experiment 2 are shown in Figure 4A. (These subjects are the same ones for which we displayed saccade velocities in Fig. 2A.) In the first two blocks, the probability of a second instruction was zero. In the subsequent four blocks, this probability increased to 0.25. At the start of the third block, Δ , representing the delay to the second instruction, was 200 ms. By the end of the sixth block, Δ had increased to

\sim 250 ms for Subject 4H, whereas it had increased to \sim 900 ms for Subject 16P.

Our model suggested two proxies for the rate of temporal discount function: change in latency early in the experiment following a failed trial and asymptotic latency late in the experiment. For change in latency, we focused on the first instruction-change block (third block overall) because latency changes became

smaller as the experiment proceeded (paired *t* test, latency change following a failed trial, first block with a second instruction vs last block with a second instruction, $t = 2.62$, $p = 0.016$) and because at this point in the experiment, subjects have similar instruction-change delays (vigor vs instruction-change cue delay, $r = -0.31$, $p = 0.15$). For every failed trial (in the first block in which the second instruction occurred), we computed the change in latency from the single-instruction trial before to the single-instruction trial after. We found that subjects who displayed more vigor in their saccades tended to have a small change in latency, as shown in Figure 4B ($r = -0.61$, $F_{(1,21)} = 12.3$, $p = 0.002$). This relationship was maintained when we normalized the change in latency with respect to successful trials (trials in which the second instruction occurred and the subject waited). We computed the change in latency following a failed trial and subtracted from it the change in latency following a successful trial. The subjects who displayed more vigor in their saccades had a strong tendency to exhibit a small change in this normalized measure of latency, as shown in Figure 4C ($r = -0.69$, $F_{(1,21)} = 19.2$, $p = 0.0003$); that is, people with high vigor were less willing to increase their latency following a failed trial, i.e., less willing to wait longer to improve their odds of success.

A second variable of interest was the asymptotic value of the instruction delay. We quantified this by taking the mean delay during the last half of the final block in which the second instruction occurred. Across our sample of volunteers, we found a negative correlation between the instruction delay and movement vigor: people who had higher saccade vigor tended to achieve a short instruction delay (Fig. 4D; $r = -0.47$, $F_{(1,21)} = 5.8$, $p = 0.025$). A similar result was found when we compared vigor with the maximum instruction delay achieved by each subject ($r = -0.49$, $F_{(1,21)} = 6.74$, $p = 0.017$).

Individual differences in valuations of immediate reward

In our model of temporal discounting, we assumed that the value of a delayed reward depended on its immediate value, multiplied by a function that discounted this value as a function of time (Equation 4). We made the assumption that two subjects differed in the rate of temporal discounting, but not the immediate value. In other words, we assumed that subjects did not differ in how they valued success in a given trial in which they did not have to wait (V in Equation 4). Is there a way to verify this assumption?

Generally, if one action is valued more than another, people (Milstein and Dorris, 2007) and animals (Tachibana and Hikosaka, 2012; Kim and Hikosaka, 2013) will react with a shorter latency or arrive at the target earlier in the more valuable scenario. Therefore, differences in the value of success may produce between-subject differences in RT or target acquisition time (RT plus movement duration) in the baseline period, i.e., in the block in which there was no instruction change. Although there were wide differences between people in the baseline blocks, such differences did not correlate with vigor (RT vs vigor, $r = -0.13$, $F_{(1,21)} = 0.36$, $p = 0.55$; target acquisition time vs vigor, $r = -0.23$, $p = 0.27$); that is, people with greater vigor were not faster in reacting to the first instruction.

There was no relationship between asymptotic delay and baseline target acquisition time ($r = 0.25$, $p = 0.25$), nor between the change in latency and baseline acquisition time (failed trials, $r = 0.06$, $p = 0.78$; failed minus successful trials, $r = 0.07$, $p = 0.75$). There was also no correlation between baseline reaction time and our measures of temporal discounting (asymptotic delay vs RT, $r = 0.23$, $p = 0.29$; change in latency, failed trials vs RT, $r = 0.08$, $p = 0.83$; failed minus successful trials vs RT, $r = 0.05$, $p = 0.83$).

People who had larger asymptotic delays or larger changes in trial-to-trial latency were not faster in reacting to the first instruction.

We demonstrated previously that differences in the implicit value of a stimulus may be reflected in the peak velocity of saccades to that stimulus (Xu-Wilson et al., 2009). However, in that work, the change in the peak velocity for stimuli of differing value was $\sim 5^\circ/\text{s}$ for a 15° saccade, more than an order of magnitude smaller than the differences in peak velocity between subjects in this work. Accordingly, the differences in vigor between subjects are unlikely to be driven by differences in stimulus value.

A difference in the value of the stimuli could also be reflected in the rate in which subjects followed the directional cue given by the first instruction. After the first block, subjects made saccades in the wrong direction only $1.0 \pm 0.2\%$ (mean \pm SEM across subjects) of the time. The accuracy of subjects was not correlated with the vigor of their movements ($r = -0.01$, $p = 0.98$). Overall, these analyses suggest that there was no systematic difference in the way subjects valued the stimuli.

Reaction to the instruction-change cue

Our task for measuring temporal discounting was similar to the stop signal reaction time (SSRT) task, a task in which subjects are provided with a “go” cue, which is occasionally followed with a “stop” cue. The objective of the SSRT task is to measure how long it takes after the occurrence of the stop signal for the subject to abort their planned movement (this latency is called the SSRT). An important difference in the SSRT task versus our task is that in the SSRT task, subjects are instructed to respond to the go cue as quickly as possible and not delay their response to await the stop cue. In our task, the subjects were told that occasionally there would be a second instruction. They were allowed to wait as long as they wished to respond to the first instruction. Despite this difference, we thought it worthwhile to analyze our data to quantify how behavior was affected in trials in which instruction-change occurred.

We began by asking whether saccade kinematics were different in the instruction-change trials (i.e., the failed trials). Our thought was that (partial) inhibition of a planned movement may produce a reduction in its amplitude. Indeed, subjects made significantly smaller saccades in instruction-change trials ($19.8 \pm 0.3^\circ$, mean \pm SEM) compared to no-change trials ($20.6 \pm 0.1^\circ$; paired *t* test, $p < 0.003$). We considered the amplitude of saccades made in those trials as a function of when the saccade was made relative to the second cue. We found that if a saccade occurred 40 ms or later after the instruction-change cue, its amplitude was reduced compared to no-change trials ($18.9 \pm 0.5^\circ$, mean \pm SEM across subjects; paired *t* test, $p < 0.0005$). However, saccades made before 40 ms after the instruction-change cue showed no amplitude differences ($20.6 \pm 0.2^\circ$, paired *t* test, $p = 0.62$). The saccades were, on average, 67 ± 1.3 ms in duration. Therefore, it took a minimum of ~ 110 ms after the instruction-change cue for the brain to alter the ongoing motor commands. This value provides an objective estimate of the lower bound on the SSRT in our task.

To estimate the SSRT for each subject, we used the approach suggested by Eagle et al. (2008) for experiments in which the timing of the second instruction is adjusted via an adaptive “staircase” procedure: we subtracted the median of the latency for the second instruction from the median of the reaction time in the trials without a second instruction. We found that, on average, the SSRT was 120.3 ± 11.3 ms (population mean \pm SEM), which agrees well with our independent measure using saccade kine-

matics (lower bound of 110 ms). A within-subject comparison of SSRT and vigor did not result in a significant correlation ($r = 0.23$, $F_{(1,21)} = 1.18$, $p = 0.29$). People who require a long time to inhibit a planned action (manifested in long SSRT) are thought to be more impulsive. Therefore, the positive value of the correlation, though not significant, is in line with our general framework. Our task, however, was not designed to measure SSRT.

Psychological profile

A commonly used method to assess decision-making characteristics of individuals is via questionnaires that measure impulsiveness. These questionnaires estimate personality traits by determining the response to queries such as “do you often buy things on impulse?” “do you mostly speak before thinking things out?” etc. Our subjects filled out two commonly used questionnaires, termed BIS and I7. Higher scores in these questionnaires suggest a psychological profile for impulsiveness. In our subjects, the score in one questionnaire was strongly correlated with the score in the other ($r = 0.65$, $F_{(1,21)} = 15.5$, $p < 0.001$). However, impulsivity as measured by these questionnaires was never a good predictor of movement vigor (I7 impulsivity subscore vs vigor, $r = 0.20$, $p = 0.35$; BIS vs vigor, $r = 0.17$, $p = 0.45$), nor of the asymptotic instruction delay (I7 impulsivity subscore vs delay, $r = -0.10$, $p = 0.66$; BIS vs delay, $r = -0.01$, $p = 0.95$). The positive correlation values for vigor and the negative correlation values for delay indicated that, in general, people who scored as slightly more impulsive on the questionnaires tended to have higher vigor and slightly shorter delays, though this tendency was not significant.

Discussion

We found consistent differences among healthy people in the speed with which they moved their eyes during a saccade. We quantified this via a measure of vigor that summarized the relationship between saccade amplitude and peak velocity. Vigor differed by as much as 50% between subjects, but was highly consistent within subjects. We hypothesized that differences in vigor may be partly due to differences in how the brain discounts reward as a function of time.

To measure temporal discounting, we considered a task in which subjects received instructions to perform an action, but improved their odds of success if they waited for a second instruction. We found that people with high vigor were less willing to increase their latency following a trial in which they failed, suggesting a higher temporal discount rate. This measure of temporal discounting in the decision-making task accounted for 48% of the between-subject variance in vigor.

To what extent can differences in vigor be explained with differences in biomechanics? In a previous study, the eyes, orbit, and extraocular muscles of healthy volunteers were imaged using MRI (Peng et al., 2012). That study concluded that the measured parameters (including muscle volume and cross-sectional area) could not account for the between-subject differences in saccade velocity. Biomechanics of the eyes are critical in describing consequences of motor commands, affecting potential costs of movements in terms of effort and variability. In our population, end point variability of saccades was unrelated to vigor, i.e., people with high vigor were not more variable, as might be expected from a signal-dependent noise perspective.

To measure temporal discounting, we designed a task in which choices had consequences (success or failure) that acted as operant reinforcers before the next choice was made. We relied on the fact that the reinforcer caused a change in behavior from

one trial to the next, and the magnitude of this change was a signature of the temporal discount function. In contrast, most experiments that measure temporal discounting in humans rely on nonoperant reinforcers in which people make choices between dollar amounts, and the consequences are either hypothetical or realized only after the experiment is over (because the delays are in days or weeks). Although both types of experiments produce measures of temporal discounting, they produce inconsistent results in the same person (Hyten et al., 1994) and produce discount rates that differ by many orders of magnitude (Navarick, 2004). The operant approach is the principal method of measuring discounting in nonhuman primates, which guided our design here.

Our task is similar to the SSRT task. In the SSRT task, subjects are provided with a stimulus that instructs a movement, but they are told to not delay their response to this instruction. In the SSRT task, the objective is to measure how quickly subjects can stop their planned movement in the case that a “stop” instruction appears. In our task, the subjects were allowed to wait as long as they wished. When the instruction changed, we estimated 110 ms as the lower bound for the time it took the brain to process the new instruction and alter the saccade. People who need a longer time to inhibit their movements exhibit impulsivity (Verbruggen and Logan, 2008), and in our sample such people tended to exhibit greater vigor, though the effect was not statistically significant.

The key variables in our task were the change in latency after an unsuccessful trial and the asymptotic latency, both of which we found to negatively correlate with vigor. A limitation of our model, however, is that the change in latency that it predicted for a given discount function was a scaled version of the values actually observed. It is unclear to us where this limitation arises from. It may indicate an asymmetry in valuation of success versus failure.

Neural basis of vigor and the link to encoding of reward

The vigor with which a saccade is performed is associated with the activity of “buildup” cells in the intermediate layers of the superior colliculus (SC; Ikeda and Hikosaka, 2007). When a saccade is planned toward a location that falls within the receptive field of an SC cell, the upcoming saccade displays greater vigor if that cell fires more strongly during the period before the saccade. This buildup activity is partly under the control of cells in an output nucleus of basal ganglia, the substantia nigra pars reticulata (SNr). SNr cells constantly inhibit the SC, but generally pause before a movement (Hikosaka and Wurtz, 1985; Handel and Glimcher, 1999). More vigorous saccades are associated with a deeper pause in the firing rates of SNr cells (Sato and Hikosaka, 2002). Indeed, saccadic vigor is increased by blocking the SNr SC inhibition (Hikosaka and Wurtz, 1985). Therefore, control of vigor is partly a function of the basal ganglia.

Within the basal ganglia, the nucleus critical for control of vigor is the external segment of globus pallidus (GPe). GPe cells inhibit the SNr and fire more strongly preceding a more vigorous saccade, and bilateral lesion of this region eliminates the ability of the animal to modulate saccade vigor in response to changes in reward (Tachibana and Hikosaka, 2012). GPe cells are inhibited by a subset of cells in the caudate (caudate cells that are part of the “indirect pathway”). Caudate cells generally fire more before a vigorous saccade (Kawagoe et al., 1998) and receive dopamine projections. Onset of a stimulus that promises reward results in a burst of dopamine (Matsumoto and Hikosaka, 2007), which is followed by a more vigorous saccade (Tachibana and Hikosaka,

2012). Indeed, chronic reduction in the concentration of dopamine in the caudate reduces saccade vigor by ~30% (Kori et al., 1995). Therefore, control of vigor is partly associated with the amount of dopamine in the basal ganglia, particularly in the indirect pathway (caudate–GPe–SNr).

Temporal discounting is also associated with release of dopamine. Consider a task in which an animal makes a decision between two stimuli, one that predicts a small reward soon and another that predicts a large reward later. Dopamine cells fire in response to each stimulus by an amount that correlates with the temporally discounted value of that stimulus (Kobayashi and Schultz, 2008). In the small number of animals for which data are available, between-subject differences in the rate of discounting as measured via dopamine discharge is a predictor of between-subject differences in decision making (Kobayashi and Schultz, 2008). Together, it appears that some of the neural circuits that are critical for control of vigor are also influenced by a neurotransmitter that has been linked to temporal discounting.

The mathematical framework of optimal control predicts a link between vigor and temporal discounting by suggesting that before a movement can be generated, there needs to be an evaluation of the reward that is expected at the end of the movement, discounted by the time it takes to complete that movement (Shadmehr et al., 2010). Of course, it is possible that there are two separate temporal discounting systems for control of movements and decision making, as the two have vastly different time scales. However, if the basis of both forms of temporal discounting is to maximize discounted rate of reward (Haith et al., 2012), then from a theoretical perspective, there is justification for the idea that there is a single temporal discounting system that affects control of movements as well as decision making. From an evolutionary perspective, control of movements may have required temporal discounting, which in turn was generalized to control of decisions.

Reaction time, impulsivity, and vigor

It is possible that during the RT period, the brain is solving the problem of “what is the best action that I can perform?” whereas during the movement, the brain is solving the problem “how do I perform this action?” Indeed, when there are two possible actions, during the RT period there is competition between the two actions: the brain accumulates evidence for each action, and the action that reaches a threshold first is selected (Gold and Shadlen, 2002). A person that has a high cost of time should, in principle, have a lower threshold, selecting actions earlier and with less evidence. If there are differences in the cost of time among people, and if these costs generalize between action selection and action execution, then there should be a negative correlation between vigor and RT.

As people wait for an expected reward, activity (as measured by fMRI) in the ventral striatum and ventromedial PFC rises, and this rise has a steeper slope for people who have a steeper temporal discount function (Jimura et al., 2013). Impulsivity is a psychological trait that is often measured via questionnaires. Impulsive people show diminished midbrain D₂/D₃ autoreceptor availability, which results in increased dopamine release in the striatum (Buckholz et al., 2010). In our sample of subjects, there was a positive, but not significant, correlation between survey-based measures of impulsivity and movement vigor. We suspect that the reason for this is that impulsivity is a complex trait that involves interactions between the basal ganglia and the frontal lobe. For example, in humans, the temporally discounted value of reward is correlated with activation in the medial prefrontal cor-

tex (Jimura et al., 2013), in addition to the dorsal and ventral striatum (Kable and Glimcher, 2007; Pine et al., 2009). The cost of time as reflected in saccade vigor may be due to the control that the basal ganglia imposes on the SC, which in turn is affected by dopamine, whereas the cost of time as reflected in decision making is a more complex process that involves interactions between the basal ganglia and the cerebral cortex.

References

- Bornstein MH, Bornstein HG (1976) The pace of life. *Nature* 259:557–559. [CrossRef](#)
- Buckholz JW, Treadway MT, Cowan RL, Woodward ND, Li R, Ansari MS, Baldwin RM, Schwartzman AN, Shelby ES, Smith CE, Kessler RM, Zald DH (2010) Dopaminergic network differences in human impulsivity. *Science* 329:532. [CrossRef](#) [Medline](#)
- Chen-Harris H, Joiner WM, Ethier V, Zee DS, Shadmehr R (2008) Adaptive control of saccades via internal feedback. *J Neurosci* 28:2804–2813. [CrossRef](#) [Medline](#)
- Eagle DM, Bari A, Robbins TW (2008) The neuropsychopharmacology of action inhibition: cross-species translation of the stop-signal and go/no-go tasks. *Psychopharmacology (Berl)* 199:439–456. [CrossRef](#)
- Eysenck SBG, Pearson PR, Easting G, Allsopp JF (1985) Age norms for impulsiveness, venturesomeness, and empathy in adults. *Personality Individual Differences* 6:613–619. [CrossRef](#)
- 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. [Medline](#)
- Gold JJ, Shadlen MN (2002) Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36:299–308. [CrossRef](#) [Medline](#)
- Haith AM, Reppert TR, Shadmehr R (2012) Evidence for hyperbolic temporal discounting of reward in control of movements. *J Neurosci* 32:11727–11736. [CrossRef](#) [Medline](#)
- Handel A, Glimcher PW (1999) Quantitative analysis of substantia nigra pars reticulata activity during a visually guided saccade task. *J Neurophysiol* 82:3458–3475. [Medline](#)
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394:780–784. [CrossRef](#) [Medline](#)
- Hikosaka O, Wurtz RH (1985) Modification of saccadic eye movements by GABA-related substances. II. Effects of muscimol in monkey substantia nigra pars reticulata. *J Neurophysiol* 53:292–308. [Medline](#)
- Hoyt DF, Taylor CR (1981) Gait and the energetics of locomotion in horses. *Nature* 292:239–240. [CrossRef](#)
- Hwang J, Kim S, Lee D (2009) Temporal discounting and inter-temporal choice in rhesus monkeys. *Front Behav Neurosci* 3:9. [Medline](#)
- Hyten C, Madden GJ, Field DP (1994) Exchange delays and impulsive choice in adult humans. *J Exp Anal Behav* 62:225–233. [CrossRef](#) [Medline](#)
- Ikeda T, Hikosaka O (2007) Positive and negative modulation of motor response in primate superior colliculus by reward expectation. *J Neurophysiol* 98:3163–3170. [CrossRef](#) [Medline](#)
- Ivry RB, Hazeltine RE (1995) Perception and production of temporal intervals across a range of durations: evidence for a common timing mechanism. *J Exp Psychol Hum Percept Perform* 21:3–18. [CrossRef](#) [Medline](#)
- 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. [CrossRef](#) [Medline](#)
- Jimura K, Chushak MS, Braver TS (2013) Impulsivity and self-control during intertemporal decision making linked to the neural dynamics of reward value representation. *J Neurosci* 33:344–357. [CrossRef](#) [Medline](#)
- Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. *Nat Neurosci* 10:1625–1633. [CrossRef](#) [Medline](#)
- Kawagoe R, Takikawa Y, Hikosaka O (1998) Expectation of reward modulates cognitive signals in the basal ganglia. *Nat Neurosci* 1:411–416. [CrossRef](#) [Medline](#)
- Kim HF, Hikosaka O (2013) Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. *Neuron* 79:1001–1010. [CrossRef](#) [Medline](#)
- Kobayashi S, Schultz W (2008) Influence of reward delays on responses of dopamine neurons. *J Neurosci* 28:7837–7846. [CrossRef](#) [Medline](#)
- Kori A, Miyashita N, Kato M, Hikosaka O, Usui S, Matsumura M (1995) Eye

- movements in monkeys with local dopamine depletion in the caudate nucleus. II. Deficits in voluntary saccades. *J Neurosci* 15:928–941. [Medline](#)
- Matsumoto M, Hikosaka O (2007) Lateral habenula as a source of negative reward signals in dopamine neurons. *Nature* 447:1111–1115. [CrossRef](#) [Medline](#)
- Millar A, Navarick DJ (1984) Self-control and choice in humans: Effects of video game playing as a positive reinforcer. *Learn Motivation* 15:203–218. [CrossRef](#)
- Milstein DM, Dorris MC (2007) The influence of expected value on saccadic preparation. *J Neurosci* 27:4810–4818. [CrossRef](#) [Medline](#)
- Myerson J, Green L (1995) Discounting of delayed rewards: Models of individual choice. *J Exp Anal Behav* 64:263–276. [CrossRef](#) [Medline](#)
- Navarick DJ (2004) Discounting of delayed reinforcers: measurement by questionnaires versus operant choice procedures. *Psychol Record* 54:85–94.
- Niv Y, Daw ND, Joel D, Dayan P (2007) Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology (Berl)* 191:507–520. [CrossRef](#)
- Opris I, Lebedev M, Nelson RJ (2011) Motor planning under unpredictable reward: modulations of movement vigor and primate striatum activity. *Front Neurosci* 5:1–12. [CrossRef](#)
- Patton JH, Stanford MS, Barratt ES (1995) Factor structure of the Barratt impulsiveness scale. *J Clin Psychol* 51:768–774. [CrossRef](#) [Medline](#)
- Peng MY, Ying S, Zee DS, Barker PB, Ying H (2012) Three-dimensional magnetic resonance analysis of extraocular muscles demonstrates gaze-specific contractile shape changes. *Assoc Vis Ophthalmol Meeting Abstr* 53:4854.
- Pine A, Seymour B, Roiser JP, Bossaerts P, Friston KJ, Curran HV, Dolan RJ (2009) Encoding of marginal utility across time in the human brain. *J Neurosci* 29:9575–9581. [CrossRef](#) [Medline](#)
- Sato M, Hikosaka O (2002) Role of primate substantia nigra pars reticulata in reward-oriented saccadic eye movement. *J Neurosci* 22:2363–2373. [Medline](#)
- Shadmehr R, Mussa-Ivaldi S (2012) *Biological learning and control: how the brain builds representations, predicts events, and makes decisions*. Cambridge, MA: MIT.
- Shadmehr R, Orban de Xivry JJ, Xu-Wilson M, Shih TY (2010) Temporal discounting of reward and the cost of time in motor control. *J Neurosci* 30:10507–10516. [CrossRef](#) [Medline](#)
- Straube A, Fuchs AF, Usher S, Robinson FR (1997) Characteristics of saccadic gain adaptation in rhesus macaques. *J Neurophysiol* 77:874–895. [Medline](#)
- Tachibana Y, Hikosaka O (2012) The primate ventral pallidum encodes expected reward value and regulates motor action. *Neuron* 76:826–837. [CrossRef](#) [Medline](#)
- 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. [CrossRef](#) [Medline](#)
- Verbruggen F, Logan GD (2008) Response inhibition in the stop-signal paradigm. *Trends Cogn Sci* 12:418–424. [CrossRef](#) [Medline](#)
- Westheimer G (1999) Discrimination of short time intervals by the human observer. *Exp Brain Res* 129:121–126. [CrossRef](#) [Medline](#)
- Willis WT, Ganley KJ, Herman RM (2005) Fuel oxidation during human walking. *Metabolism Clinical and Experimental* 54:793–799. [CrossRef](#) [Medline](#)
- Xu-Wilson M, Zee DS, Shadmehr R (2009) The intrinsic value of visual information affects saccade velocities. *Exp Brain Res* 196:475–481. [CrossRef](#) [Medline](#)