

Modulation of Saccade Vigor during Value-Based Decision Making

Thomas R. Reppert,^{1*} Karolina M. Lempert,^{2*} Paul W. Glimcher,² and Reza Shadmehr¹

¹Department of Biomedical Engineering, Johns Hopkins University School of Medicine, Baltimore, Maryland 21205, and ²Center for Neural Science, New York University, New York, New York 10003

During value-based decision-making, individuals consider the various options and select the one that provides the maximum subjective value. Although the brain integrates abstract information to compute and compare these values, the only behavioral outcome is often the decision itself. However, if the options are visual stimuli, during deliberation the brain moves the eyes from one stimulus to the other. Previous work suggests that saccade vigor, i.e., peak velocity as a function of amplitude, is greater if reward is associated with the visual stimulus. This raises the possibility that vigor during the free viewing of options may be influenced by the valuation of each option. Here, humans chose between a small, immediate monetary reward and a larger but delayed reward. As the deliberation began, vigor was similar for the saccades made to the two options but diverged 0.5 s before decision time, becoming greater for the preferred option. This difference in vigor increased as a function of the difference in the subjective values that the participant assigned to the delayed and immediate options. After the decision was made, participants continued to gaze at the options, but with reduced vigor, making it possible to infer timing of the decision from the sudden drop in vigor. Therefore, the subjective value that the brain assigned to a stimulus during decision-making affected the motor system via the vigor with which the eyes moved toward that stimulus.

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

Significance Statement

We find that, as individuals deliberate between two rewarding options and arrive at a decision, the vigor with which they make saccades to each option reflects a real-time evaluation of that option. With deliberation, saccade vigor diverges between the two options, becoming greater for the option that the individual will eventually choose. The results suggest a shared element between the network that assigns value to a stimulus during the process of decision-making and the network that controls vigor of movements toward that stimulus.

Introduction

There is some evidence that the vigor with which a movement is performed (i.e., its peak speed as a function of amplitude) is affected by the subjective value that the brain assigns to the goal of the movement. For example, Sackaloo et al. (2015) asked participants to rank order in terms of preference a number of different kinds of candy bars. When asked to reach for a

single candy bar, participants reached faster and with a shorter duration for the more preferred bar. Similarly, monkeys reached with a greater speed toward stimuli that promised higher probability of reward (Opris et al., 2011). These observations raise the possibility that movement vigor may be modulated by the subjective value that the brain assigns to the goal of the movement. Humans and other primates use saccadic eye movements to examine their available options. During deliberation, as one makes saccades to accumulate information about the available options, does saccadic vigor reflect the subjective value that the brain currently assigns to each option?

Previous work has shown that stimulus value can grossly affect the peak velocity of saccades. Monkeys exhibit a greater saccade peak velocity when the visual target is paired with food reward (Takikawa et al., 2002). Humans perform their saccades with greater peak velocity if the target is a valued stimulus, such as a face (Xu-Wilson et al., 2009). Here, we

Received July 9, 2015; revised Sept. 2, 2015; accepted Oct. 14, 2015.

Author contributions: K.M.L. designed research; K.M.L. performed research; T.R.R. analyzed data; T.R.R., K.M.L., P.W.G., and R.S. wrote the paper.

This work was supported by National Institutes of Health Grant NS37422 and the Human Frontiers Science Program.

The authors declare no competing financial interests.

*T.R.R. and K.M.L. contributed equally to this work.

Correspondence should be addressed to Thomas Reppert, 416 Traylor Building, Johns Hopkins School of Medicine, 720 Rutland Avenue, Baltimore, MD 21205. E-mail: treppert@jhu.edu.

DOI:10.1523/JNEUROSCI.2621-15.2015

Copyright © 2015 the authors 0270-6474/15/3515369-10\$15.00/0

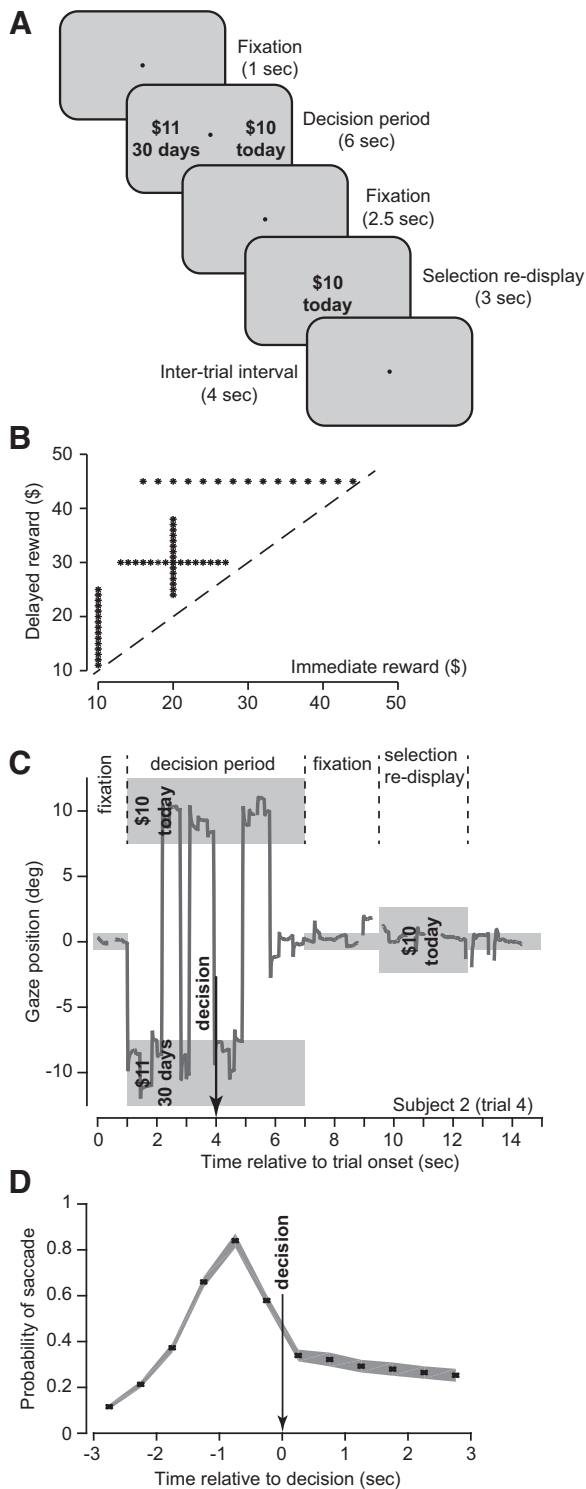


Figure 1. Experimental design. Participants completed a decision-making task in which they made a choice between a small immediate monetary reward and a larger delayed monetary reward to be attained at 30 d. **A**, Each trial began with a fixation dot displayed for 1 s, followed by display of the two reward options at either side of the fixation dot. Subjects were instructed to press a key designating their choice within 6 s. Regardless of when the participants indicated their choice, the options remained on the screen for the full 6 s. **B**, The reward pairings. Each dot represents one of the 60 reward pairings that was presented to each participant. The delayed reward was always greater than the immediate reward. Each reward pairing was presented twice. The order of presentation was random for each participant, as well as across participants. The position of presentation of each option was randomly selected to be centered at 10° to the left or right side of fixation. **C**, Exemplar trace of gaze position during a single trial. During the decision period, the arrow indicates when the participant pressed the key

designating his/her decision. Note that the participant continued to make saccades after decision time, while both options remained on the screen. In this particular trial, the participant chose the immediate reward (\$10) over the delayed option (\$11), and thus the \$10 option was redisplayed after the fixation period. **D**, Probability of saccade with respect to decision time, computed for bins of 0.5 s in duration. Shaded region is SEM.

considered a decision-making task in which participants were offered monetary rewards. We asked whether the vigor with which a saccade was performed was affected by the subjective value that the brain assigned to the potential reward. A critical component of our design was that the movement that we considered (saccade) had no bearing on the reward itself: that is, people were not rewarded for making saccades. Rather, the saccades were a mechanism with which participants acquired information for the purpose of decision-making. Our participants completed a temporal discounting task in which they chose between a small, immediate reward (in dollars) and a large reward delayed reward (to be received in 30 d). People prefer rewards sooner rather than later but vary widely in how much they are willing to wait for the larger delayed reward. We measured participants' eye movements as they considered their two choices, tracking the real-time velocity and amplitude of each saccade as they directed their gaze at each option. When the deliberation process started, the eyes moved with the same vigor toward the two options, but as the deliberation process continued, vigor became greater for the option that the subject would eventually choose. Immediately after the subject indicated a choice, vigor dropped for both options. These observations suggest that, during decision-making, the vigor with which the brain moves the eyes toward a stimulus may be a reflection of the current value that it assigns to that stimulus.

Materials and Methods

Participants. We recruited $n = 60$ healthy participants from the New York University community with no known neurologic deficits (aged 21.75 ± 3.01 years, mean \pm SD; 35 females). All were naive to the paradigm and purpose of the experiment. Each participant signed a written consent form approved by the New York University Committee on Activities Involving Human Subjects. Each was paid \$10/h in cash for participating in the study, as well as additional compensation based on their decisions in the task, as described below.

Behavioral task. Subjects sat in a darkened room in front of a CRT monitor (36.5×27.5 cm, 1024×768 pixels, light gray background, frame rate of 120 Hz), head stabilized with the use of a chinrest. The screen was placed at a distance of 55 cm from the subject's eyes. An EyeLink 1000 (SR Research) infrared camera recording system recorded movements and pupil diameter of the right eye. Gaze position and pupil diameter were recorded at 250 Hz for all subjects, with the exception of two subjects recorded at 1000 Hz and one recorded at 500 Hz. A superset of the data from this study was also examined with regard to changes in pupil diameter. A report of those findings appeared previously (Lempert et al., 2015).

We measured eye movements during a temporal discounting task. The time course of a typical trial is displayed in Figure 1A. The trial began with a 1 s fixation period (dot displayed at center of screen). Right after the fixation period, written description of the two possible rewards appeared simultaneously on the screen: (1) a text that described a small immediate monetary reward (for example, "\$10 today"); and (2) a text that described a larger delayed monetary reward (for example, "\$11 30 d"). Each text was centered at 10° to the left or right of the center fixation dot and was $4.7\text{--}7.9^\circ$ wide and 5.4° tall (as shown in Fig. 1C). The placement of the text on the left or right was chosen at random. One option was always for an immediate reward,

←

designating his/her decision. Note that the participant continued to make saccades after decision time, while both options remained on the screen. In this particular trial, the participant chose the immediate reward (\$10) over the delayed option (\$11), and thus the \$10 option was redisplayed after the fixation period. **D**, Probability of saccade with respect to decision time, computed for bins of 0.5 s in duration. Shaded region is SEM.

whereas the other option was always for a reward to be attained in 30 d.

During this decision period, participants made saccades to the stimuli. The stimuli remained on the screen for exactly 6 s, during which time the subjects indicated their decision by pressing a key (typically at ~2–3 s into the trial). Subjects were instructed to place the left hand over the 1 key and right hand over the 0 key, to select leftward and rightward rewards, respectively. Regardless of when the subjects pressed the key to designate their decision, the two options remained on the screen for the full 6 s. This was critical because this allowed the subjects to make saccades to the stimuli both before and after they made their decision.

After the completion of the 6 s decision period, the fixation dot remained on the screen for another 2.5 s. Finally, the fixation dot was removed, and the participants were presented with the option that they had chosen for 3 s. A new trial commenced after an intertrial interval of 4 s. There were 120 trials in the experiment. Our analysis focused solely on the saccades made during the 6 s decision period.

The participants were presented with 60 distinct monetary reward pairings, as shown in Figure 1B, with each pairing presented twice. The reward pairings were selected in random order such that no two subjects saw the same ordering of stimuli. On every trial, the delayed monetary reward was of greater magnitude than the immediate reward.

To increase the relevance of their choices, participants were instructed that one trial would be selected at random and they would receive the amount that they chose on that trial. That is, if they chose the immediate reward on the randomly selected trial, they would receive the money in cash after completion of the session. If they chose the larger, delayed reward, they would receive a debit card that would be activated after the delay (30 d) had elapsed.

One participant did not complete all 120 trials and was excluded from analysis. In addition, we were unable to achieve good eye calibration in six participants, which prevented measurement of saccades for those subjects. As a result, we analyzed the data from a total of $n = 53$ participants.

During each trial, we continuously recorded gaze position. Raw gaze position signals were smoothed and differentiated with the use of a Savitzky–Golay filter (second-order). The filter width was chosen as a function of the sampling rate such that each filter window encompassed 20 ms of data. We used the gaze velocity trace to determine onset and offset of saccades, with a 30°/s threshold. We used the following five criteria to identify task-relevant saccades: (1) horizontal amplitude $>2^\circ$ and $<25^\circ$; (2) vertical amplitude $<6^\circ$, with the ratio of vertical amplitude to horizontal amplitude <0.7 ; (3) peak horizontal acceleration $<35,000^\circ/s^2$; (4) skew (defined as the ratio of time from saccade start to peak velocity to saccade duration) <0.7 ; and (5) duration >20 and <120 ms. The first criterion removed $45 \pm 10\%$ (mean \pm SD) of saccades (many of the saccades were associated with the act of reading the text on the screen, a series of microsaccades). The remaining criteria together excluded $29 \pm 9\%$ of the remaining saccades. To identify an outlier saccade, we used the median absolute deviation technique (on the parameter saccade vigor) that excluded $2.7 \pm 1.5\%$ of the remaining saccades (Rousseeuw and Croux, 1993).

Data analysis: saccade vigor. During the decision period, subjects made saccades that terminated at either one of the stimuli or at the center fixation point (as illustrated in Fig. 1C). These saccades had a participant-specific velocity–amplitude relationship: some participants exhibited fast saccades, whereas others exhibited slow saccades (Choi et al., 2014). Our hypothesis was that, in a given individual, for a given saccade amplitude, the brain modulated saccade velocity as a function of reward or context (Xu–Wilson et al., 2009). To dissociate amplitude-dependent changes in velocity from reward-dependent changes in each individual, we first modeled the amplitude-dependent effects of saccade velocity for that individual and then compared changes in velocity that were present when amplitude was kept constant but reward or context changed. The result was a within-subject measure of saccade vigor, as described below.

For each participant n , we measured the amplitude of the saccade (represented by x) and its peak velocity (represented by v) in all trials.

Previous work had shown that a hyperbolic function is generally a good fit to human saccade data (Choi et al., 2014). Therefore, we fitted the data to the following function:

$$v_n = \alpha_n \left(1 - \frac{1}{1 + \beta_n x} \right). \quad (1)$$

We quantified the goodness of fit of the function for each participant using correlation coefficients. This fit produced parameter values $\hat{\alpha}_n$ and $\hat{\beta}_n$.

Given saccade amplitude x , the expected saccade velocity in subject n was represented by $\hat{v}_n(x)$. For each saccade, we computed the ratio between the measured velocity and the expected velocity: v_n/\hat{v}_n . This ratio defined a within-subject measure of saccade vigor. When this ratio was >1 , the saccade had a velocity that was larger than expected, reflecting a greater than average vigor for that subject. We used this within-subject measure of vigor to quantify changes in saccade peak velocity as a function of time during the decision-making period and as a function of the preference that the subjects exhibited toward the available options in each trial.

Data analysis: decision-making. We analyzed the decisions that each participant made by finding the value of the delayed reward that made that option equivalent to the immediate reward. For each participant, we represented the probability of choosing the delayed reward r_d as a function of the difference in the value of the delayed and immediate rewards $r_d - r_i$:

$$\Pr(\text{choice} = r_d) = \frac{1}{1 + \exp(-b(r_d - r_i - \alpha))}. \quad (2)$$

In the above expression, α represents the point of subjective equivalence between the delayed and immediate options. We fitted the above equation to the choices that the participant had made across all trials. To do so, we analyzed the trials based on the difference between the delayed and immediate rewards and then measured the probability of choosing the delayed reward in each trial. Therefore, in a trial in which $r_d - r_i = \hat{\alpha}$, the participant was equally likely to pick the delayed or the immediate option. Participants who preferred the immediate reward more often, and thus were more impulsive in their decision-making, had larger values of $\hat{\alpha}$.

To estimate the subjective value of an option for participant n , we considered a hyperbolic model of temporal discounting (Mazur, 1987; Green and Myerson, 2004; Kable and Glimcher, 2007). In this model, one assumes that people evaluate a future reward (promised to arrive after time delay t) by discounting it hyperbolically to produce a subjective value at present:

$$r_i = \frac{r_d}{1 + k_n t}. \quad (3)$$

In our experiment, the time delay t was always 1 month. For each participant n , we estimated discount factor k_n as a function of the mean ratio r_d/r_i for all trials in which the absolute difference $|r_d - r_i|$ was within \$5 of equivalence point $\hat{\alpha}$. To confirm this estimate, we also divided up the trials into four subsets (Fig. 1B, each vertical and horizontal line) and then re-estimated k_n independently for each subset of trials in each participant. This way of estimating k_n kept either the immediate or the delayed reward constant for each subset of trials. We compared the two methods and found that the two estimates correlated very well ($r^2 = 0.96$, slope of 1.174, bias of -0.14). In our results, we report the estimate arrived at using the entire dataset.

All statistical analyses were performed using SPSS (version 22; IBM) or MATLAB R2014b (MathWorks). All t tests presented are two tailed, unless specified otherwise.

Results

Saccade vigor was higher during the deliberation period

On each trial, the participants were presented with two options: (1) a monetary reward to be acquired on the day of the study; and (2) a larger reward to be acquired in 30 d. As the participants evaluated the two options and made their decision, they made saccades from one stimulus to another. On average, participants

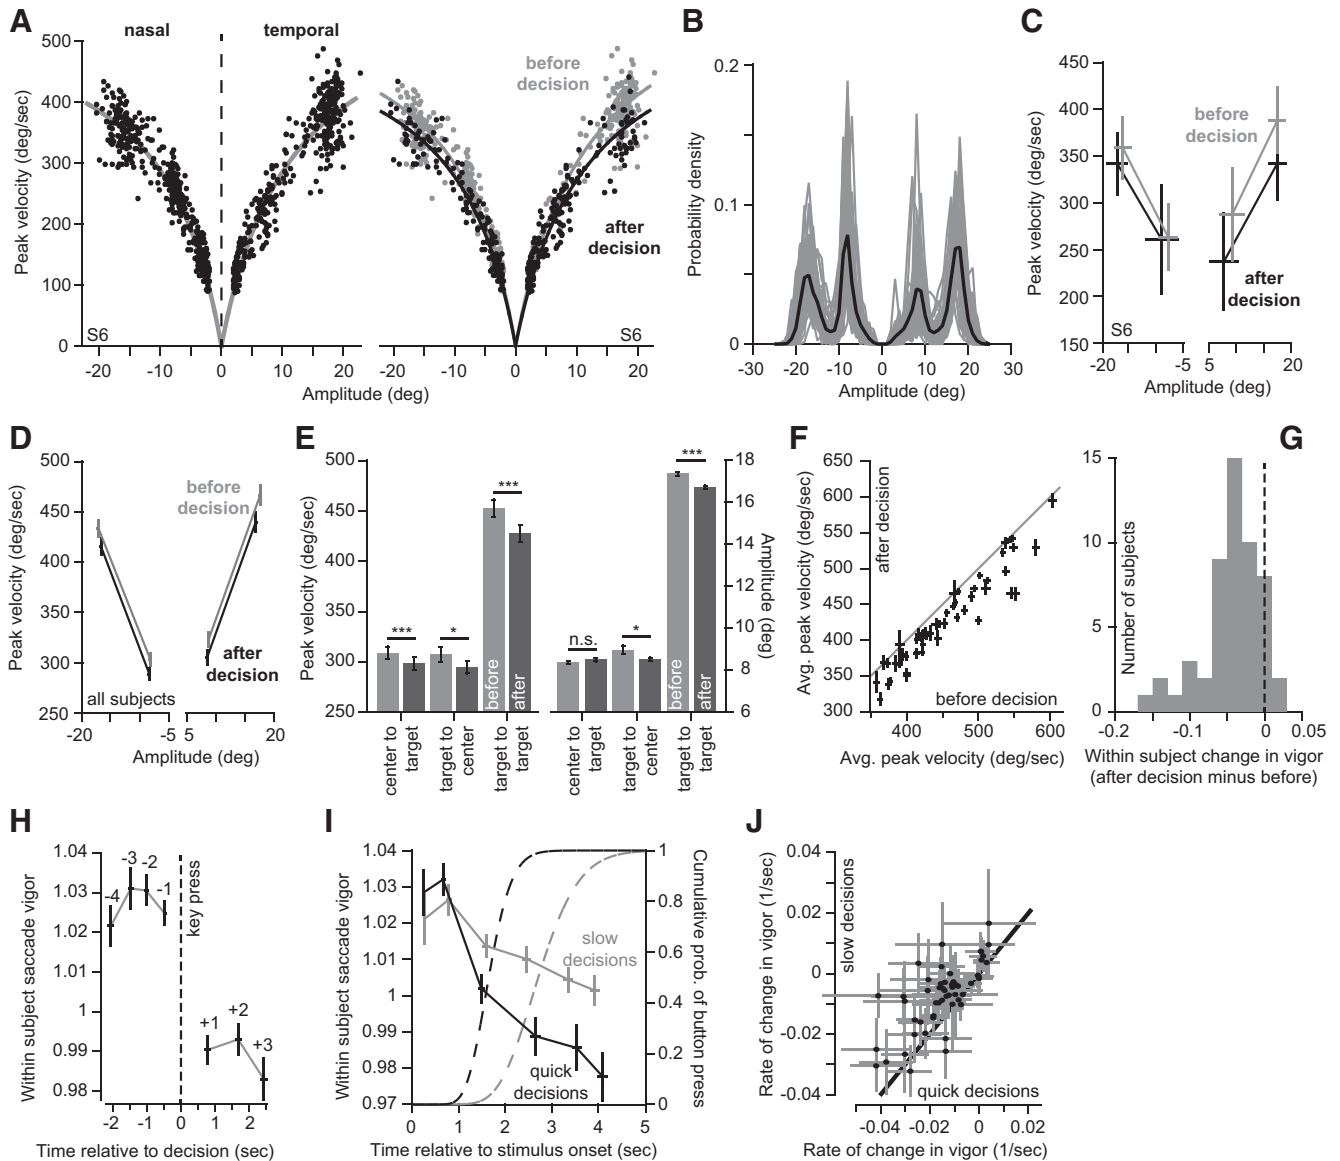


Figure 2. Saccade vigor was higher during deliberation than after decision time. **A**, Velocity–amplitude relationship for a representative participant for saccades made during the 6 s decision period. Data were fit to a hyperbolic function, separately for nasal and temporal saccades. Saccades made before decision time (gray dots) appeared to have a greater velocity than those made after decision time (black dots). **B**, Distribution of saccade amplitudes suggested that there were four groups of saccades made during the decision period: from one stimulus to another ($\pm 18^\circ$ saccades) and from center to one stimulus or back ($\pm 9^\circ$ saccades). Gray lines represent probability density for each participant. Black line represents the across-subject values. Data were binned with step size of 1° . **C**, Velocity–amplitude data from an exemplar subject split by timing of saccade relative to the key press. The saccade amplitudes were binned with bin centers located at $\pm 9^\circ$ and $\pm 18^\circ$, with bin width of 9° . Error bars represent SD. **D**, Across-subject data. The error bars represent SEM and are plotted for both amplitude and velocity. **E**, Across-subject values of amplitude and velocity for saccades made before and after decision time. Statistics refer to within-subject changes ($*p < 0.05$; $***p < 0.001$). Error bars are SEM. **F**, Average saccade peak velocities for each subject before and after decision time. Error bars are SD. **G**, Distribution of within-subject change in vigor with respect to decision time. **H**, Within-subject measure of saccade vigor as a function of timing of saccade with respect to the key press. The number on each data point represents saccade index with respect to the key press. Error bars along the x- and y-axes are SEM. **I**, Within-subject measure of saccade vigor (solid lines) as a function of saccade timing with respect to start of the decision period (stimulus onset). The plot also shows cumulative probability of key press (dashed lines) for slow-decision and quick-decision trials. Error bars are SEM. **J**, Rate of change in vigor for each participant in the quick-decision and slow-decision trials. Error bars are SD.

made 6.2 saccades per trial (25th percentile, 4.6; 75th percentile, 7.7), and they announced their decision at 1.93 ± 0.46 s into the decision period by pressing a key. However, regardless of when the decision was made (indicated by the key press), the stimuli remained on the screen for 6 s. As a result, the participants made saccades to the stimuli both before and after their decision, as illustrated in Figure 1C. To compute probability of saccades during a trial, we aligned the data to decision time and then counted number of saccades performed by a given subject in bins of 0.5 s in duration across all trials. For each bin of 0.5 s duration, we computed probability of saccade for that subject and then com-

puted the across-subject mean and SEM of that probability, as shown in Figure 1D. We found that probability of saccade reached its peak ~ 1 s before decision time but was always significantly >0 during the entire decision period (all p values $< 10^{-9}$).

For each subject, we considered each saccade that they made during the decision period and measured its amplitude and velocity (data for a typical subject are shown in Fig. 2A). Inspection of the data suggested that saccades made before decision time (i.e., during the period of deliberation before key press) may have had a higher velocity than saccades made after (Fig. 2A, right). To explore this question, we examined probability of saccade as a function of ampli-

tude and found it to have four modes (Fig. 2*B*), with peaks at $\pm 9^\circ$ and $\pm 18^\circ$ (the options were displayed at $\pm 10^\circ$ with respect to the central dot). We focused our analysis on those saccades in which one of the stimuli was the goal of the saccade (i.e., center-out or stimulus-to-stimulus saccades) or the fixation dot (stimulus to center saccades). For each saccade, we computed peak velocity as a function of amplitude. The result for a typical participant is shown in Figure 2*C*, and the population average is shown in Figure 2*D*. A within-subject comparison demonstrated that peak velocity and amplitude were significantly higher before decision time than after (Fig. 2*E*; within-subject comparison, peak velocity, $t_{(52)} = 9.49$, $p < 10^{-12}$; amplitude, $t_{(52)} = 9.06$, $p < 10^{-11}$). Indeed, for 96% of the participants, the average peak velocity of saccades was smaller in the postdecision period (Fig. 2*F*).

Because saccade velocity is a function of amplitude, the critical question was whether the higher velocities observed during the deliberation period were attributable to greater vigor or simply attributable to increased amplitude. To answer this question, we accounted for the effect of amplitude on velocity by fitting a hyperbolic function (Eq. 1) to the velocity-amplitude data of all saccades made by each participant (Fig. 2*A*, left) and then used the resulting fit to predict the expected saccade velocity at a given amplitude. The mean \pm SD r values of fits to nasal and temporal saccades were 0.94 ± 0.03 and 0.95 ± 0.03 .

For each saccade during the decision period, we measured its amplitude and computed the ratio of the measured velocity versus the expected velocity. This ratio, our proxy for a within-subject measure of saccade vigor, indicated whether the peak velocity of a given saccade was higher or lower than the expected velocity for that amplitude. For each saccade, we computed its vigor and then computed the average within-subject change in vigor from before decision time to after. We found that a significant number of subjects showed a drop in vigor after decision time (Fig. 2*G*; $t_{(52)} = 8.23$, $p < 10^{-10}$). Saccade vigor as a function of time relative to decision is plotted in Figure 2*H*, where we have numbered each saccade and plotted its timing with respect to key press. There was an $\sim 4\%$ reduction in saccade vigor after decision time (within-subject comparison, $t_{(52)} = 5.97$, $p < 10^{-6}$).

In some trials, the participants took a relatively long time to make a decision, whereas in other trials, the decision was made quickly. For each participant, we computed the median decision time and then labeled each trial for that participant as quick decision or slow decision (decision times for quick and slow trials were 1.40 ± 0.35 and 2.31 ± 0.62 s, mean \pm SD). When we plotted saccade vigor with respect to the onset of the decision period, we found that, in quick-decision trials, saccade vigor declined rapidly, whereas in slow-decision trials, saccade vigor declined gradually (Fig. 2*I*). We tested this difference in vigor as a function of saccade index with a repeated-measures ANOVA and found a significant group by saccade index interaction (Wilks' lambda = 0.604, $F_{(5,46)} = 6.041$, $p < 10^{-3}$). Indeed, a within-subject analysis revealed that the rate of decline in vigor was significantly steeper in quick-decision trials than slow-decision trials (Fig. 2*J*; within-subject t test, $t_{(52)} = 6.41$, $p < 10^{-7}$).

In summary, we found that saccade vigor (as measured via velocity of saccades normalized by amplitude for each subject) was greater during the deliberation period (before the decision was made) than immediately after. Vigor dropped quickly in trials in which participants made a quick decision but dropped slowly in trials in which they took longer to make their decision.

Saccade vigor encoded preference

On each trial, the participants pressed a key to indicate which of the two options they preferred. We asked whether saccade vigor predicted this preference. We separated the saccades based on whether they were directed toward the preferred or the nonpreferred stimulus, in which the preferred stimulus was the option that was eventually chosen by the participant on that trial. Figure 3*A* plots vigor as a function of time of saccade, indexed with respect to key press. It appeared that saccades made before decision time did not differentiate between the preferred and nonpreferred options, except for the last saccade just before key press (Fig. 3*A*). This final saccade took place at 0.520 ± 0.16 s before the key press (mean \pm SD) and had a higher vigor if it was directed to the preferred stimulus (within-subject difference in vigor between the preferred and nonpreferred options, $t_{(52)} = 3.31$, $p = 0.0017$). After the decision, the subsequent saccade also exhibited a greater vigor when it was directed to the preferred stimulus (within-subject difference in vigor, $t_{(52)} = 2.40$, $p = 0.020$). There was no difference in the vigor of saccades to preferred and nonpreferred options outside of this window, suggesting that the encoding of choice preference was a phenomenon that affected vigor only near the time of decision.

One may estimate the degree of preference for one option over the other via the difference in their subjective value. Is the difference in subjective value reflected in the difference in saccade vigor?

To compute subjective value of a given option, we analyzed the choices that the participants made. Figure 3*B* illustrates the choices made by two participants. Participant S21 (Fig. 3*B*, left subplot) often picked the delayed reward when the dollar amount of that option exceeded that of the immediate option by more than \$5. In contrast, participant S41 picked the delayed reward only when the dollar amount of that option exceeded that of the immediate option by more than \$20. We fitted these data to Equation 2, resulting in an estimate of the point of subjective equivalence for each participant (Fig. 3*B*, dashed line). For participant S21, a difference of \$4 made the delayed reward equivalent to the immediate reward. For participant S41, a difference of \$23 was required to make the delayed reward equivalent to the immediate reward.

How robust was this estimate of subjective equivalence? To answer this question, we imagined that, for each participant, the decision should be most difficult when the two options differed in value by the amount specified by the point of subjective equivalence. For example, for participant S21, the most difficult choice should be in trials in which the delayed reward was \$4 greater than the immediate reward. A proxy for this difficulty is the time that the participants needed to make their choice. We measured the time from stimulus display to key press and have plotted the results in Figure 3*C*. For each participant, we fitted their time to key press with a Gaussian and estimated its center, resulting in the difference between delayed and immediate reward that produced the longest deliberation time. As a result, the explicit choices that participants made provided one measure of subjective equivalence (Fig. 3*B*), and the time they took to make that choice provided a second measure (Fig. 3*C*). The two measures were well correlated (Fig. 3*D*; $r^2 = 0.68$, $p < 10^{-12}$). This result indicated that the point of subjective equivalence derived from the explicit choices was reasonable and robust.

We next used the decision-based estimate of subjective equivalence to compute the rate of temporal discounting (parameter k in Eq. 3), which then allowed us to compute the (relative) subjective value of the delayed reward for each participant (assuming a linear utility function). Focusing on the

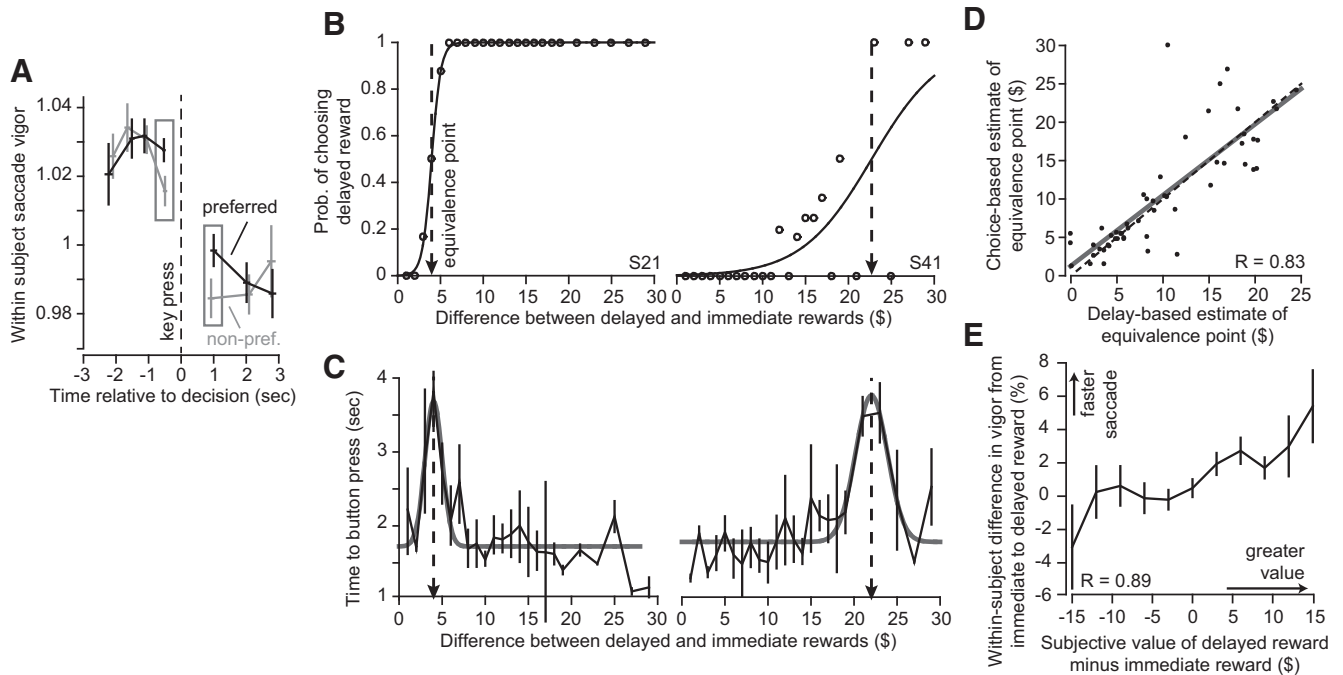


Figure 3. Saccade vigor encoded choice preference just before time of decision. **A**, For each participant, we labeled each saccade based on whether gaze was directed toward the preferred or nonpreferred option. We then indexed saccades relative to the timing of the key press. Around the time of decision, vigor of saccade made to the preferred target was higher than vigor of saccade made to the nonpreferred target. Gray boxes denote saccades immediately before and after the key press. **B**, Quantifying point of subjective equivalence via explicit decisions. The left and right columns show the probability of choosing the delayed reward option as a function of the difference in magnitude of the rewards offered for participants S21 and S41. S41 required a larger amount of delayed reward to switch preference from the immediate to delayed reward and thus tended to favor the immediate reward more than S21. We fitted a two-parameter sigmoid function to the subjects' choice data and estimated the point at which subjects switched preference from the immediate to delayed reward, which we labeled the point of subjective equivalence (\$3.97 for S21 and \$22.71 for S41). **C**, Delay to decision reflected difficulty of decision-making. The plots show the distributions of time to key press for the same two subjects. We computed the average time to key press at each difference in reward (solid black vertical lines represent 25th and 75th percentiles) and then fitted a Gaussian function to the data (gray curve). The location of the mean of the Gaussian represents the difference in reward offerings for which the subjects took the longest time to make a decision (\$4.01 and \$22.02 for the two subjects). **D**, Quantifying robustness of estimate of equivalence point. For each participant, we compared the equivalence point estimated from their explicit decisions (as in **B**) with the equivalence point estimated from their time to key press (as in **C**). Each data point is one participant. The dashed line represents equality between the two measures. Gray line is the best linear fit. **E**, For saccades made immediately before and after decision, within-subject difference in saccade vigor was related to within-subject difference in subjective value of the delayed and immediate rewards. Error bars are SEM.

two saccades made immediately before and after decision time, we measured vigor when the participants looked at the immediate reward and compared it with vigor when they looked at the delayed reward. The difference in vigor is plotted on the y -axis in Figure 3E. Vigor increased from the immediate to the delayed reward as a function of the difference in the subjective value of the delayed reward versus the immediate reward ($r = 0.89$, $p = 0.0002$). That is, around the time of decision, vigor of the saccade that directed gaze toward a stimulus was correlated with the subjective value that the brain assigned to that stimulus.

We found that the saccade made just before decision time tended to be to the preferred option. In Figure 4A, we have plotted probability that the saccade was to the preferred option, given that the participant made a saccade, computed over time bins of 0.5 s in duration. This conditional probability became significantly greater than chance ~ 1 s before decision time and reached its peak at the final time bin before decision time (within-subject comparison, $p < 10^{-11}$). Furthermore, it appeared that, as time passed after the decision, the participants were more likely to saccade to the chosen option than the nonchosen option (Fig. 4A, postdecision region).

These observations suggested that saccade patterns may be used as an implicit measure of preference. How well does this implicit measure predict the eventual choice? To check for this, we compared the choices that subjects made to the choices that

would be expected if the saccade just before decision time was used as a marker of preference. We computed an implicit equivalence point based on the option that was the target of the last saccade before decision time and found that this implicit equivalence point matched well with the explicit equivalence point as estimated from the actual choices that the subjects made (Fig. 4B; $r = 0.73$, $p < 10^{-9}$; Arieli et al., 2011; Rangel and Clithero, 2013). Thus, the target of the final saccade before decision was an excellent predictor of the explicit choices that participants made.

In summary, during the deliberation period, vigor of the saccades that placed each of the two stimuli on the fovea was similar but diverged at ~ 0.5 s before decision time, becoming larger for the preferred stimulus. As the difference between the subjective values of the delayed and immediate rewards increased, so did the difference in vigor in the movements made toward the two options. This set of findings is surprising given that saccades had no bearing on the actual outcome of the decision.

Between-subject differences in saccade vigor

In addition to within-subject changes in saccade vigor during the decision period, there were also between-subject differences in the saccadic eye movements: for a given saccade amplitude, some participants consistently moved their eyes with high velocity, whereas others consistently moved their eyes with low velocity. That is, there were between-subject differences in saccade vigor.

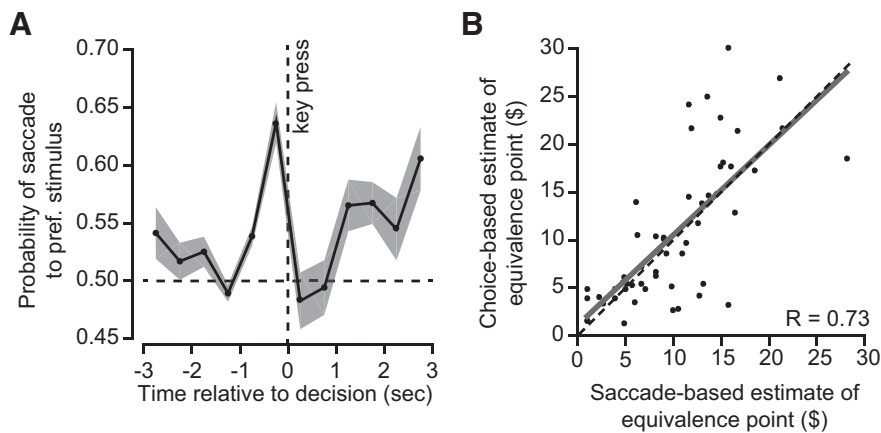


Figure 4. The stimulus that was the target of gaze just before the key press was often the option that was eventually chosen. **A**, Probability of the saccade target being the chosen option, given that a saccade was made to one of the two stimuli at that time bin. Time bins are 0.5 s in length. **B**, Quantifying robustness of estimate of equivalence point. For each participant, we used the last saccade before decision time as the predictor of the preferred option and used that result to compute an equivalence point (labeled as saccade-based estimate). Error bars are SEM.

We quantified this difference and asked whether it was related to differences in decision making.

We began by fitting Equation 1 to the velocity–amplitude data of each participant. For participant n , this produced parameter values $\hat{\alpha}_n$ and $\hat{\beta}_n$. We found the median of the $\hat{\alpha}$ and $\hat{\beta}$ distributions across all participants, producing $\bar{\alpha}$ and $\bar{\beta}$. The values of $\bar{\alpha}$ and $\bar{\beta}$ were 690.4 and 0.089, and 764.3 and 0.082, for nasal and temporal saccades, respectively. We used this estimate to produce a canonical relationship between amplitude and velocity across the entire population:

$$\bar{v} = \bar{\alpha} \left(1 - \frac{1}{1 + \bar{\beta}x} \right). \quad (4)$$

We used the above relationship to quantify the relative vigor of saccades in one participant compared with another. We followed the procedure described previously (Choi et al., 2014): we refitted each participant's saccade velocity–amplitude data to a one-parameter scaling function of the canonical function:

$$\hat{v}_n(x) = \lambda_n \bar{v}(x). \quad (5)$$

Parameter λ_n is the between-subject measure of vigor for subject n . When we have $\lambda_n > 1$, it indicates that the saccades of participant n are generally faster than the population median.

Figure 5A (left) illustrates saccade peak velocity as a function of saccade amplitude for two participants. Participant S14 had consistently faster saccades than participant S6. The right of Figure 5A shows the canonical function (dashed line, representing the population median) and the function representing the data for each participant (derived from Eq. 5). To quantify goodness of fit, we computed correlation coefficients, reflecting the ability of the one parameter model of Equation 5 to account for the saccade velocity/amplitude data of each subject. The results are illustrated in Figure 5B. For every subject, the fit between the model used to estimate between-subject saccade vigor and actual velocities was significant at a level of $p < 0.00001$.

Using this measure of between-subject saccade vigor, we asked whether individuals who moved with greater vigor were distinguishable in their patterns of decision making. We focused on between-subject differences in impulsivity (i.e., the equivalence point between the immediate and delayed reward). For example, participant S41

has a larger equivalence point than participant S21 (Fig. 3B). This translates into a larger temporal discount rate, implying a greater impulsivity. However, we did not find a significant relationship between vigor and impulsivity (Fig. 5C; $\rho = -0.24$, $p = 0.078$), nor did we find any relationship between vigor and discount factor k ($\rho = -0.25$, $p = 0.074$). Therefore, in this task, the between-subject differences in saccade vigor were not a predictor of differences in decision-making.

Discussion

We examined saccades that participants made as they considered two monetary options: a small reward to be obtained immediately versus a larger reward to be obtained at 30 d. We found that saccade vigor, a within-subject measure of peak velocity normalized by amplitude, was greater during the deliberation period than immediately after. Vigor dropped

rapidly in trials in which participants made a quick decision but dropped slowly in trials in which they took longer. Among the saccades made just before and just after the decision, saccades to the preferred option exhibited a greater vigor than saccades to the nonpreferred option. The participants signaled their decision with a key press ~ 0.5 s after saccade vigor diverged between the two options. The disparity between vigor of saccades to the two options became larger as the difference in the subjective values of the two options increased. Therefore, during decision-making, the subjective value that the brain assigned to a stimulus influenced the vigor with which the eyes moved toward that stimulus.

Neural basis of vigor

The vigor with which a saccade is performed is associated with 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 movement 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, substantia nigra pars reticulata (SNr). SNr cells constantly inhibit 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), and reward modulates the depth of this pause (Handel and Glimcher, 1999). 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, some cells in the caudate nucleus influence the discharge of SNr neurons directly, whereas other cells do so indirectly via their projections to the external segment of globus pallidus (GPe). Caudate cells receive dopamine projections and generally fire more before a rewarding saccade (Kawagoe et al., 1998). 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 $\sim 30\%$ (Kori et al.,

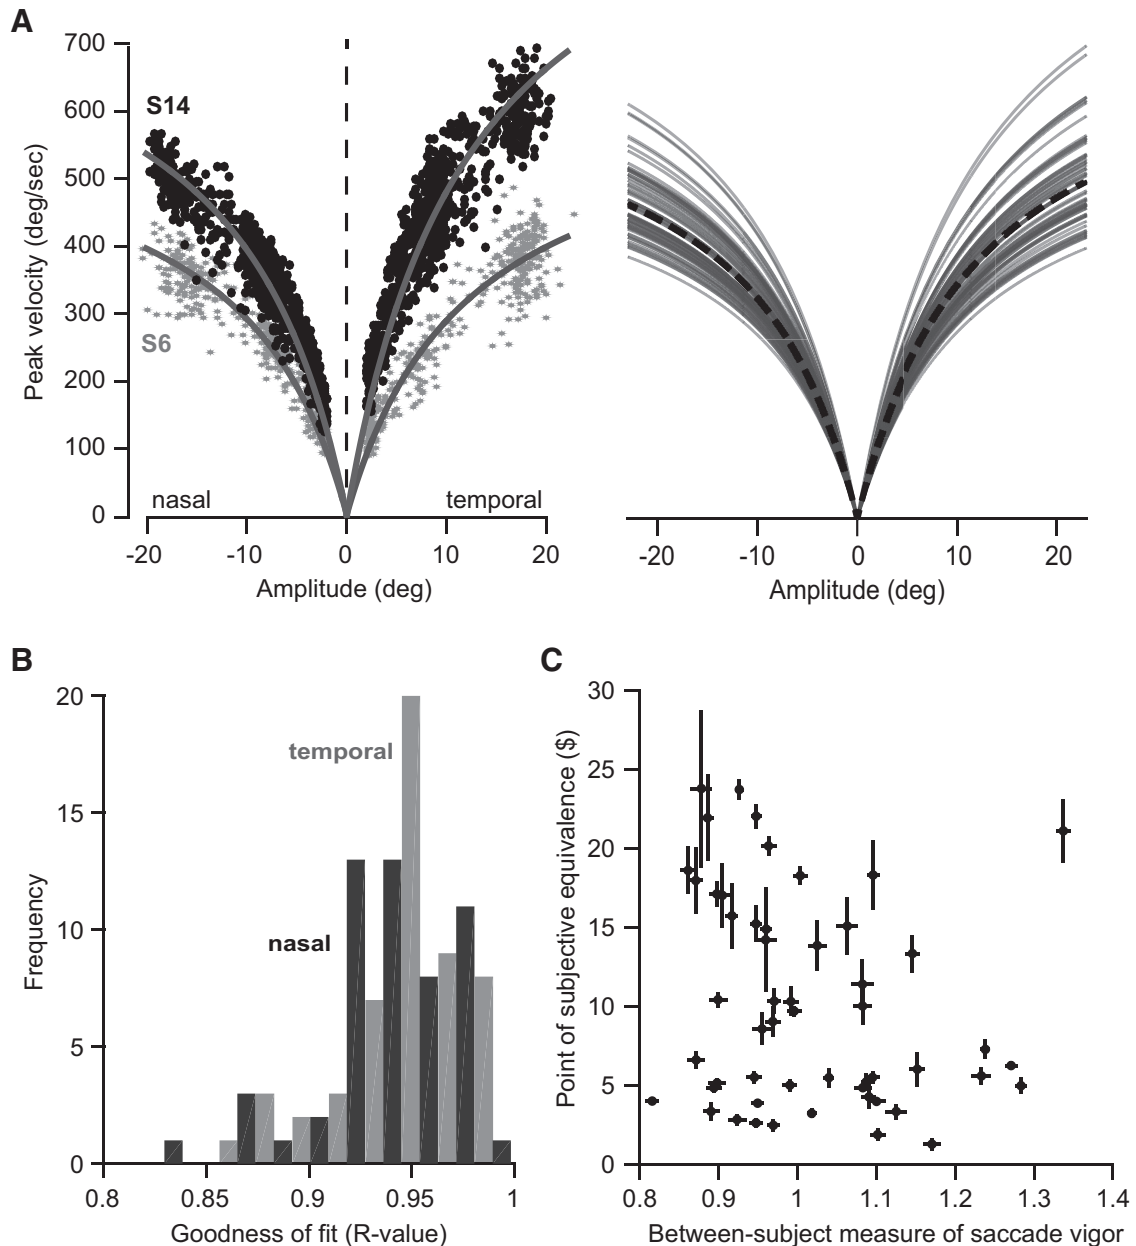


Figure 5. Between-subject differences in vigor. **A**, Peak velocity as a function of amplitude for two representative participants. Participants S14 made consistently faster saccades than participant S6. For each subject, we fitted the velocity–amplitude data for nasal and temporal saccades separately to a two-parameter hyperbolic function. Using these parameters, we computed the average relationship between velocity and amplitude across the population (dashed lines, right). Finally, we computed the between-subject measure of vigor as the scaling of the population average velocity–amplitude relationship. **B**, Goodness of fit of the hyperbolic model to the velocity–amplitude relationship in each subject. Dark and light gray bars represent fits to nasal and temporal saccades, respectively. **C**, Between-subject measure of saccade vigor was not a predictor of subjective equivalence between immediate and delayed rewards.

1995). GPe cells inhibit 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). Therefore, control of vigor is partly associated with the amount of dopamine in the basal ganglia, modulating activity of caudate, affecting the depth of pause in the SNr.

During decision-making, temporal discounting is also associated with release of dopamine. When an animal makes a decision between a small magnitude, short-delay reward and a large magnitude, long-delay reward, dopamine cells fire in response to each stimulus by an amount that correlates with the subjective value of that stimulus (Kobayashi and Schultz, 2008). Together, it appears

that some of the circuits that are critical for control of vigor are also influenced by a neurotransmitter that has been linked to subjective valuation of reward. This link, we speculate, may be the reason for the modulation of saccade vigor during the deliberation process.

Subjective value versus motivational salience

We found that vigor reflected the subjective value of the stimulus that acted as the goal of the movement. However, an alternate hypothesis is that vigor is a reflection of the motivational salience of the stimulus, predicting that because motivational salience associated with loss of \$10 is greater than loss of \$5, vigor will be greater toward $-\$10$ than $-\$5$, despite the fact that the subject-

tive value of $-\$5$ is greater than $-\$10$. The subjective value hypothesis predicts the opposite: vigor should be higher for $-\$5$ than $-\$10$.

Kobayashi et al. (2006) asked monkeys to view a cue that determined whether the upcoming saccadic eye movement was a reward trial (apple juice), punishment trial (air puff), or neutral (sound). Motivational requirements of the trials were highest for air puff and juice and lowest for the neutral condition, as evidenced by the fact that correct performance rates were highest in the reward and air-puff trials and lowest in the neutral trials. However, the subjective values of the trials were highest for juice, lowest for air puff, and in between for neutral. The authors observed that saccade peak velocity was highest for the reward trials, lowest for air puff, and in between for neutral trials. This implies that saccade vigor is a reflection of subjective value, not motivational salience.

In a temporal discounting task conducted in monkeys and with lateral intraparietal sulcus (LIP) recordings, we reported that the activity of LIP neurons varied with the subjective value of the delayed reward (Louie and Glimcher, 2010). By varying the delay to the time of reward acquisition, we found that the subjective value of the delayed reward could be reduced by up to 60% in one monkey and up to 40% in the other monkey. During the delay period in both the forced-choice and free-choice trials, activity of LIP neurons was modulated as a function of the subjective value of the stimulus, with a gain of nearly 1. In contrast, here we found that change in saccade vigor was a maximum of 7% (Fig. 3E) compared with a change in subjective value of $\sim 35\%$, a gain of 0.2. Therefore, we speculate that activity of LIP is more closely related to the utility of the action compared with the vigor of that action.

Modulation of vigor during decision-making

In our task, saccades were not associated with reward but were a means by which the brain acquired information for the purpose of making the decision. This is in contrast to many previous experiments in which the act of making a saccade was itself associated with reward (Kawagoe et al., 1998; Takikawa et al., 2002; Kobayashi and Schultz, 2008; but see Thura and Cisek, 2014; Thura et al., 2014). Despite this, in our task, saccade vigor was modulated by subjective value of the stimulus. Our results suggest that, during decision-making, actions that acquire information relevant to the eventual decision have a subjective value associated with them, as evidenced by the vigor of that action.

This view provides a potential explanation as to why vigor dropped after the choice. We speculate that saccades that were made during the deliberation period had a greater vigor because each movement acquired information relevant to the eventual reward. Once the choice had been indicated, the same actions no longer acquired relevant information. In this sense, the subjective values of the movements performed during deliberation were higher than those performed after.

A recent experiment by (Thura et al., 2014) noted that urgent decisions were followed by more vigorous movements. The authors suggested that a rising urgency signal combined with the process of evidence accumulation, forcing a hastier decision in some circumstances and a more deliberate decision in other circumstances. Vigor was affected by the rate of rise of this urgency signal. These results complement our findings by demonstrating that, in addition to stimulus value, other contextual factors, such as rate of reward, can affect decision-making and movement vigor (Haith et al., 2012).

Between-subject differences in vigor

As in our previous work (Choi et al., 2014), here we found that there were between-subject differences in saccade vigor: some individuals consistently moved their eyes with greater velocity than other individuals. Previously, we found that individuals who had greater saccade vigor were less willing to wait to increase their probability of success. However, here in a value-based decision-making task, we found no relationship between temporal discounts rates and saccade vigor.

There are a number of reasons that could underlie this disparity. To measure temporal discounting, in the previous study (Choi et al., 2014), we designed a task in which each choice had an immediate consequence, acting as an operant reinforcement on the next choice. In contrast, here we measured temporal discounting in a task in which choices had consequences that were not experienced until after experiment completion. Although both types of approaches produce measures of temporal discounting, they produce inconsistent results in the same person (Hyten et al., 1994) and produce greatly differing discount rates (Navarick, 2010). Therefore, fundamental differences in how one measures temporal discounting during decision-making may underlie differences in the two studies.

Another possibility is that, in a value-based decision-making task without immediate consequences, participants may have more control over their explicit decisions, a phenomenon commonly referred to as impulse control (Ainslie, 1975; Bechara, 2005). For example, it has been shown that Parkinson's disease patients who have been treated with a dopamine agonist have both increased saccade vigor (Nakamura et al., 1991) and a higher propensity for impulse control disorders (Weintraub et al., 2010). Thus, it seems possible that modulation of movement vigor is a measure that can be used to ascertain choice preference, even when subjects may be concealing their true preferences.

Conclusions

During the deliberation period of a decision-making task, vigor was similar as saccades were made between the two options but diverged ~ 0.5 s before decision time, becoming greater for the option that was eventually chosen. Therefore, vigor of the movement that brought the gaze toward an option was affected by the value (or salience) that the brain assigned to that option. Overall, our results suggest a link between the neural mechanism that assigns value (or salience) to a stimulus and the mechanism that controls vigor of movements toward that stimulus.

References

- Ainslie G (1975) Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychol Bull* 82:463–496. [CrossRef Medline](#)
- Arieli A, Ben-Ami Y, Rubinstein A (2011) Tracking decision makers under uncertainty. *Am Econ J Microecon* 3:68–76. [CrossRef](#)
- Bechara A (2005) Decision making, impulse control and loss of willpower to resist drugs: a neurocognitive perspective. *Nat Neurosci* 8:1458–1463. [CrossRef Medline](#)
- Choi JE, Vaswani PA, Shadmehr R (2014) Vigor of movements and the cost of time in decision making. *J Neurosci* 34:1212–1223. [CrossRef Medline](#)
- Green L, Myerson J (2004) A discounting framework for choice with delayed and probabilistic rewards. *Psychol Bull* 130:769–792. [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](#)
- 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](#)
- 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](#)
- 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](#)
- Kobayashi S, Schultz W (2008) Influence of reward delays on responses of dopamine neurons. *J Neurosci* 28:7837–7846. [CrossRef Medline](#)
- Kobayashi S, Nomoto K, Watanabe M, Hikosaka O, Schultz W, Sakagami M (2006) Influences of rewarding and aversive outcomes on activity in macaque lateral prefrontal cortex. *Neuron* 51:861–870. [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](#)
- Lempert KM, Glimcher PW, Phelps EA (2015) Emotional arousal and discount rate in intertemporal choice are reference dependent. *J Exp Psychol Gen* 144:366–373. [CrossRef Medline](#)
- Louie K, Glimcher PW (2010) Separating value from choice: delay discounting activity in the lateral intraparietal area. *J Neurosci* 30:5498–5507. [CrossRef Medline](#)
- Matsumoto M, Hikosaka O (2007) Lateral habenula as a source of negative reward signals in dopamine neurons. *Nature* 447:1111–1115. [CrossRef Medline](#)
- Mazur JE (1987) An adjusting procedure for studying delayed reinforcement. In: *The effect of delay and of intervening events on reinforcement value* (Commons ML, Mazur JE, Nevin JA, Rachlin H, eds), pp 55–73. Hillsdale, NJ: Erlbaum.
- 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* 111:559–562. [CrossRef Medline](#)
- Navarick D (2010) Discounting of delayed reinforcers: measurement by questionnaires versus operant choice procedures. *Psychol Rec* 54:85–94.
- Opris I, Lebedev M, Nelson RJ (2011) Motor planning under unpredictable reward: modulations of movement vigor and primate striatum activity. *Front Neurosci* 5:61. [CrossRef Medline](#)
- Rangel A, Clithero J (2013) The computation of stimulus values in simple choice, Chap 8. In: *Neuroeconomics: decision making and the brain*, Ed 2, pp 125–147. San Diego: Academic.
- Rousseeuw PJ, Croux C (1993) Alternatives to the median absolute deviation. *J Am Stat Assoc* 88:1273–1283. [CrossRef](#)
- Sackaloo K, Strouse E, Rice MS (2015) Degree of preference and its influence on motor control when reaching for most preferred, neutrally preferred, and least preferred candy. *OTJR (Thorofare N J)* 35:81–88. [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](#)
- 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](#)
- Thura D, Cisek P (2014) Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* 81:1401–1416. [CrossRef Medline](#)
- Thura D, Cos I, Trung J, Cisek P (2014) Context-dependent urgency influences speed–accuracy trade-offs in decision-making and movement execution. *J Neurosci* 34:16442–16454. [CrossRef Medline](#)
- Weintraub D, Koester J, Potenza MN, Siderowf AD, Stacy M, Voon V, Whetteckey J, Wunderlich GR, Lang AE (2010) Impulse control disorders in Parkinson disease: a cross-sectional study of 3090 patients. *Arch Neurol* 67:589–595. [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](#)