

## RESEARCH ARTICLE | *Higher Neural Functions and Behavior*

# Saccade vigor and the subjective economic value of visual stimuli

Tehrim Yoon,<sup>1</sup> Afareen Jaleel,<sup>1</sup> Alaa A. Ahmed,<sup>2</sup> and  Reza Shadmehr<sup>1</sup>

<sup>1</sup>Laboratory for Computational Motor Control, Department of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, Maryland; and <sup>2</sup>Departments of Integrative Physiology and Mechanical Engineering University of Colorado, Boulder, Colorado

Submitted 12 November 2019; accepted in final form 29 April 2020

**Yoon T, Jaleel A, Ahmed AA, Shadmehr R.** Saccade vigor and the subjective economic value of visual stimuli. *J Neurophysiol* 123: 2161–2172, 2020. First published May 6, 2020; doi:10.1152/jn.00700.2019.—Decisions are made based on the subjective value that the brain assigns to options. However, subjective value is a mathematical construct that cannot be measured directly, but rather is inferred from choices. Recent results have demonstrated that reaction time, amplitude, and velocity of movements are modulated by reward, raising the possibility that there is a link between how the brain evaluates an option and how it controls movements toward that option. Here, we asked people to choose among risky options represented by abstract stimuli, some associated with gain (points in a game), and others with loss. From their choices we estimated the subjective value that they assigned to each stimulus. In probe trials, a single stimulus appeared at center, instructing subjects to make a saccade to a peripheral target. We found that the reaction time, peak velocity, and amplitude of the peripherally directed saccade varied roughly linearly with the subjective value that the participant had assigned to the central stimulus: reaction time was shorter, velocity was higher, and amplitude was larger for stimuli that the participant valued more. Naturally, participants differed in how much they valued a given stimulus. Remarkably, those who valued a stimulus more, as evidenced by their choices in decision trials, tended to move with shorter reaction time and greater velocity in response to that stimulus in probe trials. Overall, the reaction time of the saccade in response to a stimulus partly predicted the subjective value that the brain assigned to that stimulus.

**NEW & NOTEWORTHY** Behavioral economics relies on subjective evaluation, an abstract quantity that cannot be measured directly but must be inferred by fitting decision models to the choice patterns. Here, we present a new approach to estimate subjective value: with nothing to fit, we show that it is possible to estimate subjective value based on movement kinematics, providing a modest ability to predict a participant's preferences without prior measurement of their choice patterns.

decision making; motor control; subjective value; vigor

---

“A true theory of economy can only be attained by going back to the great springs of human action—the feelings of pleasure and pain.” William Stanley Jevons (1866)

## INTRODUCTION

Theory of subjective value was introduced in the 19th century to account for the fact that in voluntary transactions, each party values the goods, labor, or money that they receive

more than the goods, labor, or money that they provide (Jevons 1866; Menger 1871). The theory posited that subjective value is not specified by an objective property of the good, but rather the incremental increase in pleasure that an individual assigns to acquisition of that good (Jevons 1866). Although subjective valuation is an important aspect of behavioral economics, it is an abstract quantity that cannot be measured directly. Rather, it must be inferred from decisions that individuals make (von Neumann and Morgenstern 1944), often in scenarios involving lotteries and risky options.

A serendipitous discovery in motor neuroscience has been the observation that factors that affect preference, such as reward and effort, also affect movements (Shadmehr et al. 2010, 2019). For example, in goal-directed movements, people and other primates move with a shorter reaction time and greater velocity toward stimuli that they associate with greater gain (Kawagoe et al. 1998; Milstein and Dorris 2007; Summerside et al. 2018; Xu-Wilson et al. 2009; Yoon et al. 2018). Recent work (Sedaghat-Nejad et al. 2019) has shown that if presentation of a stimulus results in a reward prediction error, the movement that ensues tends to be expressed with greater vigor (defined as the reciprocal of reaction time plus movement duration). Reward prediction error is the principal variable that modulates dopamine release (Bayer and Glimcher 2005; Schultz et al. 1997), and, intriguingly, stimulation of dopamine around movement onset tends to increase movement acceleration (da Silva et al. 2018). Thus, both the process of learning subjective value from reward prediction error and control of movement speed depend on dopamine, raising the possibility that the vigor with which an individual moves toward an option is partly influenced by the subjective value that they assign to that option.

Previous work has established that when people are presented with a decision between two options, their deliberation time is a measure of their strength of preference: participants typically decide sooner if they prefer one stimulus much more than another (Konovalov and Krajbich 2019; Spiliopoulos and Ortmann 2018). Thus, these works have demonstrated that certain aspects of behavior during decision making are related to the difference in the subjective value of the two options.

Here, we asked a different question: suppose one could only observe movements during presentation of a single stimulus, but not during decision making. Can one infer subjective value from the movements in response to single stimuli A and B and then predict choice when the participant decides between A

---

Correspondence: R. Shadmehr (shadmehr@jhu.edu).

and B? To do this, one would have to predict choice despite having a model that has never observed choice (thus, nothing to fit). If this were possible, how well might movement kinematics in single stimulus trials allow one to predict subjective value, and thus choice in decision trials?

It is possible that this approach will fail because movements may not reflect subjective value, but rather a measure of attention. For example, both the stimulus that promises a gain and the stimulus that foretells a penalty are important and will garner more attention than stimuli that promise smaller gain and loss. In this scenario, movement kinematics will not vary monotonically with subjective value, but rather produce a U-shaped function, becoming large for both gains and losses. In such a scenario, a kinematics-based model of subjective value will fail to predict choice.

There are plausible neural mechanisms that support this alternate hypothesis. Saccade reaction time and velocity are partly modulated by the excitatory inputs that the superior colliculus receives from the cortical regions which compute subjective value: the frontal eye field (FEF) (Glaser et al. 2016; Hanes and Schall 1996; Heitz and Schall 2012) and the lateral intraparietal area (LIP) (Louie and Glimcher 2010; Platt and Glimcher 1999s). LIP neurons that encode stimulus value exhibit greater activity both when the stimulus promises a large reward and when the stimulus promises a large penalty (Leathers and Olson 2012). Some of these neurons exhibit sensitivity to both novelty and value (Foley et al. 2014). Furthermore, some dopamine neurons increase their activity when the stimulus promises reward, whereas others increase their activity for both punishment and reward (Matsumoto and Hikosaka 2009). Thus, the neural activity that could modulate saccade kinematics shows positive sensitivity to gain, as well as loss. This leads us to the question of whether saccade kinematics monotonically reflects valuation over a range that includes both losses and gains, or is kinematics a U-shaped function of value.

Here, participants learned to associate value to 10 abstract stimuli, each paired with a different magnitude of loss or gain. Because the task involved learning, it induced between-participant variability in assignment of subjective value. In decision trials, the participants deliberated between various stimuli and made a choice, from which we also inferred the subjective value that they assigned to each stimulus. In probe trials, we presented a randomly chosen stimulus and measured saccade kinematics in response to it. We found that in probe trials, saccade reaction time was lowest for stimuli that promised a loss and highest for stimuli that promised a gain. Thus, subjective value could be estimated from reaction time. This kinematic based estimate correctly predicted ~60% of the choices made in decision trials.

## MATERIALS AND METHODS

Healthy participants ( $n = 24$ ,  $26.3 \pm 8.2$  yr old, mean  $\pm$  SD, 8 women) with no known neurological disorders and normal color vision sat in a well-lit room in front of an LED monitor ( $59.7 \times 33.6$  cm,  $2560 \times 1440$  pixels, light gray background, frame rate 144 Hz) placed at a distance of 35 cm. Their head was restrained using a bite bar. They viewed visual stimuli on the screen, and we measured their eye movements using an EyeLink 1000 (SR Research) infrared recording system (sampling rate 1 kHz). Only the right eye was tracked. All participants were naïve to the paradigm. The experiments

were approved by the Johns Hopkins University School of Medicine Institutional Review Board, and all participants signed the written consent form approved by the board. Participants were paid \$15/hour regardless of any behavioral outcome. One participant was excluded from the results presented here because their performance in the task was at chance level, suggesting that they did not learn to assign value to the various stimuli.

**Stimulus properties.** We performed an experiment in which people learned the value of 10 abstract visual stimuli (Fig. 1A). Each stimulus was a  $2^\circ \times 2^\circ$  colored box, designated with a + or – (Fig. 1B). Each stimulus was randomly assigned to a point distribution, with a mean that ranged from loss of 5 points to gain of 5 points. The points associated with each color were selected randomly on each trial from a beta distribution with parameters  $\alpha = \beta = 2$ , scaled so that each color was associated with a single mean:  $-5, -4, \dots, +5$ . The plus and minus indicator at the center of the stimulus noted the sign of the mean of the distribution. The color-to-point relationship was selected randomly for each participant, but remained consistent throughout the experiment. For example, the plus yellow box in Fig. 1B was associated with a distribution with mean equal to gain of 4 points, and the minus yellow box was associated with mean equal to loss of 4 points. In addition to these 10 colored boxes, a black box with “0” at its center was associated with exactly 0 points. Thus, the experiment design employed abstract stimuli that the participants learned to associate with points. We hoped that this would produce a wide diversity in subjective values that the participants assigned to a given stimulus, allowing us to test whether movement kinematics was a predictor of the between participant differences in subjective value.

**Decision trials.** The experiment contained two types of trials, randomly intermixed. Both types of trials (Fig. 1A) began with a center fixation period that lasted for 1 s and ended with a beep (1 kHz).

In decision trials, the fixation point was replaced with three different colored boxes (stimuli). Importantly, all three colored stimuli appeared within  $2^\circ$  of the central location. In addition to the colored stimuli at center, there were two dots ( $0.5 \times 0.5^\circ$ ), one at  $+20^\circ$  and the other at  $-20^\circ$  along the horizontal axis.

The sure bet was the colored box that appeared alone on one side of fixation (Fig. 1A). The risky bet was the pair of boxes that appeared together on the other side of fixation. If the participant chose the sure bet, she would indicate that choice by making a saccade to the dot that was on the same side as the single stimulus. She would then receive with 100% probability the points associated with that stimulus. If the participant chose the risky bet, she would indicate that choice by making a saccade to the dot that was on the same side as the double stimulus. She would receive the points associated with one of the two stimuli (50% probability).

The participant had 5 s to indicate her choice (the sure bet or the risky bet) by making a saccade to one of the dots. Once the saccade concluded, the stimuli at center were erased and the trial consequences were displayed for 1 s: the earned stimulus was displayed at the dot location along with text that indicated the number of points acquired. The points were drawn from the random distribution associated with the colored box. Failure to make a choice within the time limit resulted in loss of 10 points. The trial ended with the display of the colored box and the amount of points gained or lost for that trial (duration of 1 s).

In a decision trial, we randomly picked 3 stimuli from among the 11 stimuli. We presented the medium-valued stimulus as the sure bet and the other two stimuli (one loss, and the other gain) as the risky bet. Participants were not provided any information about the value of the stimuli and thus had to make their decisions solely based on consequences of previous trials. The side that represented the sure bet was random and chosen with equal left-right frequency for each block.

**Probe trials.** Probe trials were randomly intermixed with decision trials. In probe trials, the fixation point was removed, a single stimulus (chosen at random from the 10 colored boxes) was displayed at center,

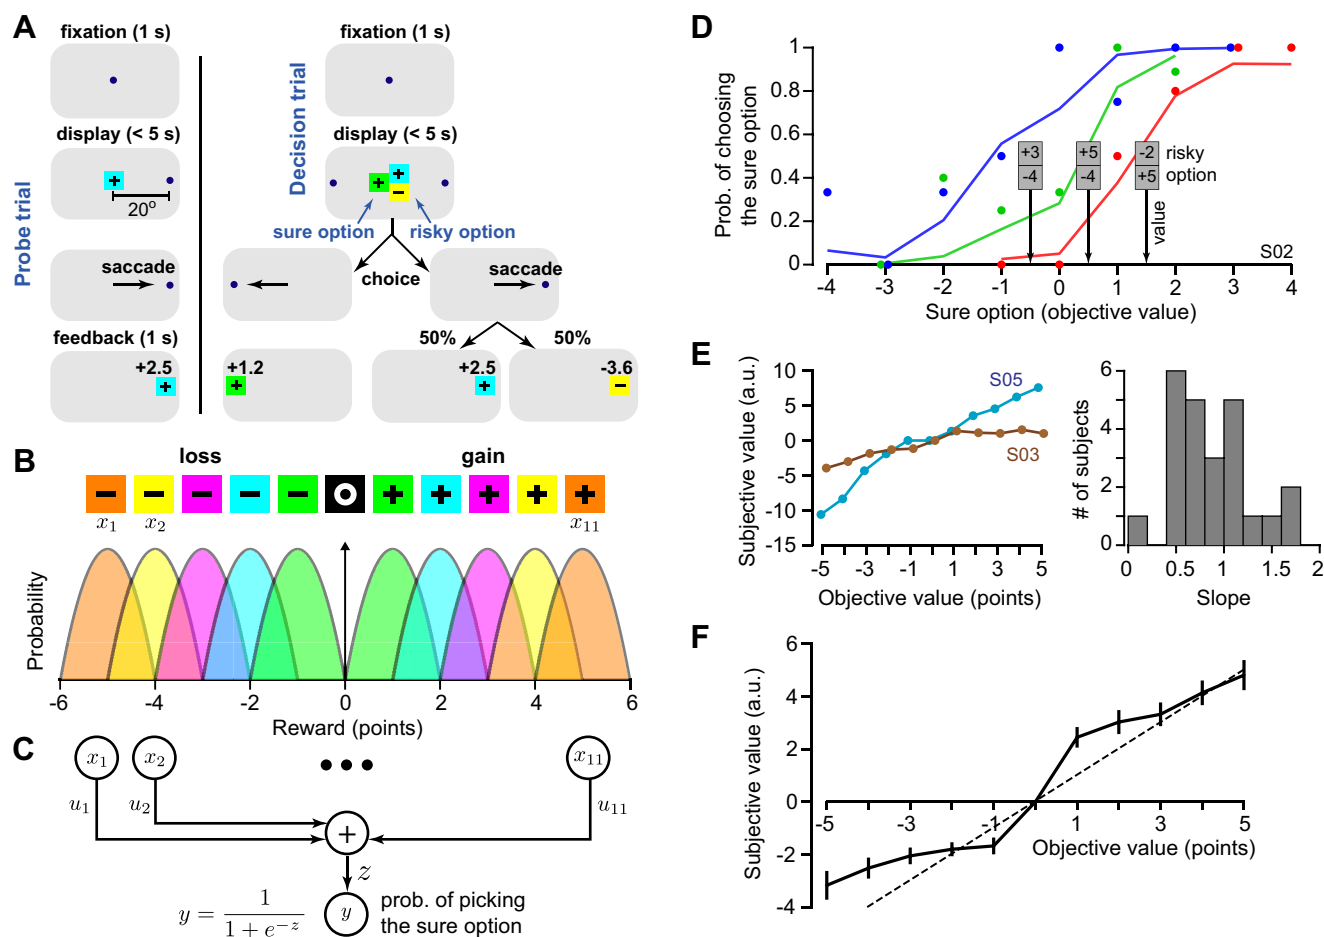


Fig. 1. Estimating subjective value of abstract stimuli. **A**: in probe trials, a single stimulus was presented at center, and a dot was presented as saccade target at  $\pm 20^\circ$ . By making a saccade, the participants earned the points associated with that stimulus (gain or loss). In decision trials, a single stimulus representing a sure bet and two stimuli representing a risky bet appeared at center. In addition two dots appeared at  $\pm 20^\circ$ . The participants made a choice by making a saccade to one of the dots. **B**: the stimuli consisted of 11 boxes. The colored stimuli were associated with gain or loss (indicated with the plus or minus), each with a distribution as shown. The black stimulus was always associated with zero points. **C**: we used a neural network to model the decision-making process. The input  $\mathbf{x}$  was an 11-element vector, with each element representing one of the stimuli  $x_1, \dots, x_{11}$  starting from the most negative to the most positive, and the black box (0 points) being the sixth element. On each trial, the input vector  $\mathbf{x}$  was set so that one element had value of  $-1$  for the sure stimulus, two elements had value of  $+0.5$  for the pair of risky stimuli, and 0 for the remaining elements. The weight vector  $\mathbf{u}$  represented the subjective value of each stimulus. Variable  $z$  was determined by Eq. 1, and the output of the network was the probability (prob.) of picking the sure option. **D**: some of the choices made by a participant and the output of the network. The colored dots indicate the stimuli that were presented for the risky option. For example, the red dots indicate trials in which  $-2$  and  $+5$  were presented as the risky option. The x-axis is the point value of the sure option. The y-axis is the probability of picking the sure option by this participant. The line traces connect the output of the network for each decision. For example, the green dots show the probability that the participant picked the sure option when the risky option was  $+5$  and  $-4$  stimuli. The expected value of the risky option was 0.5. This participant tended to pick the sure option if that option had a value greater than 0.5. **E**: there was diversity in the subjective valuation that the participants learned to assign to the stimuli. **Left**: subjective values that two participants learned. **Right**: the distribution of the slope of the subjective versus objective values across the participants. **F**: subjective values across all participants. Dashed line is identity. Error bars are SE. a.u., Arbitrary units; S02, S03, S05, individual subjects.

and a single dot appeared on the horizontal axis (at either  $+20^\circ$  or  $-20^\circ$ ). This was the instruction for the participant to make a saccade to the dot. Once the saccade concluded, the stimulus at center was erased and displayed at the dot location, along with text that indicated the number of points that the participant had gained or lost for the trial. As in the decision trials, the points were drawn from the random distribution associated with the colored stimulus. The subject had 5 s to complete the saccade. If no saccade was made, the subject was penalized with  $-10$  points.

We were concerned that the asymmetry in velocity of temporal and nasal saccades could affect our ability to measure the relationship between saccade velocity and subjective value. Therefore, we designed our experiment so that, in probe trials, for each colored stimulus (displayed at center) the peripheral dot was placed an equal number of times at  $+20^\circ$  and  $-20^\circ$ . This ensured that any differences in saccade velocity between stimuli were not due to differences in

direction of saccades in probe trials, but rather differences in valuation of the stimulus. In addition, the 10 colored boxes were presented with equal frequency within each block, distributed randomly in the probe trials.

In summary, probe trials included boxes at the central location that were associated with gain or loss. By making a saccade to a peripheral dot, the participant earned that gain or loss. Failure to make a saccade resulted in loss of 10 points.

**Experiment design.** Before the start of the experiment, the participants were instructed that there were 10 stimuli consisting of two sets of 5 colored boxes that represented points that could be gained or lost on each trial. "Each color will indicate how many points you will gain or lose. Black box will always give zero points when chosen. Boxes with plus signs will add to your score, while boxes with minus signs will decrease your score. For example, if orange box with plus sign indicates gain of 10, orange box with minus sign will indicate loss of 10."



The experiment consisted of 11 blocks, each with 100 trials. The first block was a training block and began with 100 points and included only probe trials. This first block served to teach the participants the points associated with the various stimuli. The remaining 10 blocks each had 40 probe trials and 60 decision trials, distributed randomly. The total score was reset to 100 at the start of the second block. Following completion of the second block, the final score of the previous block was carried over as the starting score of the next block. At the conclusion of every fourth trial, the total score earned was displayed at center fixation.

**Data analysis.** Eye position data were filtered with a second-order Savitzky-Golay filter (frame size 11, degree 3). Saccade onset and offset were determined in real time with 20°/s threshold. We identified valid saccades as those that occurred between stimuli with start and end points that were within 5° of the boundaries of the start and end images (to account for the fact that participants were not specifically instructed to fixate on a precise location). For probe trials, we excluded reaction times that were larger than 1 s.

Our goal was to test whether behavior in probe trials reflected the subjective value that we had estimated from decision trials. Thus, we focused on saccade kinematics in probe trials and inferred subjective value based on choices made in decision trials. Statistical testing relied on linear mixed-effect models. In each model, the dependent variables were saccade peak velocity, reaction time, and amplitude. Fixed effects were stimulus objective value and subjective value, and random effects were individuals. Dependent variables were normalized for each individual by dividing the measured value by the within-subject mean. Statistics were performed on normalized dependent variables.

**Estimating subjective value of stimuli.** The objective value of each stimulus was set by the mean of the point distribution associated with each colored stimulus (Fig. 1B). The participants formed subjective values, and we inferred these values based on the choices that they made in decision trials.

In a decision trial, the choice was between a sure option (a single stimulus) and a risky option (two stimuli, 50% chance of each). To model the choices that participants made, we designed a one-layer perceptron network that had as its input the three stimuli that were available on each trial. The input to the network was an 11 element vector  $\mathbf{x}$ , with each element representing one of the stimuli  $x_1, \dots, x_{11}$  starting from the most negative to the most positive, and the black box (0 points) being the sixth element. The output of the network was the probability that the participant would pick the sure option (Fig. 1C). To train the network, on each decision trial the input vector  $\mathbf{x}$  was set so that one element had value of  $-1$  for the sure stimulus, two elements had value of  $0.5$  for the pair of risky stimuli, and  $0$  for the remaining elements. The weight vector  $\mathbf{u}$  represented the subjective value of each stimulus and was also an 11-element vector. A linear combination of the available stimuli were represented with variable  $z$ :

$$z = \mathbf{u}^T \mathbf{x} \quad (1)$$

For example, if in a given trial the sure option was stimulus  $x_4$ , and the risky option was stimuli  $x_2$  and  $x_7$ , then  $z = 0.5(x_2 + x_7) - x_4$ . In other words, the variable  $z$  represented the difference between the subjective values of the two options. This was then transformed via a logistic function that produced an output  $y$  that represented the probability of picking the sure option:

$$y = \frac{1}{1 + \exp(-z)} \quad (2)$$

Our goal was to estimate the subjective value that the participant had assigned to each stimulus, represented via the weight vector  $\mathbf{u}$ . We assumed that the subjective value of the zero stimulus (the sixth element of  $\mathbf{u}$ ) was exactly zero. To find the remaining weights, we used a binary cross-entropy loss function:

$$J = -\frac{1}{N} \sum_{n=1}^N (t^{(n)} \log y^{(n)} + (1 - t^{(n)}) \log(1 - y^{(n)})) \quad (3)$$

In the above equation,  $N$  is the total number of decision trials (600). Binary variable  $t^{(n)}$  represented the actual decision of the participant on trial  $n$ :  $t^{(n)} = 1$  for choosing the risky option, and  $t^{(n)} = 0$  for the sure option. To find the value of  $\mathbf{u}$  that minimized Eq. 3, we used stochastic gradient descent and differentiated Eq. 3 with respect to  $\mathbf{u}$ . We have:

$$\begin{aligned} \frac{dJ}{dy^{(n)}} &= \frac{t^{(n)} - y^{(n)}}{(1 - y^{(n)})y^{(n)}} \\ \frac{dy^{(n)}}{dz} &= \frac{\exp(-z)}{(1 + \exp(-z))^2} \end{aligned} \quad (4)$$

In addition, we have  $dz/d\mathbf{u} = \mathbf{x}$ . On each trial, we updated estimate of  $\mathbf{u}$  as follows:

$$\mathbf{u}^{(n+1)} = \mathbf{u}^{(n)} - \alpha \frac{dJ}{dy^{(n)}} \frac{dy^{(n)}}{dz} \mathbf{x}^{(n)} \quad (5)$$

We stopped the algorithm when the norm of change of the subjective value  $\Delta \mathbf{u}$  was less than  $10^{-4}$ , which usually happened after 1,000 iterations. We used the full data set as one batch for each iteration. The learning rate was set to 10 (as loss was divided by the number of inputs).

**Normalizing saccade velocity for changes in amplitude.** In probe trials, participants were presented with a colored stimulus at center and a dot in the periphery. They made a saccade to the dot. We found that the reaction time, peak velocity, and amplitude of this saccade varied with the subjective value that the participant had assigned to the colored stimulus: velocity and amplitude both increased with stimulus value. Because velocity naturally increases with saccade amplitude, we wondered whether all the velocity changes were driven by amplitude, or whether velocity was specifically driven by subjective value over and beyond what would be expected because of amplitude changes.

To normalize for amplitude-dependent changes in velocity, we used a procedure described earlier (Reppert et al. 2015). For each participant, and in all trials (probe and decision trials combined), we measured the amplitude of each saccade along with its peak velocity. Previous work had shown that a hyperbolic function is generally a good fit to human saccade data (Choi et al. 2014). The resulting hyperbolic function produced participant-specific expected saccade velocity as a function of amplitude. In probe trials, for each saccade we computed the ratio between the measured peak velocity and the expected velocity. This ratio produced a within-subject measure of amplitude-normalized peak velocity. We then looked to see whether this amplitude-normalized velocity varied with stimulus subjective value.

**Using vigor in probe trials to predict choice in decision trials.** Once we determined reaction time, peak velocity, and saccade amplitude associated with a given stimulus in the probe trials, we asked whether these variables could serve as a proxy for subjective value. To evaluate the accuracy of such a policy we used two different approaches: a winner-take-all approach that predicted choice in the decision trials, and a likelihood-estimate approach that predicted probability of choice in decision trials.

In the winner-take-all approach, for each stimulus we computed the mean saccade velocity, reaction time, and amplitude from the probe trials. We defined subjective value of each stimulus by these kinematic variables. Then on each decision trial we used these kinematic measures to predict choice. For example, for a given participant, we assigned subjective value to the 11 stimuli based on reaction time on their probe trials. That is, we set vector  $\mathbf{u}$  to be equal to the mean probe trial reaction time for the various stimuli. We then used these values to predict choice of that participant in each decision trial: pick

the option that has the highest subjective value. This was computed as 100% of the reaction time for the sure option stimulus, versus sum of 50% of reaction time for each of the risky option stimuli. Thus, this model had no free parameters. Rather, it simply picked the option that was associated with the lowest average reaction time (as measured in probe trials).

We compared the accuracy of this kinematic-based policy with a policy that made choices based on subjective values that were estimated based on the actual decisions of each participant (the 10-parameter model of Fig. 1C). To predict outcomes, we used the actual options faced by each participant. We used Wilcoxon signed-rank test to compare performance of the various policies.

In the likelihood-estimate approach, we began by setting the vector  $\mathbf{u}$  to be equal to the mean reaction time (or peak velocity or amplitude) for the various stimuli in the probe trials. We then used this kinematic based estimate of subjective value to predict the probability that the participant would pick the sure option in a given decision trial:

$$y = \frac{1}{1 + \exp(-\beta z)} \quad (6)$$

In the above formulation, the term  $\beta$  appears because unlike Eq. 2,  $z$  in Eq. 4 has units that are different than subjective value. We used Eq. 3 to guide the gradient descent procedure for finding  $\beta$  for each participant. Thus, we used Eq. 4 to predict the probability that on a given trial the participant would pick the sure option. We then evaluated the goodness of this policy by computing the log-likelihood. We compared the kinematic based policy to a random policy. Log-likelihood for the random policy was one in which the probability of choosing the sure option was 0.5 (equivalent to having elements of  $\mathbf{u}$  equal to each other).

The neural network had 10 free parameters that were fit to the actual choices that the participants made, resulting in a logistic model that provided an upper bound on how well one could predict choices of each participant. In contrast, the kinematic-based approach made predictions regarding choice patterns using a model that had zero free parameters. To provide a more fair comparison of the two models, for the neural network, for each participant we fit the model using 90% of the decision trials and then tested its accuracy for the remaining 10%. We repeated this 10 times for each participant and report this cross-validation decision accuracy.

## RESULTS

In a baseline set of probe trials, the participants learned to associate abstract stimuli (colored boxes) with a gain or loss (Fig. 1B). In these trials, a single stimulus appeared at center, and a dot appeared to one side. Once the saccade moved the eyes to the dot, the value of the stimulus was revealed. In a subsequent set of decision trials, participants were presented with a sure option and a risky option (Fig. 1A, right column). They expressed their choice by making a saccade to one side or another. We asked whether we could use saccade kinematics from the probe trials to estimate the subjective value that the participant had assigned to each stimulus and then predict their choices in decision trials.

Some of the choices made by a representative participant are shown in Fig. 1D. The  $x$ -axis of this plot presents the objective value of the sure option on various trials (mean of the point distribution associated with that stimulus, Fig. 1B). The  $y$ -axis of Fig. 1D presents the probability that the participant selected the sure option. The various colored dots indicate the stimuli that were presented for the risky option. For example, the green dots show the probability that the participant picked the sure option when the risky option consisted of the colors associated

with +5 and -4 points. The expected value of the risky option was 0.5. Indeed, this participant tended to pick the sure option if that option had a value greater than 0.5. In another example, the red dots show the probability that the participant picked the sure option when the risky option was -2 and +5 points. In this case, the risky option had an expected value of 1.5. The participant now picked the sure option when that option had a value that was greater than 1.5.

We used a one-layer neural network to model the choices that each participant made in the decision trials and infer the subjective value that they had assigned to each stimulus (Eq. 2). On each trial, given the sure and two risky stimuli, the network predicted the probability that the participant would pick the sure option (lines in Fig. 1D). We trained the network with the actual choices that the participant had made. The loss function (Eq. 3) was guaranteed to minimize the difference between choices observed and choices predicted. After fitting the network to the data of each participant, the weight vector  $\mathbf{u}$  provided the estimate of the participant's subjective value for each stimulus.

Examples of subjective values inferred from choices made by two participants are illustrated in Fig. 1E, left. Participant S03 learned a shallow function, distinguishing between positive and negative valued stimuli, but not distinguishing well within stimuli that were positive or negative. In contrast, participant S05 learned a steep function that distinguished well within negative and positive stimuli. Thus, among the participants there was diversity in the value that they assigned each stimulus, as shown by the distribution of subjective versus objective value slope ( $0.88 \pm 0.40$ , mean  $\pm$  SD, Fig. 1E, right). As we will see, this diversity played an important role in the question of whether vigor was driven by subjective value.

The average pattern of subjective value is presented across the participants in Fig. 1F. We found that subjective value strongly correlated with objective value of the stimuli ( $r^2 = 0.72$ ;  $P < 10^{-30}$ ). Within participant analysis of subjective value revealed a main effect of objective value [ $F(1,262) = 585.7$ ,  $P < 10^{-30}$ ], demonstrating that as the objective value of the stimuli increased, so did the subjective value that the participants had assigned to them.

Overall, our method for estimating subjective values correctly predicted choices that the participants made in  $80.4 \pm 1.5\%$  of the trials. In comparison, when we assumed that the objective values were known a priori, such a model correctly predicted choices in  $81.1 \pm 1.4\%$  of the trials. Thus, the participants learned the task. However, there were also differences among participants, with some learning steep value functions, while others learning shallow functions.

*Velocity increased and reaction time decreased with stimulus value.* In probe trials a single colored stimulus appeared at center, indicating the value of that trial, and simultaneously a dot appeared on the periphery, indicating the saccade target. Presentation of the colored stimulus served as the go cue. The dependent variables were reaction time, velocity, and amplitude of the ensuing saccade. Participants had 5 s to make a saccade to the peripheral dot and by doing so earned the loss or gain that was associated with the stimulus. Notably, the loss that was indicated by the negative valued stimuli was always less than the large penalty (10 points) that would be applied if the participants did not make the correct saccade.

Figure 2A illustrates saccade amplitude, velocity, and reaction times for one participant in probe trials for +5 and −5 stimuli. In response to the higher valued stimulus, this participant produced a saccade that had a shorter reaction time, and a higher peak velocity, but with little change in saccade amplitude.

To examine these trends across the participants, we normalized peak speed, reaction time, and amplitude for each individual with respect to their own mean as measured across probe trials (Haith et al. 2012). Data from two participants are presented in Fig. 2B. Participant S05 exhibited peak velocity and reaction time that correlated strongly with stimulus objective value (velocity:  $r = +0.93$ ,  $P < 10^{-4}$ ; reaction time:  $r = -0.96$ ,  $P < 10^{-5}$ ). In contrast, in participant S03 velocity and reaction time were poorly correlated with objective value (velocity:  $r = +0.39$ ,  $P = 0.26$ ; reaction time:  $r = -0.48$ ,  $P = 0.16$ ). In contrast, saccade amplitude did not vary significantly with objective value in these two participants (S03,  $r = 0.55$ ,  $P = 0.10$ ; S05,  $r = 0.62$ ,  $P = 0.055$ ).

The distribution of correlation coefficients for all participants is plotted in Fig. 2C. The correlation coefficients between reaction time and subjective value had a mean of  $-0.38 \pm 0.076$  (Wilcoxon signed-rank test,  $P = 2.1 \times 10^{-4}$ ). With respect to objective value, this distribution had a mean of  $-0.42 \pm 0.072$  (Wilcoxon signed-rank test,  $P = 3.2 \times 10^{-4}$ ). The distribution of correlation coefficients between peak velocity and subjective value had a mean of  $0.29 \pm 0.082$  (Wilcoxon signed-rank test,  $P = 0.0051$ ). With respect to objective

value, this distribution had a mean of  $0.28 \pm 0.085$  (Wilcoxon signed-rank test,  $P = 0.0032$ ). The distribution of amplitude and subjective value had a mean of  $0.25 \pm 0.094$  (Wilcoxon signed-rank test,  $P = 0.0192$ ). With respect to objective value, this distribution had a mean of  $0.23 \pm 0.09$  (Wilcoxon signed-rank test,  $P = 0.0225$ ). Thus, these data suggested that, across participants, kinematics in probe trials might serve as a sensitive proxy for subjective value.

To examine the results together, we binned the vigor data based on the objective value of the stimulus (10 bins, one per stimulus, Fig. 3A) and found that saccade velocity and amplitude increased with objective value of the stimulus [within-subject effect, velocity  $F(1,228) = 28.6$ ,  $P = 2.2 \times 10^{-7}$ , amplitude  $F(1,228) = 17.97$ ,  $P = 3.3 \times 10^{-5}$ ]. Similarly, reaction time decreased with objective value of the stimulus [within-subject effect,  $F(1,228) = 50.6$ ,  $P = 1.4 \times 10^{-11}$ ]. Thus, although there was diversity among the participants, saccade kinematics in probe trials was significantly affected by objective value of the stimulus.

The reduction in reaction time despite increased saccade amplitude is noteworthy because in horizontal saccades of  $10^\circ$  or larger, reaction time tends to increase with increased amplitude (Reppert et al. 2018). Here, as objective value increased, amplitude increased by  $1.1 \pm 0.35\%$ , peak velocity increased by  $3.1 \pm 0.97\%$ , and reaction time declined by  $7.35 \pm 2.1\%$ .

We next tested the effects of subjective value on saccade parameters. We observed that as subjective value increased,

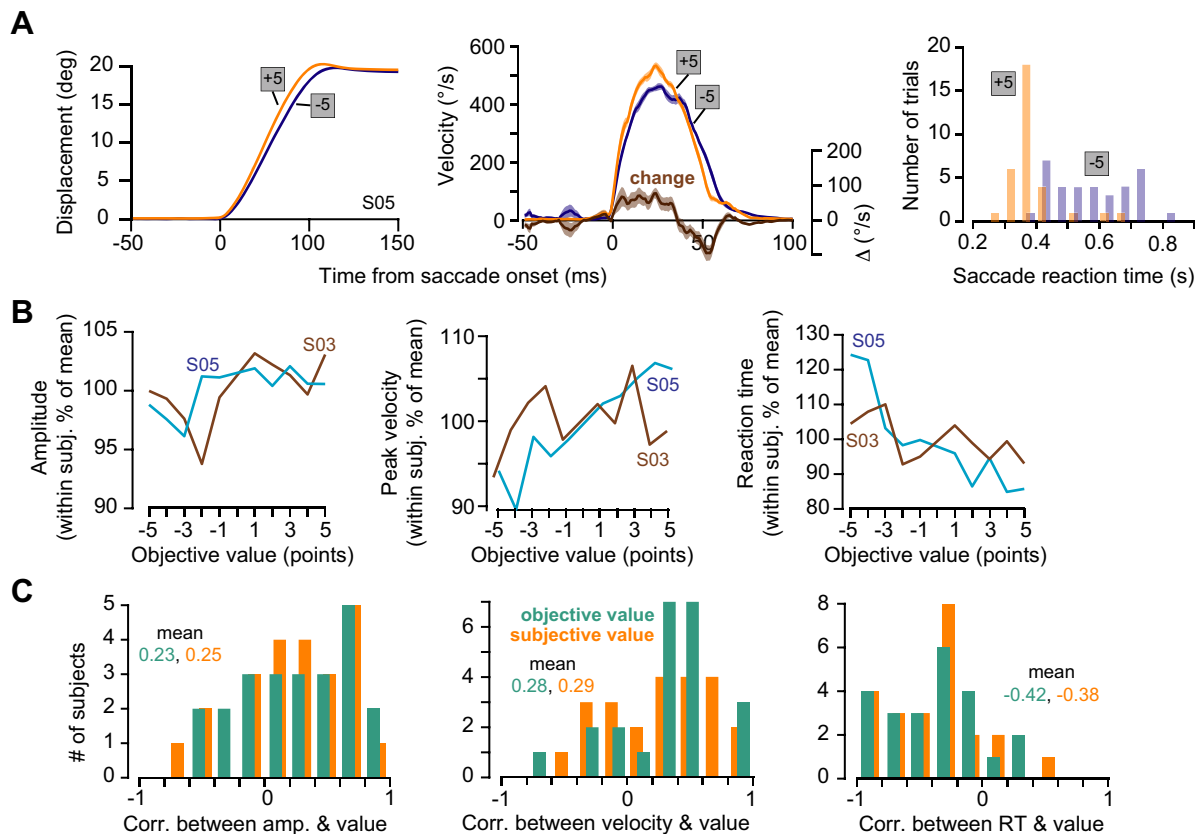


Fig. 2. Saccade characteristics in probe trials. A: data from a single participant. Plots show saccade displacement, velocity, and reaction time when the stimulus was +5, or −5. Error bars are within-subject SE. B: amplitude, peak velocity, and reaction time in probe trials as a function of objective value of the stimulus for two participants (same individuals, S03 and S05, as those shown in Fig. 1E). C: the distribution of correlation coefficients between kinematic variables and stimulus value. Mean value refers to the mean of subjective and objective values. amp., Amplitude; Corr., correlation; RT, reaction time; subj., subject.

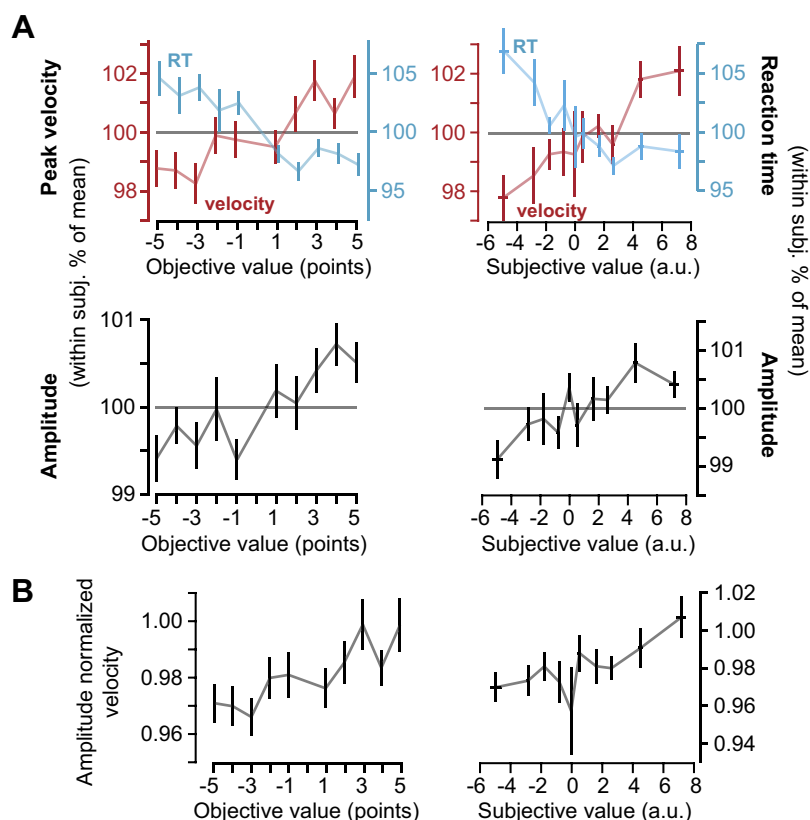


Fig. 3. Relationship between saccades and stimulus value in probe trials. **A:** peak velocity, reaction time, and amplitude of saccades as a function of objective and subjective values. Each variable is normalized for each participant with respect to their own mean in probe trials. **B:** amplitude-normalized change in peak velocity as a function of objective and subjective value of the stimulus in probe trials. Normalization built a within-subject model of amplitude-velocity relationship based on all saccades, including those in decision trials. In probe trials, velocities were slightly lower than in decision trials, thus the fact that the values here are generally smaller than 1. Error bars are between-subject SE. a.u., Arbitrary units; RT, reaction time; subj., subject.

saccade velocity and amplitude increased [within-subject effect, velocity  $F(1,228) = 33.6$ ,  $P = 2.3 \times 10^{-8}$ , amplitude  $F(1,228) = 23.2$ ,  $P = 2.7 \times 10^{-6}$ ], and reaction time decreased [within-subject effect,  $F(1,228) = 62.2$ ,  $P = 1.3 \times 10^{-13}$ ], as shown in the right part of Fig. 3A. To make this plot, we began with the distribution of subjective values across all participants, and then sampled that distribution into 10 bins of equal probability. Thus, the bins have error bars in both  $x$ - and  $y$ -dimensions. Together, these data demonstrated that vigor in probe trials was not a U-shaped function of stimulus value. Rather, vigor tended to be smallest for stimuli that were associated with loss, and largest for stimuli that were associated with gain.

Because peak velocity tends to increase with saccade amplitude, we wondered whether the velocity increase associated with stimulus value was driven solely by amplitude changes, or whether velocity increased over and beyond the expected change with amplitude. For each participant, using data from saccades in both probe and decision trials, we built a mathematical model of the participant's amplitude-velocity relationship (Reppert et al. 2015). This model produced the expected saccade velocity as a function of amplitude. For a given stimulus in probe trials, we measured the resulting saccade peak velocity and amplitude and represented velocity as a ratio with respect to the expected velocity for that amplitude. This produced an amplitude-normalized measure of velocity. We then looked to see whether this amplitude-normalized measure changed with stimulus value (Fig. 3B). In general, probe trial saccades had velocities that were somewhat smaller than decision trial saccades (thus, the values in Fig. 3B are not centered at 1). However, amplitude-normalized velocity increased with both objective and subjective value [objective

value:  $F(1,228) = 18.7$ ,  $P = 2.31 \times 10^{-5}$ , subjective value:  $F(1,228) = 23.7$ ,  $P = 2.15 \times 10^{-6}$ ].

Given the between-subject diversity in the relationship between vigor and stimulus value in probe trials (Fig. 2C), we wondered whether there was some characteristic of participants in decision trials that dissociated their kinematic modulation in probe trials. One clue was that some participants learned a steep value function, while others learned a shallow function (Fig. 1E). Indeed, we found that the slope of subjective to objective values was modestly correlated with the slope of saccade velocity with respect to subjective values (slope of velocity versus subjective value compared with slope of subjective value versus objective value,  $r = 0.49$ ,  $P = 0.019$ ). That is, the participants whose saccade velocity was more strongly modulated by stimulus value in probe trials tended to have learned a steeper value function, as inferred from their choices in decision trials.

In summary, we observed that in probe trials saccades had reaction times that decreased with subjective value, and peak velocities and amplitudes that increased with subjective value. However, there was diversity in the strengths of these relationships. It appeared that saccade kinematics was more strongly modulated by stimulus value in those participants who had also learned a steeper value function.

*Between-subject differences in subjective value influence between-subject differences in modulation of saccade kinematics.* Some participants learned to assign a large subjective value to a stimulus, while others assigned a lower value to the same stimulus. Could this between-subject difference in valuation be gleaned from the kinematic patterns?

To examine this question, we described our hypothesis via a graphical model (Fig. 4A). In this model, choice depended on



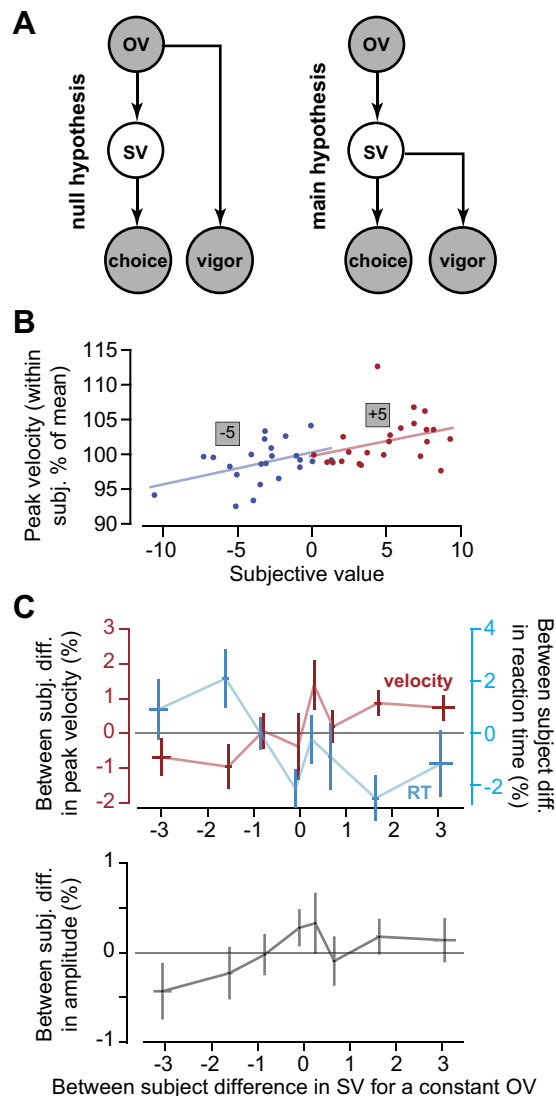


Fig. 4. Statistical relationship between saccade kinematics, objective value (OV), and subjective value (SV). **A**: graphical model representing two hypotheses. Each circle is a random variable. OV is objective value, and SV is subjective value. The filled circles are measured variables. The unfilled circle is not measured but estimated. In the null hypothesis, choice depends on SV, and both SV and kinematics depend on OV. In the main hypothesis, SV affects both choice and kinematics. **B**: to evaluate merits of the hypotheses, we kept OV constant and measured variability in velocity and reaction time as a function of variability in SV. In this plot blue dots are peak saccade velocity as a function of subjective value (each dot is a participant), for the fixed objective value of  $-5$ . The red dots are for the fixed objective value of  $+5$ . For a given stimulus, some participants assigned a high SV, while others assigned a low SV. Velocity appeared to be higher and reaction time lower when the participant assigned a high SV to the stimulus. There was no significant effect on amplitude. **C**: to consider the data in part B across stimuli, we found the mean of the dot distribution for each stimulus (for example, the mean for the stimulus  $OV = +5$ ), and then represented each dot with respect to the within stimulus mean. The result revealed that, for a constant OV, a change in SV produced a change in reaction time (RT) and velocity, thus rejecting the null hypothesis. diff., Difference; subj., subject.

subjective value, which in turn depended (through learning) on the objective value of the stimulus. In our null hypothesis (H0, Fig. 4A), the objective value affected kinematics, whereas subjective value affected choice. In our main hypothesis (H1, Fig. 4A), objective value affected subjective value, which in turn affected both choice and kinematics. Under H1, if a

participant had learned to associate a small subjective value with a stimulus, then their vigor would be low in response to that stimulus. However, if that same stimulus was valued highly by another participant, then their vigor would be high. Thus, to test this hypothesis, we kept objective value constant and asked whether changes in subjective value across participants modulated saccade kinematics.

To help explain how we tested this hypothesis, Fig. 4B illustrates peak velocity in probe trials as a function of subjective value for two different stimuli. For the  $+5$  stimulus, some participants assigned a large value, while others assigned a small value. Similarly, for the  $-5$  stimulus, there was diversity in assignment of subjective values. However, individuals that assigned larger subjective value to a given stimulus also appeared to move with greater velocity in response to that stimulus (similar positive slopes of the red and blue lines in Fig. 4B).

To test for the consistency of this relationship, for each stimulus (constant objective value) we measured the kinematic variable for a participant (i.e., the y-value of a point in Fig. 4B with respect to the mean of the points with the same color), and the subjective value that they had assigned (i.e., the x-value of a point in Fig. 4B with respect to the mean of the points with the same color). Thus, given a constant objective value, we measured how the between-subject differences in subjective value affected between-subject differences in kinematics (Fig. 4C). We found that given a constant objective value, an increase in subjective value produced a reduction in reaction time  $[F(1,228) = 8.7, P = 0.0036]$ , and an increase in peak velocity  $[F(1,228) = 8.1, P = 0.0047]$ . However, there was no effect of subjective value on saccade amplitude  $[F(1,228) = 2.62, P = 0.11]$ .

These results suggest that between-subject differences in valuation of a stimulus can be partially inferred from the between-subject differences in saccade kinematics: participants who learned to associate a greater value to a given stimulus also tended to exhibit a greater modulation of reaction time and velocity (but not amplitude) in response to that stimulus.

**Predicting preference from kinematics only.** We next asked how well kinematic measurements in probe trials could predict choices that individuals made in decision trials. To predict choice, we used only the kinematic data in probe trials. Thus, our approach relied on an a priori model that had no free parameters.

The kinematic measurements produced two policies: a policy that assigned subjective value based on reaction time in probe trials, and another policy that assigned subjective value based on peak velocity in the same trials (combining the two policies only marginally improved performance). For example, given the probe trial data for a participant, the reaction time policy assigned a subjective value to the various stimuli, which we then used to predict choice in decision trials for that participant. This served as the winner-take-all approach. In addition, we considered a likelihood approach in which we predicted the probability that the participant would pick the sure option based on their kinematic patterns in the probe trials (this model had one free parameter).

For the winner-take-all approach (nothing to fit), we divided the decision trials into easy and hard based on the difference in the objective value of the sure and risky options: easy trials were denoted by objective value difference of 1 point or more,



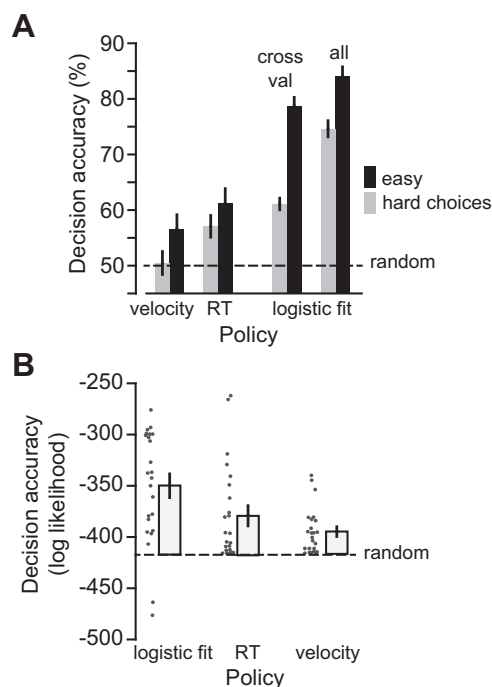


Fig. 5. The ability to predict choice in decision trials from kinematic variables in probe trials. **A**: hard choices are those in which the difference between the objective values of the two options was 1 point or less. Easy choices are those in which this difference was greater than 1 point. The velocity and reaction time (RT) policies computed subjective values based on probe trials and then predicted choices in decision trials. Logistic fit policy fit the all data in decision trials and predicted choice in the same data set, thus representing the ceiling for a model based prediction. Logistic fit cross validation (cross val) fit 90% of the data in decision trials for each participant and predicted choice in the remaining 10%. **B**: decision accuracy as measured via a log-likelihood estimate. From the kinematic data in probe trials we predicted the probability that the participant would pick the sure option in a given decision trial (Eq. 4) and then evaluated the goodness of the policy by computing the log-likelihood. The dashed line indicates performance of a random policy.

and hard trials were denoted by objective value difference of less than 1 point. We quantified accuracy of the kinematic policies based on the number of correct predictions that the reaction time and the velocity policies made.

To define an upper bound on prediction accuracy, we also quantified performance of a policy that relied on the neural network that was fit to the actual choices (termed logistic fit). In comparison to our kinematic based model (0 or 1 free parameter), the neural network had 10 free parameters. We found that when the neural network was trained on all the actual choices, the upper bound on decision accuracy was around 80% (bars labeled “all”; Fig. 5A). When we trained the network on 90% of each participant’s choice data and then predicted choices on the remaining 10%, decision accuracy of the network dropped to around 70% (bars labeled cross validation, “cross val”; Fig. 5A).

The results of the kinematic based policy are shown in Fig. 5A. We found that for hard choices, a velocity-based policy performed no better than chance (Fig. 5A, Wilcoxon signed-rank test,  $P = 0.99$ ). However, for the same hard choices a reaction time policy performed significantly better than chance (Wilcoxon signed-rank test,  $P = 0.0042$ ). For the easier choices, both the velocity based policy and the reaction time policies performed significantly better than chance (Fig. 5A, left subplot, velocity Wilcoxon signed-rank test  $P = 0.0225$ ;

reaction time Wilcoxon signed-rank test  $P = 5.5 \times 10^{-4}$ ). Reaction time policy produced a policy that had a decision accuracy of roughly 60%.

In addition to predicting choice via winner take all, we also used kinematic based estimates of subjective value to compute the probability of choosing the sure option (Eq. 4). We estimated the goodness of the vigor based predictions via log likelihood and compared it to the likelihood from a random policy (Fig. 5B). The value of the velocity and reaction time policies were better than choosing randomly (Wilcoxon signed-rank test, velocity:  $P = 1.22 \times 10^{-4}$ ; reaction time:  $P = 2.93 \times 10^{-4}$ ).

Overall, using kinematics in probe trials as a proxy for subjective value was informative, as the results were better than chance. Reaction time was a better predictor than velocity, allowing one to predict with roughly 60% accuracy the choices made by the participants. This compares with the cross-validation performance of roughly 70% accuracy when subjective value was estimated from 90% of the actual choices and then tested on the remaining 10%.

**Kinematic patterns in decision trials.** In decision trials the participants expressed their choices with a saccade. We asked whether vigor patterns in the decision trials carried information about the contents of the trial.

The top plot of Fig. 6 displays time to decision (deliberation time) as well as saccade velocity as a function of trial difficulty. To quantify trial difficulty, we measured the subjective value of the chosen option minus the subjective value of the




Fig. 6. Behavior during decision trials. Decision time (DT) refers to time from trial onset to the saccade onset that indicated choice. Peak velocity refers to the velocity of the saccade that indicated choice. **Top**: trial difficulty was measured for each participant via the difference between-subjective value of the chosen option minus the value of the other option. Hard choices are those in which this difference is less than or equal to 1. Easy choices are those in which this difference is greater than 1. **Bottom**: trial value was measured for each participant via the sum of the subjective values of the two options. When this sum was large and positive, both options were good, predicting gain regardless of the chosen option. When this sum was large and negative, both options were bad, thus predicting loss. Saccade velocity that indicated choice increased with trial value. subj., Subject; SV, subjective value.

alternative option. For example, if the participant chose the sure option, trial difficulty was the subjective value of the chosen stimulus minus 0.5 times the sum of subjective values of the two stimuli in the risky option. As this difference became more positive, the decision became easier. Indeed, easier decisions coincided with reduced deliberation time [labeled DT in Fig. 6, *top*,  $F(1,149) = 68.3$ ,  $P = 7.2 \times 10^{-14}$ ]. Saccade velocity that reported the choice tended to be low in the most difficult trials (yellow region, Fig. 6, *top*), possibly indicating that reward was uncertain. This is consistent with earlier work that reported low saccade velocity in trials with increased reward uncertainty (Seideman et al. 2018).

The random nature of the stimuli resulted in some trials that predicted a large loss regardless of choice and other trials that predicted a large gain. For example, when the sure stimulus was associated with a gain, the risky option often also included stimuli that summed to a gain. Thus in this case both options were associated with gain, producing a large trial value. When we considered trial value as the sum of the subjective values of both options, a pattern emerged: when the trial predicted a loss (because both options were bad), reaction times tended to be long and saccade velocities tended to be low (Fig. 6, *bottom*). As trial value increased, saccade velocity increased [ $F(1,275) = 20.3$ ,  $P = 9.8 \times 10^{-6}$ ] while decision time decreased [ $F(1,275) = 133$ ,  $P = 2.1 \times 10^{-25}$ ]. The dependence of saccade velocity on trial value was present even after we normalized trials based on their difficulty [effect of sum of subjective values on decision time in easy trials:  $F(1,457) = 130.5$ ,  $P = 9.3 \times 10^{-27}$ ; in hard trials:  $F(1,457) = 7.56$ ;  $P = 0.0062$ ].

In summary, when the two options were both bad, forecasting a loss, velocity of the saccade that reported the choice was low. As the value of the trial improved, forecasting a gain, saccade velocity increased. Thus, in both probe and decision trials, saccade velocity tended to be low for trials that predicted loss and high for stimuli that predicted gain.

## DISCUSSION

The brain makes decisions based on subjective valuation of the available options. Yet, how we value an option is a hidden variable that cannot be measured directly. Rather, it must be inferred from our decisions. Is there a component of behavior other than choice that can serve as a proxy for subjective valuation?

Here, we presented abstract visual stimuli that participants learned to associate with gains or losses. We inferred the value that each participant assigned to each stimulus from their choices in decision trials. As expected, some participants learned a steep value function that strongly differentiated the various stimuli, whereas others learned a shallow function. In probe trials, we presented the participants a single stimulus at a central location, indicating the value of that trial, and asked them to make a saccade to a peripheral target. The reaction time, peak velocity, and amplitude of that saccade carried information about the subjective value of the stimulus. For example, reaction time was highest for stimuli that forecasted loss, lowest for stimuli that predicted gain. Even after normalizing for amplitude, velocity varied with stimulus value. Saccade kinematics were more strongly modulated in those participants who had also learned a steeper value function.

As expected, some participants valued a given stimulus more than other participants. A critical question was whether between-subject differences in valuation could be gleaned from the between-subject differences in their patterns of saccade kinematics. We found that for a given stimulus (thus a constant objective value) there was a relationship between subjective value and kinematics: individuals that assigned larger subjective value to a stimulus also tended to move with greater velocity and smaller reaction time in response to that stimulus.

We asked whether kinematic measurements in probe trials could act as a proxy for subjective value and thus predict choices that individuals made in decision trials. We found that reaction time was a better estimator of subjective value than peak velocity, allowing one to predict with roughly 60% accuracy the decisions that were made by participants. This compares to 70% cross-validation accuracy, as described by a 10-parameter model that was fitted to the actual choices. Thus, kinematics in probe trials was modulated by subjective value. Kinematics produced an a priori model (no free parameters) that could provide moderately accurate predictions about individual preferences.

*Estimating subjective value.* Estimating subjective value usually relies on a concept called certainty equivalence (CE): if the option is a risky one that has 50% probability of producing one of two results, then the CE will be the mean of the subjective values of the two results. CE could be directly reported by the participants (Grether and Plott 1979), or via fitting of a logistic function between the a fixed risky option and the variable sure option, or even using a psychological adaptive method such as PEST (parameter estimation by sequential testing) in which each option depends on the choice made in the previous option (Bostic et al. 1990; Christopoulos et al. 2009; Stauffer et al. 2014).

Here we employed a different approach: we implemented a simple neural network, which we found to be an efficient way to infer subjective values from the patterns of choice. Our specific learning rule relied on a loss function that guaranteed that the result would produce the optimum prediction of choices made by each participant. Our approach had the advantage that it allowed us to use a relatively small number of decision trials (600) in which all stimuli were chosen at random. The method produced reasonable results: subjective valuation correlated strongly with objective value (Fig. 1*F*), producing correct prediction of choice on roughly 81% of the trials.

However, we analyzed the data based on an assumption of stationarity of subjective values. That is, we assumed that subjective value was constant throughout the decision trials. We provided 100 baseline trials that provided information about value of each stimulus to the participants before the decision trials began, but our assumption is clearly a simplification. Unfortunately, it is difficult to analyze the data without the assumption of stationarity because in that case one must assume a learning model, which introduces further unknown parameters that require fitting to behavior. However, if such an approach could be pursued, then one could estimate subjective value as a function of time and look for correlations with kinematics.

It is noteworthy that, on average, the participants placed as much importance on the difference between gain and loss (e.g., subjective difference between 0 and +1) as they did on the

exact quantity of gain or loss (e.g., subjective difference between +1 and +5) (Fig. 1*F*). That is, subjective value rose rapidly, then slowly. We think that this may be due to the fact that the participants learned the value of the stimuli within the positive (and negative) range, whereas they were explicitly instructed, via the positive and negative labels inside each stimulus, of a qualitative difference in the values of the labeled stimuli. It seems likely that with further training the subjective values would acquire a larger range over the domain of objective values, perhaps resulting in greater modulation of vigor.

*Subjective value monotonically varies with saccade kinematics.* The main question that we wished to answer was whether vigor of a movement was a monotonic function of its subjective value across the range that spanned loss to gain. While subjective value may be lower for a loss, the stimulus that predicts a loss may gather equal or greater attention than the stimulus that predicts gain. The neural circuits that influence saccade vigor are affected by both subjective valuation (Platt and Glimcher 1999) and attention (Leathers and Olson 2012), making it unclear whether vigor would be influenced by one or the other.

Reaction time, amplitude, and velocity of a saccade are variables that are controlled by activity of neurons in the superior colliculus (Dorris and Munoz 1995; Dorris et al. 1997; Ratcliff et al. 2003; Smalianchuk et al. 2018; Sparks and Hu 2006). Collicular activity is in turn influenced by the excitatory inputs it receives from the cerebral cortex, and the inhibitory inputs that it receives from the basal ganglia. The cortical inputs include projections from the frontal eye field (FEF) and lateral intraparietal area (LIP), both of which house neurons that tend to respond more strongly to stimuli that predict greater reward (Glaser et al. 2016; Louie and Glimcher 2010; Platt and Glimcher 1999). The basal ganglia projections are from the substantia nigra reticulata (SNr), which houses inhibitory neurons that change their discharge in response to magnitude of reward (Sato and Hikosaka 2002; Yasuda et al. 2012; Yasuda and Hikosaka 2017). Thus, subjective valuation of a rewarding stimulus could, in principle, be reflected in an increase in the excitatory inputs to the colliculus from the cortex and a decrease in the inhibitory inputs from the SNr, resulting in a saccade that has a shorter reaction time and greater velocity.

However, the cortical and basal ganglia inputs to the colliculus are also affected by attentional demands of the stimulus. For example, firing rates of neurons in the regions that project to the colliculus are not monotonically driven by value of the stimulus. For example, LIP neurons that respond with greater activity to more rewarding stimuli also respond more strongly to stimuli that predict a stronger loss (Leathers and Olson 2012). In the basal ganglia, activity of SNr neurons is controlled directly and indirectly by neurons in the striatum, which in turn are modulated by dopamine. Dopamine regulates how the striatal neurons respond to cortical inputs. However, while increased dopamine release before onset of a movement tends to invigorate that movement (Kawagoe et al. 2004), some dopaminergic neurons show increased activity in response to a reward predicting stimulus, while others respond with greater activity to both reward and punishment predicting stimuli (Matsumoto and Hikosaka 2009).

Thus, assuming that saccade kinematics is a reflection of excitatory cortical and inhibitory basal ganglia inputs to the colliculus, the current neurophysiological data do not specify whether vigor should be a monotonic function of stimulus value, growing from loss to gain, or whether vigor should be a U-shaped function, showing increased activity both for large gains and large losses.

Here our results unequivocally demonstrate that saccade amplitude, velocity, and reaction time grow monotonically with subjective value across the range that spans from loss to gain. It is possible that in some of the earlier studies in which cortical and dopaminergic activity increased with punishment, the movement that followed may have been expressed with greater vigor (for example, increased rate of blinking).

Furthermore, we found that it was possible to infer some of the between-subject differences in valuation from the between-subject differences in patterns of kinematics. This is reminiscent of an earlier study that found a monkey that did not show velocity and reaction time sensitivity to reward also lacked dopaminergic sensitivity to stimuli that predicted reward (Kawagoe et al. 2004).

*Limitations.* In our experiment the participants learned the value of the stimuli through observation (probe trials) and choice (decision trials), but we analyzed the data as if the subjective values were constant throughout the decision trials. A better approach would be to have a real-time estimate of subjective values during the task. However, such an approach would require fitting behavior to a learning model, which introduces new parameters in the estimation problem. That approach remains to be developed.

In probe trials, the stimulus predicted a loss or gain if the participant performed the correct action (saccade to target). However, if the participant performed an incorrect action (or no action), the consequence was a large loss. Thus, in probe trials the participant could prevent a large loss by performing the correct action but could not prevent the smaller loss associated with the stimulus. In a different design in which the stimulus predicts a loss, but the correct action can prevent it, vigor of that action will likely grow with magnitude of loss. That is, if the correct action can aid in prevention of a loss, then we speculate that vigor would no longer exhibit the pattern we found here. This conjecture remains to be tested.

To test whether subjective value affects kinematics, we relied on the fact that among participants, a given stimulus was associated with a range of subjective values. This between-subject analysis revealed that individuals who valued a stimulus more tended to also exhibit greater changes in velocity and reaction time. However, to conclusively infer a causal relationship between subjective value and a kinematic variable we would need to test whether within participant changes in subjective value produce changes in that variable. One way with which subjective valuation may be increased is via expenditure of effort: individuals who expend effort to acquire a particular reward tend to increase the value that they assign that reward. With saccades, effort expenditure can be modulated via eccentricity (Yoon et al. 2018). Future work is needed to explore the within-subject changes in valuation with their vigor.



## GRANTS

The work was supported by grants from the NIH (5-R01-NS078311, 1-R01-NS096083), the Office of Naval Research (N00014-15-1-2312), and the National Science Foundation (CNS-1714623).

## DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

## AUTHOR CONTRIBUTIONS

T.Y. and R.S. conceived and designed research; T.Y. and A.J. performed experiments; T.Y. analyzed data; T.Y. and R.S. interpreted results of experiments; T.Y. and R.S. prepared figures; T.Y., A.A.A., and R.S. edited and revised manuscript; T.Y., A.J., A.A.A., and R.S. approved final version of manuscript; R.S. drafted manuscript.

## REFERENCES

- Bayer HM, Glimcher PW. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47: 129–141, 2005. doi:10.1016/j.neuron.2005.05.020.
- Bostic R, Herrnstein RJ, Luce RD. The effect on the preference-reversal phenomenon of using choice indifference. *J Econ Behav Organ* 13: 193–212, 1990. doi:10.1016/0167-2681(90)90086-S.
- Choi JE, Vaswani PA, Shadmehr R. Vigor of movements and the cost of time in decision making. *J Neurosci* 34: 1212–1223, 2014. doi:10.1523/JNEUROSCI.2798-13.2014.
- Christopoulos GI, Tobler PN, Bossaerts P, Dolan RJ, Schultz W. Neural correlates of value, risk, and risk aversion contributing to decision making under risk. *J Neurosci* 29: 12574–12583, 2009. doi:10.1523/JNEUROSCI.2614-09.2009.
- da Silva JA, Tecuapetla F, Paixão V, Costa RM. Dopamine neuron activity before action initiation gates and invigorates future movements. *Nature* 554: 244–248, 2018. doi:10.1038/nature25457.
- Dorris MC, Munoz DP. A neural correlate for the gap effect on saccadic reaction times in monkey. *J Neurophysiol* 73: 2558–2562, 1995. doi:10.1152/jn.1995.73.6.2558.
- Dorris MC, Paré M, Munoz DP. Neuronal activity in monkey superior colliculus related to the initiation of saccadic eye movements. *J Neurosci* 17: 8566–8579, 1997. doi:10.1523/JNEUROSCI.17-21-08566.1997.
- Foley NC, Jangraw DC, Peck C, Gottlieb J. Novelty enhances visual salience independently of reward in the parietal lobe. *J Neurosci* 34: 7947–7957, 2014. doi:10.1523/JNEUROSCI.4171-13.2014.
- Glaser JI, Wood DK, Lawlor PN, Ramkumar P, Kording KP, Segraves MA. Role of expected reward in frontal eye field during natural scene search. *J Neurophysiol* 116: 645–657, 2016. doi:10.1152/jn.00119.2016.
- Grether DM, Plott CR. Economic theory of choice and the preference reversal phenomenon. *Am Econ Rev* 69: 623–638, 1979.
- Haith AM, Reppert TR, Shadmehr R. Evidence for hyperbolic temporal discounting of reward in control of movements. *J Neurosci* 32: 11727–11736, 2012. doi:10.1523/JNEUROSCI.0424-12.2012.
- Hanes DP, Schall JD. Neural control of voluntary movement initiation. *Science* 274: 427–430, 1996. doi:10.1126/science.274.5286.427.
- Heitz RP, Schall JD. Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76: 616–628, 2012. doi:10.1016/j.neuron.2012.08.030.
- Jevons WS. Brief account of a general mathematical theory of political economy. *J Roy Stat Soc XXIX*: 282–287, 1866.
- Kawagoe R, Takikawa Y, Hikosaka O. Expectation of reward modulates cognitive signals in the basal ganglia. *Nat Neurosci* 1: 411–416, 1998. doi:10.1038/1625.
- Kawagoe R, Takikawa Y, Hikosaka O. Reward-predicting activity of dopamine and caudate neurons—a possible mechanism of motivational control of saccadic eye movement. *J Neurophysiol* 91: 1013–1024, 2004. doi:10.1152/jn.00721.2003.
- Kononov A, Krajibich I. Revealed strength of preference: inference from response times. *Judgm Decis Mak* 14: 381–394, 2019.
- Leathers ML, Olson CR. In monkeys making value-based decisions, LIP neurons encode cue salience and not action value. *Science* 338: 132–135, 2012. doi:10.1126/science.1226405.
- Louie K, Glimcher PW. Separating value from choice: delay discounting activity in the lateral intraparietal area. *J Neurosci* 30: 5498–5507, 2010. doi:10.1523/JNEUROSCI.5742-09.2010.
- Matsumoto M, Hikosaka O. Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature* 459: 837–841, 2009. doi:10.1038/nature08028.
- Menger C. *Principles of Economics*. Vienna, Austria: Braumuller, 1871.
- Milstein DM, Dorris MC. The influence of expected value on saccadic preparation. *J Neurosci* 27: 4810–4818, 2007. doi:10.1523/JNEUROSCI.0577-07.2007.
- Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. *Nature* 400: 233–238, 1999. doi:10.1038/22268.
- Ratcliff R, Chierian A, Segraves M. A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *J Neurophysiol* 90: 1392–1407, 2003. doi:10.1152/jn.01049.2002.
- Reppert TR, Lempert KM, Glimcher PW, Shadmehr R. Modulation of saccade vigor during value-based decision making. *J Neurosci* 35: 15369–15378, 2015. doi:10.1523/JNEUROSCI.2621-15.2015.
- Reppert TR, Rigas I, Herzfeld DJ, Sedaghat-Nejad E, Komogortsev O, Shadmehr R. Movement vigor as a traitlike attribute of individuality. *J Neurophysiol* 120: 741–757, 2018. doi:10.1152/jn.00033.2018.
- Sato M, Hikosaka O. Role of primate substantia nigra pars reticulata in reward-oriented saccadic eye movement. *J Neurosci* 22: 2363–2373, 2002. doi:10.1523/JNEUROSCI.22-06-02363.2002.
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science* 275: 1593–1599, 1997. doi:10.1126/science.275.5306.1593.
- Sedaghat-Nejad E, Herzfeld DJ, Shadmehr R. Reward prediction error modulates saccade vigor. *J Neurosci* 39: 5010–5017, 2019. doi:10.1523/JNEUROSCI.0432-19.2019.
- Seideman JA, Stanford TR, Salinas E. Saccade metrics reflect decision-making dynamics during urgent choices. *Nat Commun* 9: 2907, 2018. doi:10.1038/s41467-018-05319-w.
- Shadmehr R, Orban de Xivry JJ, Xu-Wilson M, Shih TY. Temporal discounting of reward and the cost of time in motor control. *J Neurosci* 30: 10507–10516, 2010. doi:10.1523/JNEUROSCI.1343-10.2010.
- Shadmehr R, Reppert TR, Summerside EM, Yoon T, Ahmed AA. Movement vigor as a reflection of subjective economic utility. *Trends Neurosci* 42: 323–336, 2019. doi:10.1016/j.tins.2019.02.003.
- Smalianchuk I, Jagadisan UK, Gandhi NJ. Instantaneous midbrain control of saccade velocity. *J Neurosci* 38: 10156–10167, 2018. doi:10.1523/JNEUROSCI.0962-18.2018.
- Sparks DL, Hu X. Saccade initiation and the reliability of motor signals involved in the generation of saccadic eye movements. *Novartis Found Symp* 270: 75–88, 2006.
- Spiliopoulos L, Ortmann A. The BCD of response time analysis in experimental economics. *Exp Econ* 21: 383–433, 2018. doi:10.1007/s10683-017-9528-1.
- Stauffer WR, Lak A, Schultz W. Dopamine reward prediction error responses reflect marginal utility. *Curr Biol* 24: 2491–2500, 2014. doi:10.1016/j.cub.2014.08.064.
- Summerside EM, Shadmehr R, Ahmed AA. Vigor of reaching movements: reward discounts the cost of effort. *J Neurophysiol* 119: 2347–2357, 2018. doi:10.1152/jn.00872.2017.
- von Neumann JV, Morgenstern O. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press, 1944.
- Xu-Wilson M, Zee DS, Shadmehr R. The intrinsic value of visual information affects saccade velocities. *Exp Brain Res* 196: 475–481, 2009. doi:10.1007/s00221-009-1879-1.
- Yasuda M, Hikosaka O. To wait or not to wait—separate mechanisms in the oculomotor circuit of basal ganglia. *Front Neuroanat* 11: 35, 2017. doi:10.3389/fnana.2017.00035.
- Yasuda M, Yamamoto S, Hikosaka O. Robust representation of stable object values in the oculomotor Basal ganglia. *J Neurosci* 32: 16917–16932, 2012. doi:10.1523/JNEUROSCI.3438-12.2012.
- Yoon T, Geary RB, Ahmed AA, Shadmehr R. Control of movement vigor and decision making during foraging. *Proc Natl Acad Sci USA* 115: E10476–E10485, 2018. [Erratum in *Proc Natl Acad Sci USA* 115: E11884, 2018]. doi:10.1073/pnas.1812979115.