Behavioral/Systems/Cognitive

Separating Value from Choice: Delay Discounting Activity in the Lateral Intraparietal Area

Kenway Louie and Paul W. Glimcher
Center for Neural Science, New York University, New York, New York 10003

The mathematical formulations used to study the neurophysiological signals governing choice behavior fall under one of two major theoretical frameworks: "choice probability" or "subjective value." These two formulations represent behavioral quantities closely tied to the decision process, but it is unknown whether one of these variables, or both, dominates the neural mechanisms that mediate choice. Value and choice probability are difficult to distinguish in practice, because higher-valued options are chosen more frequently in free-choice tasks. This distinction is particularly relevant for sensorimotor areas such as parietal cortex, where both value information and motor signals related to choice have been observed. We recorded the activity of neurons in the lateral intraparietal area while monkeys performed an intertemporal choice task for rewards differing in delay to reinforcement. Here we show that the activity of parietal neurons is precisely correlated with the individual-specific discounted value of delayed rewards, with peak subjective value modulation occurring early in task trials. In contrast, late in the decision process these same neurons transition to encode the selected action. When directly compared, the strong delay-related modulation early during decision making is driven by subjective value rather than the monkey's probability of choice. These findings show that in addition to information about gains, parietal cortex also incorporates information about delay into a precise physiological correlate of economic value functions, independent of the probability of choice.

Introduction

Decision making involves the transformation of information into a behavioral choice. In perceptual decision tasks, this flow of information links sensory processing to the selection of an action. Lateral intraparietal area (LIP) neurons are hypothesized to mediate direct sensorimotor transformations, responding to stimuli in a selective region of visual space and firing before a saccadic eye movement to the same location (Gnadt and Andersen, 1988; Andersen and Buneo, 2002). In addition to spatially congruent sensorimotor activity, LIP neurons also represent sensory information that can be used to specify saccade metrics even when the stimulus is not colocalized with the reinforced movement. For example, in different experimental paradigms, LIP activity has been shown to reflect accumulated motion evidence, target color, temporal information, and probabilistic cues (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Toth and Assad, 2002; Leon and Shadlen, 2003; Yang and Shadlen, 2007). Typically, decision-related activity in these experiments has been taken to represent the probability of choice: increased activation correlates with a higher likelihood of a correct decision, and thus a higher probability of behavioral selection (Sugrue et al., 2005; Gold and Shadlen, 2007).

While the choice probability framework has proven powerful, decision processes can also incorporate nonsensory, internally derived information such as value, strategic planning, and attention. In particular, recent experiments suggest that LIP activity reflects value information such as the probability or magnitude of reinforcement, reward history, and strategic game valuation (Platt and Glimcher, 1999; Dorris and Glimcher, 2004; Sugrue et al., 2004) even though these properties are not instantiated as sensory signals at the level of single trials. These findings have led to the proposal that decision-related activity may represent the subjective value of a specific action (Glimcher et al., 2005). In this framework, LIP activity combines all relevant reward information and sensory evidence into a single decision variable that reflects the overall subjective value of the saccade (or attentional target) encoded by the neuron under study. Like choice probability coding, subjective value is presumed to act via modulation of the spatially tuned response fields widely observed in parietal cortex.

Thus LIP neurons could be modulated by the behavioral probability that a response field movement would be chosen for execution or by the underlying subjective value of the action. Additionally, it has been suggested that both subjective value and choice probability may be represented, in sequential stages, during the decision process (Sugrue et al., 2005). Importantly, although subjective value and choice probability are separable quantities in principle, value and choice can be difficult to disambiguate: higher-valued options are chosen more frequently. For example, in experiments using matching law behavior, choice behavior is by definition directly proportional to the relative values of the options, and LIP activity correlates with both signals (Platt and Glimcher, 1999; Sugrue et al., 2004). In general, this correlation between choice probability and value confounds many free-choice paradigms, particularly those in

Received Nov. 19, 2009; revised March 4, 2010; accepted March 9, 2010.

We thank B. Lau, J. Kable, A. Caplin, G. Frechette, and E. DeWitt for discussions and reading of the manuscript, and M. Grantner and E. Ryklin for technical support.

Correspondence should be addressed to Kenway Louie, Center for Neural Science, New York University, 4 Washington Place, Room 809, New York, NY 10003. E-mail: klouie@cns.nyu.edu.

DOI:10.1523/JNEUROSCI.5742-09.2010

0270-6474/10/305498-10$15.00/0

which the differences in (or ratios between) the option values under consideration span only a narrow experimental range. The result is that existing studies cannot distinguish between these two representations, and it is unclear whether parietal cortex carries a value signal distinguishable from the probability of choice.

To investigate these issues, we recorded the activity of LIP neurons while monkeys performed an oculomotor intertemporal choice task between a small immediately available reward and a larger delayed reward. Using a traditional psychophysical approach, we measured the choice probabilities of two monkeys as they made these decisions, and using an experimental economics approach, we quantified the individual subjective value of rewards as a function of delay, enabling us to examine how each variable controls LIP activity.

Materials and Methods

Subjects and task. Two male rhesus monkeys (*Macaca mulatta*; monkey D, ~8.6 kg; monkey W, ~6.0 kg) were used as subjects. All experimental procedures were performed in accordance with the Public Health Service's *Guide for the Care and Use of Laboratory Animals* and approved by the New York University Institutional Use and Care Committee.

Experiments were conducted in a dimly lit sound-attenuated room. The monkeys were head-restrained and seated in a Plexiglas enclosure that permitted arm and leg movements. Visual stimuli were generated using an array of tristate light-emitting diodes (LEDs) situated on a tangent screen 145 cm from the eyes of the monkey. The LEDs formed a grid with points spaced at 2° intervals, spanning 40° horizontally and 40° vertically. Eye movements were monitored using the scleral search coil technique (Fuchs and Robinson, 1966) with horizontal and vertical eye position sampled at 500 Hz using a quadrature phase detector (Riverbend Electronics). Presentation of visual stimuli and water reinforcement delivery were controlled with an integrated software and hardware system (Gramalkn; Ryklin Software).

Each trial began with the monkey fixating a central fixation target. Two peripheral targets were then presented, a red target associated with a small immediate reward and a green target associated with a larger delayed reward. After 800 ms, the fixation target was dimmed for 200 ms, followed by the presentation of a central instruction cue for 500 ms. In forced-choice trials, the color of the central cue specified the saccade target; in free-choice trials, a yellow cue indicated that a saccade to either target would be rewarded. At 1500 ms, the central fixation cue was extinguished, indicating that the monkey was permitted to initiate a saccade; peripheral target cues were extinguished after the monkey completed a saccade to one of the presented targets. Rewards were delivered either immediately or after the designated delay; monkeys were not required to maintain fixation over the delay interval. In immediate reward trials, an additional interval was imposed after the immediate reward was delivered, thus equalizing the duration of immediate and delayed reward trials. Each session was conducted in blocks of 40 forced-choice followed by 20 free-choice trials. Delay and reward magnitudes were held constant across a block. Delays were varied between blocks and chosen to span choice threshold in behavior-only sessions. In electrophysiology sessions, delays were 0, 1, 2, 4, 8, and 12 s, presented in randomized order (3–6 blocks). The immediate reward was 0.130 ml of water; the delayed reward was constant in a session and either 0.143, 0.163, 0.196, or 0.260 ml of water.

Electrophysiological recording. Monkeys were implanted with a Cilux recording chamber (Crist Instruments) targeting the lateral bank of the intraparietal sulcus, centered 3 mm caudal and 12 mm lateral to the intersection of the midsagittal and interaural planes in either the left hemisphere (monkey D) or the right hemisphere (monkey W). Chamber location was verified using anatomical magnetic resonance imaging (3 T; Siemens). At the start of each recording session, a 23 gauge guide tube was positioned in a support grid (1 mm spacing; Crist Instruments) and inserted through intact dura. A tungsten steel electrode (8–10 MΩ; FHC) was lowered through the guide tube using a computer-controlled micropositioner (EPS; Alpha-Omega). Electrophysiological signals were amplified, bandpass filtered, and digitized, and individual neurons were isolated based on waveform characteristics (MAP; Plexon).

Within a given session, recording was initiated once stable electrophysiological signals were obtained from a depth corresponding to LIP according to the magnetic resonance images. Single intraparietal neurons were identified and response fields were characterized as previously described (Platt and Glimcher, 1999). Once a stable response field was estimated, the intertemporal choice task was run with the delayed reward target location placed within the estimated response field, and the immediate reward target placed outside the response field, typically in the opposite hemifield and at an equivalent distance from fixation. Neurons were recorded while monkeys performed 3–6 blocks of the intertemporal choice task, with randomized selection of delays between blocks. For neural analyses, the first two forced trials in a block to each target were excluded to minimize block transition effects, while all free-choice trials were included.

Behavioral analyses. The intertemporal choice task was conducted under four different conditions of delayed reward magnitude (0.143, 0.163, 0.196, and 0.260 ml), randomized across sessions, to quantify the discount function (or more precisely, the discounted "utility" function). Four choice curves and a discount function were fit to the complete binary choice dataset using a two-parameter binary logit model, with a separate fit for each monkey. The choice function:

$$p_{\mathrm{L}} = \frac{1}{1 + \exp(\beta(\mathrm{SV}_{\mathrm{L}} - \mathrm{SV}_{\mathrm{S}}))},$$

where p_{L} is the probability of choosing the delayed reward as a function of the difference between the subjective values of the two options (SV_{L}, SV_{S}) and a noise parameter β, and the discount function:

$$\mathrm{SV} = \frac{A}{1 + kD},$$

where the decline in subjective value SV is a function of delay D, amount A, and a discount parameter k, were simultaneously fit by maximum likelihood estimation. Bootstrap distributions were obtained for each discount factor k by resampling the sample distribution of behavioral data, treating individual blocks of free-choice data as samples. A bootstrap sample k was produced for each resample procedure and repeated for a total of 2000 iterations, and 95% percentile confidence intervals were quantified for significance testing.

Neural analyses. Individual LIP neurons were run under a single delayed reward magnitude condition, with neurons from monkey W collected under two conditions (0.143, 0.260 ml) and from monkey D under three conditions (0.143, 0.163, 0.260 ml). Because discounting data, normalized to the zero delay condition for that magnitude condition (see below), were not significantly different between different magnitude conditions, normalized neural data recorded under the magnitude conditions were combined together for each monkey.

For population neural analyses, each neuron was normalized by its mean neural firing rate across the zero delay condition trial. To construct the neural discount function, the mean normalized population activity (0–200 ms) was quantified for each delay condition and normalized to the population zero delay condition mean. This procedure produces a delay-dependent activity function, relative to the zero delay condition; this function allows comparison to the behavioral discount function, which describes relative subjective value as a function of delay. For forced-choice trial data, the first two trials to each target after a block transition were excluded from analysis so as to examine data only after the animal had already sampled the new reward contingencies. In free-choice trial blocks, the number of trials with saccades into the response field was delay dependent and therefore varied by block, resulting in some blocks with few or no trials where the monkey chose the delayed reward target (i.e., at long delays). For the neural discount function analyses, free-choice neural data were included only from blocks with a minimum of eight trials with a saccade into the response field.

To examine the relative contribution of subjective value and choice to LIP activity, univariate and multiple regression analyses were performed in non-overlapping 100 ms windows across the duration of free-choice trials. For each block of free-choice data, neural activity was quantified as mean firing rate across the block; because individual neurons were not recorded under all delay conditions, single-neuron firing rates were normalized by the average activity of the neuron in the zero delay condition to allow comparison across delays and neurons. Each block of free-choice data was associated with a subjective value, calculated from the delay to the delayed reward implemented in that block and the individual monkey discount function, and a choice probability, quantified as the average probability of delayed reward target choice during that particular block. Regression analyses were conducted on the population of block data points, combining across monkeys (n = 279 blocks). In each temporal window, univariate linear regression was performed with either subjective value or observed choice probability as predictors, and multiple linear regression was performed including both subjective value and choice probability. Neural data from all free-choice trials were included in this analysis.

Results

We trained two monkeys (*Macaca mulatta*) to perform an oculomotor version of an intertemporal choice task (Fig. 1*a*). In each trial the monkey viewed two targets, one associated with a small immediate reward and the other with a larger delayed reward. In forced-choice trials, a change in the color of the central fixation cue instructed the monkey to make a saccade to either the target that yielded the smaller immediate reward or the target that yielded the larger delayed reward; in free-choice trials, the monkey could select either reward. Each block of trials began with forced-choice trials, in which both targets had an equal probability of instruction, followed by a series of free-choice trials. Reward contingencies (delay and magnitude) were fixed during a block, so that for any given block monkeys first learned the values of the two alternatives and subsequently expressed their preference between the two options. The total trial duration was also fixed, regardless of the choices of the monkey, to ensure that selecting the smaller immediate reward could not lead to higher overall reward rates. To examine the effect of delay on subjective value, we varied between blocks the delay required to receive the larger reward, holding both reward magnitudes constant over the course of any single session.

Delay discounting behavior

Intertemporal choice behavior is governed by the delay to reward in a wide array of species (Mazur, 1987; Rachlin et al., 1991; Myerson and Green, 1995; Laibson, 1997; Kim et al., 2008). In the present task, monkeys chose the saccade yielding the larger reward when both rewards were offered immediately, but as the delay to the larger reward increased they eventually preferred the smaller but more immediate option (Fig. 1*b*). Because trial duration within a block was identical regardless of the monkey's choice, preference for the smaller but more immediately available option reflects a true subjective preference rather than an underlying rate-maximization strategy.

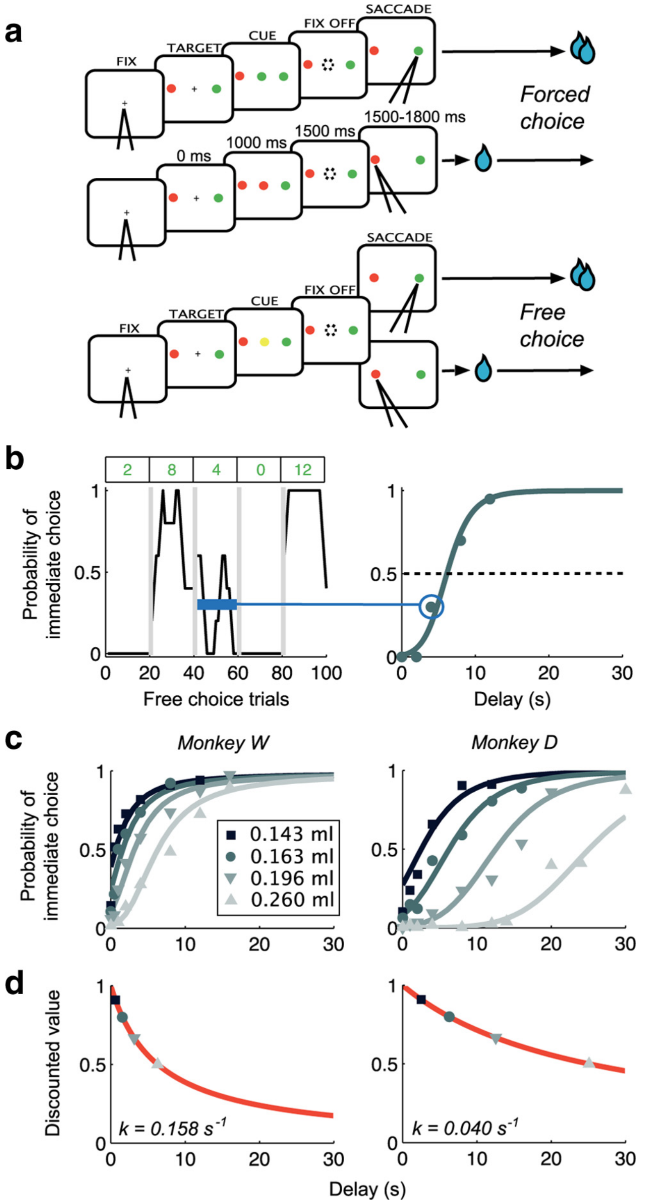

Figure 1. Temporal discounting behavior in the intertemporal choice task. ***a***, Trial structure in forced and free-choice variants of the task. Visual targets were associated with either a smaller reward delivered immediately (red) or a larger reward delivered after a delay (green). Trials with immediate rewards were followed by an additional wait interval so that overall trial duration was identical for both immediate and delayed reward trials in a given block. ***b***, Example single session free-choice data (left) and preference curve (right). Position–reward contingencies were fixed in a block of trials (40 forced followed by 20 free), and delay to the larger reward was varied between blocks (green numbers, in seconds). Preference curves are generated from average block choice data (example, blue line). In the preference curve, the indifference point (choice probability = 0.5) represents the delay at which the subjective value of the delayed reward is equivalent to that of the immediate reward. ***c***, Monkey choice behavior, analyzed separately for individual animals. Choice data were collected for four different magnitudes of delayed reward. The preference curves and discount function were simultaneously fit to the choice data using a binary logit model (see Materials and Methods). ***d***, Behavioral discount functions. The best fit hyperbolic discount function for each monkey is shown in red. Symbols indicate the four preference curve indifference points, which by construction lie on the discount function. Each discount function is characterized by its discount factor k, which quantifies how steeply subjective value declines as a function of delay.

Logistic choice functions fit to these data quantify an indifference point (a point of subjective equality) for each magnitude condition: the delay at which the monkey showed equal preference for the small immediate and larger delayed rewards. Figure 1*c* shows delay-dependent choice behavior for four different magnitudes of delayed reward. Changing the size of the delayed reward across days correspondingly shifted the position of the indifference point; monkeys would wait longer for larger rewards. Together, these temporally defined indifference points for different magnitudes of reward describe a function, closely related to the discounted utility function of neoclassical economic theory, which we term the behavioral discount function (Fig. 1*d*). Consistent with previous behavioral studies (Mazur,

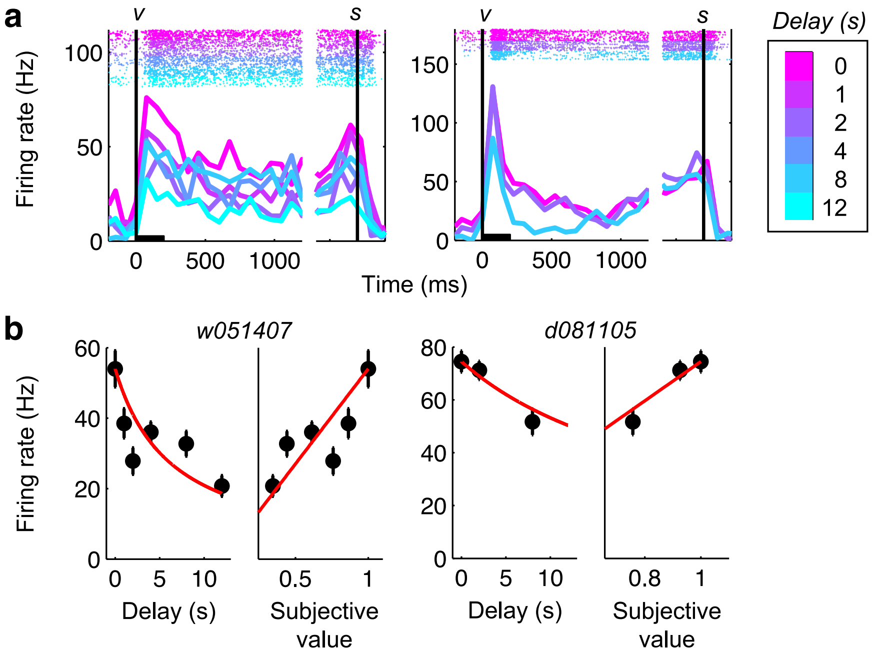

Figure 2. Activity of single LIP neurons during the intertemporal choice task. ***a***, Activity of two example single neurons from different monkeys in forced trials with saccades into the response field. Firing rate histograms and rasters are color-coded by delay to the delayed reward, and shown aligned to both the time of visual target onset (*v*) and the time of saccade (*s*). Instruction to saccade was always presented 1500 ms after visual target onset. ***b***, Firing rate as a function of delay and subjective value. Average firing rates (±SEM) are shown for the 200 ms epoch immediately following visual target presentation. The activity of these neurons is strongly modulated by delay, as evident when neural activity is plotted as a function of delay (left) and delay-dependent subjective value (right) derived from the behavioral discount function. Red lines represent the best fit hyperbolic discount function to the neural data.

1987; Rachlin et al., 1991; Myerson and Green, 1995; Laibson, 1997), this decline in subjective reinforcer value (SV) as a function of delay is well described by a simple hyperbolic equation:

$$\mathrm{SV} = \frac{A}{1 + kD},$$

where A is the reward magnitude, D the delay to reinforcement, and k the discount factor quantifying how steeply the discount function declines. For each animal, the four preference curves and the discount function were fit simultaneously using a binary logit model, which assumes a minimal number of free parameters (two). Importantly, while these discount functions were stable for each individual animal (see supplemental Fig. S1, available at www.jneurosci.org as supplemental material), there is a significant difference in the rates of discounting between the two monkeys we studied (monkey D: $k = 0.040\ \mathrm{s}^{-1}$; monkey W: $k = 0.158\ \mathrm{s}^{-1}$; $p < 10^{-45}$, permutation test).

Delay modulation of neural activity

Given these choice and behavioral discount functions, is the activity of neurons in LIP better correlated with the objective value (magnitude) of the offered reward, the subjective value of the offered reward, or the subsequent choice behavior? To answer this question, we recorded the activity of 71 LIP neurons while monkeys performed the intertemporal choice task described above. Neurons in LIP are spatially tuned, increasing their firing rate when a visual stimulus appears in a circumscribed region of space termed the response field. Consistent with a decision-related visuomotor transformation, many of these neurons also show presaccadic activity specifically for eye movements that carry gaze into the response field. To examine the effect of delay on neural activity, we placed the target yielding the larger delayed reward in the response field of each recorded neuron and placed the target yielding the small immediate reward outside the response field. We then monitored neuronal activity as the delay to reward was changed across blocks.

Figure 2 shows example activity from single LIP neurons during forced-choice trials in which the monkey was instructed to make a saccade to the delayed reward target. In all such trials, the monkey views the same visual stimuli and performs the same saccadic movement; only the delay to reinforcement after the saccade is complete varies between blocks of trials. As previously reported (Gnadt and Andersen, 1988), the spiking activity of these neurons evolves throughout the trial, with activation typically highest immediately after target onset, then maintained above baseline throughout the trial, and finally rising before a saccade into the response field (Fig. 2*a*). We found that under longer delay conditions, LIP neurons showed lower firing rates throughout much of the trial despite the presence of identical reward magnitudes. These neurons were thus sensitive to delay to reward, a variable that influences both subjective value and choice but not objective value in our task (Fig. 2*b*). The majority of sampled LIP neurons (47/71, 66.2%) showed significant modulation by delay (regression analysis, activity in the epoch 0–200 ms after target onset), and delay strongly modulates the population response of these neurons (Fig. 3*a*).

Neural representation of subjective value

Does this effect of delay represent the monkeys' subjective valuation of delayed reward? To answer this question, we computed for each animal a neural discount function, defined—analogously to the traditional behavioral discount function described above—as the best hyperbolic fit to the population firing rate as a function of delay. We examined activity in the epoch immediately after target onset (0–200 ms), pooling data for all of the neurons studied in each animal, at all magnitudes, normalized by average response to the zero delay condition. Thus activity is represented as a fraction of the neuronal response to a given immediate reward magnitude. Importantly, because the behavioral discount functions differed significantly between monkeys, we analyzed the individual monkey neural data separately. Because the behavioral (red line) and neural (black line) discount functions are defined in the same units (discounted value as a function of delay), we can directly compare delay-discounted subjective value and LIP activity (Fig. 3*b*). We found that for each monkey the neural discount function matches the behavioral discount function with surprising precision (monkey W: $k_{\mathrm{neural}} = 0.157\ \mathrm{s}^{-1}$, $k_{\mathrm{behav}} = 0.158\ \mathrm{s}^{-1}$; monkey D: $k_{\mathrm{neural}} = 0.038\ \mathrm{s}^{-1}$, $k_{\mathrm{behav}} = 0.040\ \mathrm{s}^{-1}$; 95% bootstrap confidence intervals shaded in Fig. 3*b*). Furthermore, each neural discount function differs significantly from both the behavioral and neural discount functions of the other monkey (see supplemental material, available at

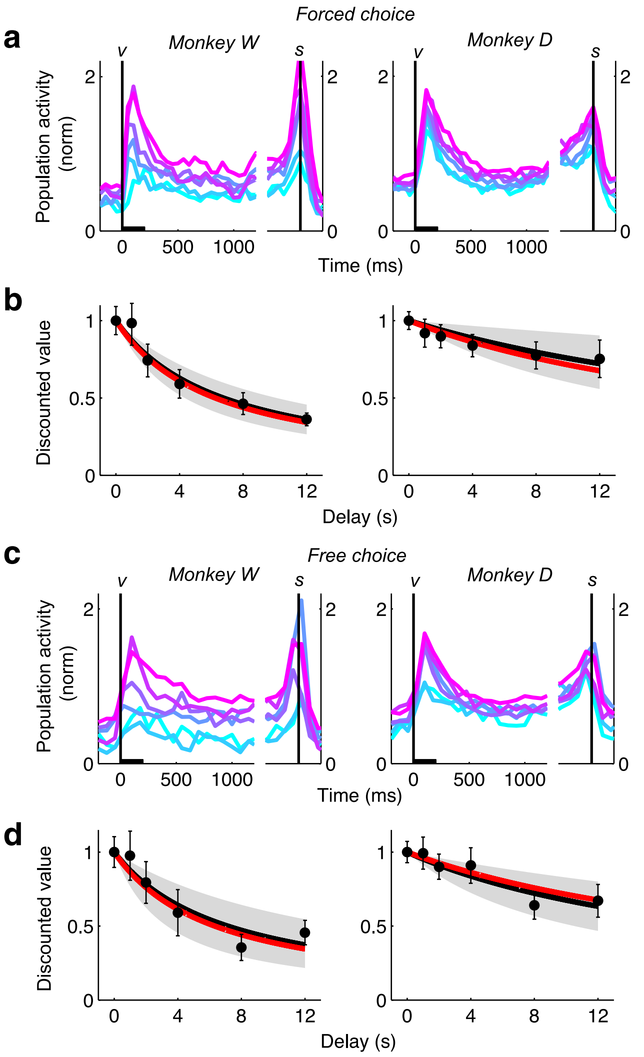

Figure 3. Population analyses: neural and behavioral discount functions. ***a***, Average forced-choice activity, segregated by delay. Firing rate histograms are shown for trials with saccades into the neural response field, color-coded by delay and aligned to visual target onset (*v*) and saccade initiation (*s*) as in Figure 2. Firing rate activity is presented separately for the two monkeys (monkey W, $n = 23$ neurons, monkey D, $n = 48$ neurons). ***b***, Forced-choice neural and behavioral discount functions. Mean (±SEM) normalized LIP activity is shown as a function of delay, relative to activity at zero delay. The neural discount function is the best fit hyperbolic curve to the neural activity as a function of delay (black line); 95% bootstrap confidence intervals are plotted as shaded gray regions. Each neural discount function matches the respective monkey's behavioral discount function (red lines). ***c***, Average free-choice activity, segregated by delay. Data are presented identically to forced-choice trials, for all trials in which the monkey chose to saccade into the neuron's response field. ***d***, Free-choice neural and behavioral discount functions.

www.jneurosci.org), suggesting a specific psychometric–neurometric match between perceived value and neuronal activity in each individual.

The preceding data, however, only reflect neural activity during forced-choice trials. If these representations drive decision-making processes, then subjective value should modulate LIP neurons during free-choice trials as well. We therefore examined neural activity during free-choice trials, restricting our analysis to the subset of trials in which the monkey chose the target in the neuron's response field (RF). Despite a smaller number of sampled trials imposed by the subject's preferences, it is clear that LIP population activity during free choice is strongly modulated by delay (Fig. 3*c*, displayed for trials with saccades into the RF), and the free-choice neural discount functions also match the behavioral discount functions (Fig. 3*d*, monkey W: $k_{\text{neural}} = 0.140\ \text{s}^{-1}$, $k_{\text{behav}} = 0.158\ \text{s}^{-1}$; monkey D: $k_{\text{neural}} = 0.048\ \text{s}^{-1}$, $k_{\text{behav}} = 0.040\ \text{s}^{-1}$; 95% bootstrap confidence intervals shaded).

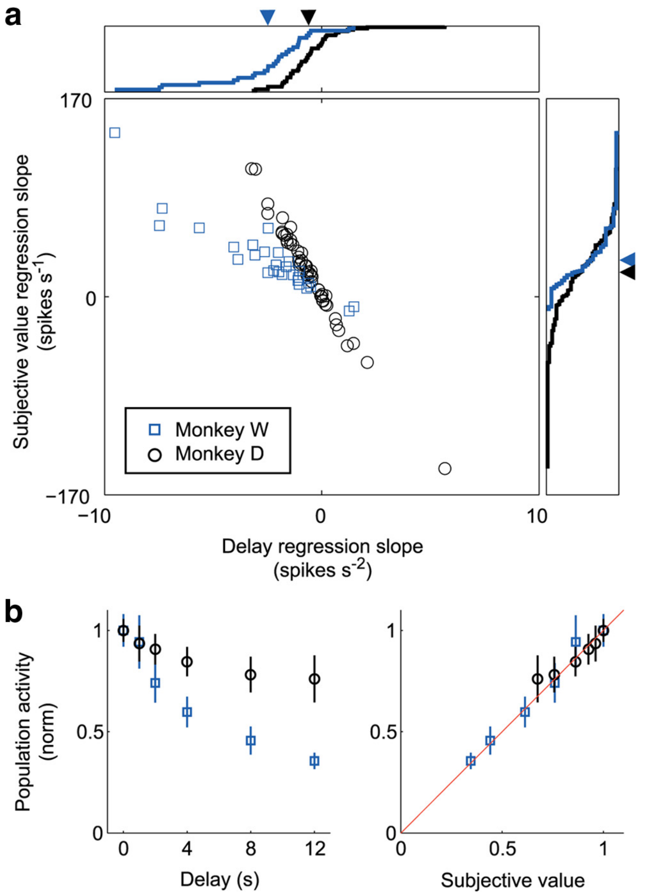

Figure 4. LIP modulation reflects delay-discounted subjective value. ***a***, Linear regression slopes quantifying LIP modulation by either delay or subjective value. Each data point represents an individual LIP neuron from monkey W (blue) or monkey D (black). Cumulative marginal distributions are plotted for delay (above) and subjective value (right); arrows indicate mean slope values. Subjective value slopes are similar ($p = 0.70$), but delay slopes differ significantly between monkeys ($p = 0.0008$), suggesting that subjective value drives LIP modulation in both animals. ***b***, Normalized population activity as a function of delay (left) and subjective value (right). The combined data are much better characterized by subjective value ($R^2 = 0.95$) than delay ($R^2 = 0.58$).

Comparison of neural activity across both monkeys confirms that subjective value is a more parsimonious explanation of LIP activity than delay. For each neuron, we quantified the influence of either delay or subjective value using separate linear regression models. Figure 4 plots the regression slopes relating delay to firing rate versus the regression slopes relating subjective value to firing rate for all neurons, separated by monkey, as well as the cumulative marginal regression weight distributions. This figure indicates that while LIP neurons in both monkeys show similar regression slopes for subjective value ($p = 0.70$, Wilcoxon rank-sum test), delay regression slopes differ significantly between monkeys ($p = 0.0008$, Wilcoxon rank-sum test). Thus, while subjective value controls LIP activity in the same manner in both monkeys, delay to reward more strongly modulates neural firing rates in monkey W than monkey D. Importantly, these neural results are consistent with the behavioral data, in which subjective value also declines more quickly as a function of delay in monkey W ($k = 0.158$) than in monkey D ($k = 0.040$)—when subjective value is a steeper function of delay, neural firing rates are also more strongly delay modulated.

This encoding of subjective value is evident in the strong correspondence between the behavioral and neural discount functions seen in Figure 3. To examine this directly, we plot in Figure 4*b* the normalized population neural activity of both monkeys as a function of delay (left) and subjective value (right). Note that the computation of subjective value relies solely on choice behavior in the discounting task, and does not rely on neuronal data. Nevertheless, compared to delay, subjective value explains more of the variance in the population data in both monkey W ($R_{\mathrm{d}}^2 = 0.87$, $R_{\mathrm{SV}}^2 = 0.96$) and monkey D ($R_{\mathrm{d}}^2 = 0.87$, $R_{\mathrm{SV}}^2 = 0.93$). Furthermore, combining all data, population LIP activity is much better characterized as a function of subjective value ($R^2 = 0.95$) than delay ($R^2 = 0.58$).

Distinct representations of value and choice

The preceding data show that immediately after target onset, LIP activity precisely covaries with delay-discounted subjective value. Decision areas such as LIP are also known to ultimately signal the chosen action (Gnadt and Andersen, 1988; Andersen and Buneo, 2002), which is a function of subjective value—subjectively higher-valued targets are by definition chosen over lower-valued ones (Platt and Glimcher, 1999; Sugrue et al., 2004). Thus LIP activity appears to represent both input and output variables necessary in decision making: value and selected action. It has, however, been proposed that option values are transformed into choice probability functions as an intermediate step in generating stochastic choice behavior, and that LIP activity may actually reflect these underlying choice probabilities (Sugrue et al., 2005). Could the modulation we observe simply reflect the animals' upcoming probability of choice rather than a distinct representation of subjective value per se?

To examine the relative influence of subjective value and choice probability on LIP activity, we quantified single-neuron firing rates for individual blocks of free-choice trials ($n = 279$ blocks); each block was associated with a subjective value (determined by the specified delay to reward and the individual-specific discount function) and a mean choice probability (averaged over the monkey's choices in that block). Though subjective value is directly calculated from overall choice behavior, two properties of this dataset allow us to effectively disassociate subjective value and choice probability. First, choice behavior at the block level exhibits variability between blocks for identical subjective value conditions. Second, the subjective value of the delayed target continues to diminish even after choice probabilities have reached asymptotically high levels (Fig. 1). As evident in Figure 5, there is considerable variation between these two parameters, particularly when data are grouped across magnitude conditions (variance inflation factor = 1.26). Utilizing this relative independence, we used regression analyses to determine whether neural activity is better correlated with choice probability or subjective value, and how this encoding changes during decision making.

We performed a sliding window analysis across the duration of all free-choice trials to examine modulation of LIP activity over the decision process. In each non-overlapping 100 ms bin, we quantified the influence of subjective value or choice probability on normalized block-averaged firing rates by univariate linear regression (individual neuronal firing rates normalized by mean zero delay activity; see Materials and Methods). As shown in Figure 6*a*, subjective value (blue line) explains a significant proportion of LIP population variability across the length of free-choice trials, with a peak in the coefficient of determination (R^2) immediately after target presentation. The small but significant value modulation during fixation likely reflects the task design, in which target locations and rewards are fixed within blocks, making information about the value of saccades available to the animals before target onset. In contrast to strong value modulation, neural activity early in the trial is minimally explained by choice probability (red line). R^2 values from multiple regression analysis (black line) confirm that including choice probability as a factor provides little additional explanatory power beyond that of subjective value alone.

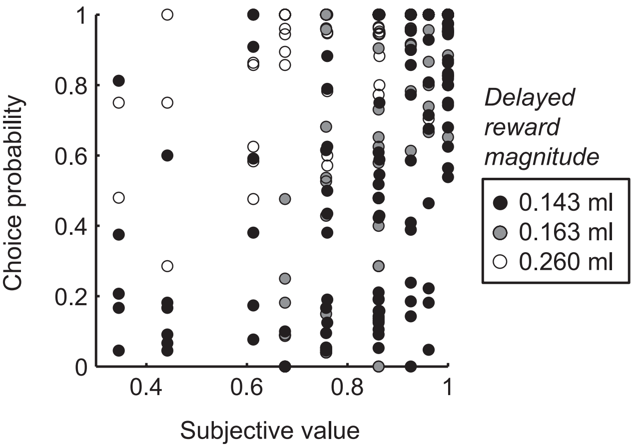

Figure 5. Dissociation of subjective value and choice probability. Shown is a scatter plot of block-averaged choice probability and subjective value. Each point represents a single block of free-choice trials during neural recording sessions; all blocks with neural data are displayed ($n = 279$ blocks). Subjective value was determined by the delay to reward of the response field target and the individual monkey discount function. Choice probability was quantified as the mean probability that the monkey chose the delayed reward target in that particular block. For clarity, data collected under different magnitude conditions are colored as indicated. The grouped data show considerable dissociation between subjective value and choice probability (variance inflation factor = 1.26).

To examine these results more directly, we plot in Figure 6*b* population average firing rates as a function of either subjective value or choice probability for the 0–1000 ms period following target presentation. Consistent with the sliding window analysis, firing rate is significantly dependent on subjective value ($R_{\mathrm{SV}}^2 = 0.151$, $p < 10^{-11}$, *F* test) but not choice probability ($R_{\mathrm{CP}}^2 = 0.005$, $p = 0.23$, *F* test). Furthermore, when data are restricted to blocks where choice probability was equal to 1 (i.e., the monkey chose the delayed option exclusively), LIP activity is still a significant function of subjective value ($R_{\mathrm{SV}}^2 = 0.163$, $p < 0.00001$, *F* test; data not shown). We note that the use of local, block-average choice probability represents a conservative approach to estimating the influence of subjective value: under the alternative hypothesis that firing rates are driven by choice probability and not subjective value, using block-level probabilities more closely tied to daily variations in behavior and global subjective values would make it more difficult to detect subjective value-related modulations. To ensure that our results did not depend on the particular definition of choice probability we used, we repeated the univariate regression analysis with global, experiment-averaged choice probabilities. Using this formulation of choice probability subjective value ($R_{\mathrm{SV}}^2 = 0.151$, $p < 10^{-11}$, *F* test) is still a stronger predictor of LIP activity than choice probability ($R_{\mathrm{CP}}^2 = 0.035$, $p = 0.002$, *F* test). The higher coefficient of determination for global versus local choice probability likely arises from the relationship between global choice probability and subjective value (in our analysis, average choice probability is a function of sub-

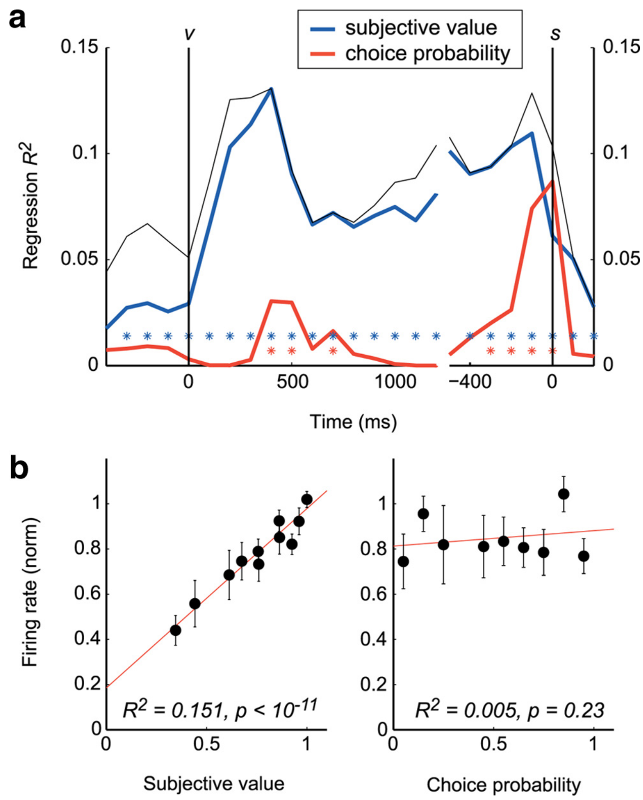

Figure 6. Value representation independent of choice probability. ***a***, Relationship between LIP neural activity and either subjective value (blue) or choice probability (red) during free-choice trials, plotted as the coefficient of determination (R^2). Neural activity was aligned to target onset (*v*) for the beginning of the trial and saccade initiation (*s*) for the end of the trial. Regression analyses were performed in contiguous non-overlapping 100 ms windows across the duration of the trial using either subjective value or choice probability as the independent factor. R^2 values for multiple regression analyses using both subjective value and choice probability as independent factors are also shown (black). Asterisks indicate time windows when subjective value or choice explained a significant proportion of population activity ($p < 0.05$, *F* test). ***b***, Mean LIP activity as a function of subjective value (left) or choice probability (right) during the interval 0–1000 ms after target presentation. Regression lines (red) indicate that LIP population activity is significantly dependent on subjective value ($R^2_{SV} = 0.151, p < 10^{-11}$, *F* test) but not choice probability ($R^2_{CP} = 0.005, p = 0.23$, *F* test).

jective value); when both global choice probability and subjective value are included in a multiple regression analysis, LIP activity is only dependent on subjective value (SV regression slope = 0.96, 95% C.I. [0.65 1.24], global CP regression slope = 0.13, 95% C.I. [−0.28 0.03]). Thus, LIP neurons are dependent on subjective value and not choice probability, defined either locally or globally, early in the decision process.

LIP activity does not show peak modulation by choice probability in this task until late in trials, in the interval immediately preceding saccadic eye movement (Fig. 6*a*, red). Thus, there is a shift in the population response: at the onset of each trial, activity reflects the subjective value of the option in the response field, regardless of the subsequent choice behavior; as the trial progresses, activation associated with the monkey's choice grows, peaking immediately before movement onset. These results were not driven by either magnitude or individual animal effects; additional regression analyses on data segregated by either reward magnitude or individual subject (see supplemental material, available at www.jneurosci.org) demonstrated equivalent results in both animals: LIP activity reflects subjective value and not choice probability.

Discussion

We examined the relationship between parietal neuron activity and subjective values for actions during value-guided decision making. Monkeys were trained to choose between rewards differing in magnitude and delay to reinforcement in a delay-discounting task, enabling a precise behavioral quantification of the subjective values of saccadic targets. We found that LIP activity is tightly correlated to the delay-discounted value of a saccade, independent of the underlying probability of choice. While previous studies have demonstrated subjective value-related activity in LIP (Platt and Glimcher, 1999; Dorris and Glimcher, 2004; Sugrue et al., 2004), our results extend these findings in three significant ways. First, the sensitivity of parietal activity to delay indicates that this value representation extends beyond manipulations of expected reward like probability of reinforcement and magnitude to include exclusively subjective components of value such as delay. Neural signals related to reinforcement delay have been observed in other brain regions, notably the frontal cortices (Roesch and Olson, 2005; Kim et al., 2008), but typically only modulate a minority of neurons. The present discounting modulation was observed to strongly influence a majority of parietal neurons, suggesting the importance of subjective value (action value) coding at this stage of visuomotor processing (Glimcher et al., 2005). Second, precise quantification of delay-discounted subjective value demonstrates a surprisingly accurate match between the behavioral and neural value representations. This relationship is similar to the psychometric–neurometric correspondence of sensory signals (Newsome et al., 1989), now extended to parietal cortex and the domain of value representation. This precise neural representation of subjective value suggests that parietal activity is not simply modulated by reward-related variables, but instead may reflect the underlying neural value signal guiding choice, a neural correlate of the economic discounted utility function. Finally, our comparison of choice and value signals demonstrates for the first time that LIP neurons carry a subjective value signal that can be separated from signals encoding the probability of choosing an option. We find that value and choice signals are temporally dissociated, with subjective value representation early in the decision process giving way to representation of the chosen action near the time of saccade; this value-to-choice transformation in neural activity may represent the critical input and output stages hypothesized in standard models of the value-guided decision process.

Choice probability and subjective value

In studies of perceptual decision making rooted in signal detection theory, choice probabilities are typically constructed by presenting a perceptually ambiguous stimulus that varies from trial to trial and measuring the aggregate probability that a subject will make one of two evaluative responses. In the classic paradigm of this type, monkeys view a random-dot stimulus that contains net image motion in one of two possible directions and choose one of two responses; if the monkey selects the response associated with the correct motion signal, it receives a reward (Newsome et al., 1989; Britten et al., 1992). The goal is to demonstrate that neural activity (the neurometric function) is correlated with measured response probabilities (the psychometric function) across different stimulus conditions, and evidence for such neurometric–psychometric matches exists in multiple cortical areas.

Neoclassical economists developed an alternative approach to the behavioral study of choice, hypothesizing that subjects choosing between two gains behave as if those gains were represented on an internal scale (Von Neumann and Morgenstern, 1944;

Samuelson, 1947; Savage, 1954). For this reason, a number of neurobiological studies have proposed that the brain must contain neural signals representing the subjective values of options in a way that is at least partially independent of observed choice probabilities (Dorris and Glimcher, 2004; Glimcher et al., 2005; Padoa-Schioppa and Assad, 2006, 2008). Employing economic theory, these studies argue that the largely transitive nature of monkey choices necessitates an underlying representation of subjective value that is distinct from choice (or choice probabilities). The correlation between neural activity and external reward value is taken as evidence that these signals encode subjective value, with an implied physiological mapping rule from external to mean internal value.

In practice, however, distinguishing these two frameworks is difficult because choice probability and subjective value are often tightly correlated. Consider for instance the recent neurophysiological decision-making studies using the matching law (Sugrue et al., 2004; Lau and Glimcher, 2008), in which reward magnitude was systematically manipulated while choices were observed. Under the specified experimental conditions, the animals distributed their choices according to Herrnstein's matching law:

$$\frac{R_1}{R_2} = \alpha \frac{C_1}{C_2},$$

where R_1/R_2 is the ratio of reward magnitudes, C_1/C_2 is the observed ratio of choices, and α is a fixed constant. When behavior follows the matching law, choice probabilities and relative reward values are directly related. Thus, while there is increasing evidence for decision-related neurophysiological signals, these studies cannot discriminate between the choice probability and value frameworks.

In this experiment, we took advantage of two characteristics of choice behavior to effectively dissociate subjective value and choice probability. First, stochastic choice behavior is a function of the values of both options in a decision. Critically, this means that the subjective value of a single option (the RF target) can vary widely without a change in the associated choice probability, if the value of the other option remains much lower (or higher). Second, choice behavior at the block level exhibits considerable variability between blocks for identical subjective value conditions, suggesting a role for additional, value-independent sources of variance in choice. Together, these characteristics produced an effective experimental dissociation of value and choice probability, enabling a comparison of their relative influences on LIP activity.

Subjective value and decision making

We examined temporal discounting behavior, which displays a clear, well established behavioral subjective value signal: the discount function. Detailed quantification of the hyperbolic form of the discount function showed that parietal activity encodes a precise neural correlate of subjective value. Furthermore, comparison of neural activity to both the discount function and observed choice probabilities revealed that subjective value is the primary influence on early LIP activity, independent of the ultimate choice behavior. The striking correspondence between LIP activity and the discount function across animals suggests that neural activity is linearly related to the internal, subjective experience of value.

Classically, neurons in LIP respond to a spatially restricted subset of visual stimuli and saccadic eye movements, suggesting that these neurons mediate the visuomotor transformations underlying saccade selection and attentional allocation (Andersen and Buneo, 2002; Goldberg et al., 2002). However, LIP activity has also been shown to be modulated by additional classes of information, including more abstract, nonspatial task variables ranging from color or shape, to elapsed time, to reward probabilities, to the accumulation of sensory signals (Platt and Glimcher, 1999; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Toth and Assad, 2002; Leon and Shadlen, 2003; Sugrue et al., 2004; Janssen and Shadlen, 2005; Yang and Shadlen, 2007). While each of these responses can be characterized with a task-specific model, one is led to wonder whether a unifying framework exists that could relate these various findings.

We show here that the initial activity of a given LIP neuron during intertemporal choice is tightly correlated with the delay-dependent subjective value of the associated saccade. This representation of subjective value may allow a reinterpretation of many previous LIP studies within a broader common framework. For example, early studies showed that LIP activity in a motion discrimination task reflects the integral of the motion signal, a quantity encoding the accumulated evidence for motion in a particular direction (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002). Because reward was contingent on the monkey correctly indicating the true motion direction with a saccade, the integral of the motion signal was very closely related to the probability of reward—the more evidence for a particular direction of motion, the higher the subjective value of the associated saccade. Similarly, LIP activity in tasks involving the perception of elapsed time could also be seen as reflecting not time per se, but rather how such temporal information affects the instantaneous subjective value the subject places on a particular eye movement (Leon and Shadlen, 2003; Janssen and Shadlen, 2005). More generally, LIP neurons have been shown to encode stimulus attributes such as color only when such features are behaviorally relevant for obtaining rewards (Toth and Assad, 2002); such activity may not reflect color, but rather the information color carries about the value of making a particular saccade. Given the relative specificity of LIP for eye movements compared to other actions such as reaches (Snyder et al., 1997; Andersen and Buneo, 2002), the representation of saccadic subjective value we observe in LIP may be paralleled by similar subjective value coding for different types of actions in adjacent regions of parietal cortex (e.g., the parietal reach region).

Information about the subjective values of actions is certainly not unique to the parietal cortex, but present in a larger brain network that processes rewards, updates stored value information, and guides behavior (Schultz, 2004). One would expect the nature of such value information to vary from region to region in the brain, in a way that corresponds to the specific role of a given region in learning values and guiding behavior. For example, activity in dopaminergic nuclei, postulated to guide the learning of stimulus and action values, has been hypothesized to encode the difference between predicted and received rewards (Schultz et al., 1997). In decision-related areas, value information would be expected to closely approximate subjective valuations because such valuations are quantified directly from choice behavior. This representation of subjective value should thus combine all relevant information guiding choice, ranging from reward characteristics like magnitude to cost information such as required delay or effort.

We thus propose that subjective value representation in LIP operates within the existing parietal spatial framework: the subjective value associated with a given visuospatial location modulates the corresponding response field activity. It should be noted that while most LIP activity is spatially tuned, consistent with intentional or attentional activity, recent studies have also demonstrated nonspatial modulation by information such as learned categorical membership, effector limb usage, stimulus shape characteristics, and cognitive task rules (Sereno and Maunsell, 1998; Stoet and Snyder, 2004; Freedman and Assad, 2006; Oristaglio et al., 2006). The value framework may be difficult to reconcile with these nonspatial functions of LIP, and such nonspatial processing may represent an additional role for parietal cortex in visuomotor processing.

Parietal representations of value and attention

In addition to its role in oculomotor decision making, the parietal cortex is also modulated by both top-down and bottom-up attentional processing (Gottlieb et al., 1998; Goldberg et al., 2002), raising the question of whether the signals we observed in area LIP reflect the allocation of attention independent of any movement-related phenomena. It is often difficult to separate the effects of value and attention, since these concepts are closely coupled in the real world: attention is naturally directed toward more valuable objects or locations (Maunsell, 2004). Several details tentatively suggest, however, that the subjective value model may account for LIP activity in a way that is dissociable from general models of attentional allocation in this particular experiment. First, LIP firing rates are strongly correlated with subjective value even during the cue presentation period of forced trials, when the monkeys might be expected to direct their attention toward the central instruction cue. This finding is analogous to data from motion discrimination experiments, where the activity of LIP neurons reflects the accumulated motion information for or against a particular saccade, even though the monkeys are presumed to be attending the central motion stimulus (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002). Together these findings suggest that the locus of spatial attention does not uniquely specify LIP activity. Second, we observed reaction times that were nonmonotonic functions of delay (supplemental Fig. S8, available at www.jneurosci.org as supplemental material). Given the general relationship between attention and reaction times (Posner, 1980), these data also tentatively suggest that in this task attention may not be strongly correlated with delay to reward or subjective value. However, without direct behavioral measures of attentional allocation, we cannot exclude the possibility that our data reflect a delay- and reward-dependent allocation of spatial attention; this could be explicitly addressed in future work by using a nonspatial choice mechanism, such as a lever release. Given the strong one-to-one correspondence between LIP single unit activity and delay-discounted value we observed, if attention mediates this parietal modulation then these findings would imply the novel conclusion that subjective value serves as a primary and precise determinant of attentional allocation in this task.

Conclusions

We find that neural activity in the posterior parietal cortex is linearly related to the private, idiosyncratic experience of subjective value. We inferred this from a novel type of psychometric–neurometric match, one that specifically relates a subjective internal percept of value to a neural activation; such a physiological variable provides an empirical link between brain function and existing theoretical models of value, such as utility. Over the course of the decision process, this close match between LIP activity and subjective value evolves into a correlation between activity and choice in these same neurons. Both the unexpectedly linear mapping between activity of LIP neurons and subjective preference and the transition these neurons undergo during choice are precisely the kind of signals expected in decision-making circuits, and may provide avenues for future studies at the intersection of valuation and decision making.

References

Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. Annu Rev Neurosci 25:189–220.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci 12:4745–4765.

Dorris MC, Glimcher PW (2004) Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. Neuron 44:365–378.

Freedman DJ, Assad JA (2006) Experience-dependent representation of visual categories in parietal cortex. Nature 443:85–88.

Fuchs AF, Robinson DA (1966) A method for measuring horizontal and vertical eye movement chronically in the monkey. J Appl Physiol 21:1068–1070.

Glimcher PW, Dorris MC, Bayer HM (2005) Physiological utility theory and the neuroeconomics of choice. Games Econ Behav 52:213–256.

Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. Exp Brain Res 70:216–220.

Gold JI, Shadlen MN (2007) The neural basis of decision making. Annu Rev Neurosci 30:535–574.

Goldberg ME, Bisley J, Powell KD, Gottlieb J, Kusunoki M (2002) The role of the lateral intraparietal area of the monkey in the generation of saccades and visuospatial attention. Ann N Y Acad Sci 956:205–215.

Gottlieb JP, Kusunoki M, Goldberg ME (1998) The representation of visual salience in monkey parietal cortex. Nature 391:481–484.

Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. Nat Neurosci 8:234–241.

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

Laibson D (1997) Golden eggs and hyperbolic discounting. Q J Econ 112:443–477.

Lau B, Glimcher PW (2008) Value representations in the primate striatum during matching behavior. Neuron 58:451–463.

Leon MI, Shadlen MN (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. Neuron 38:317–327.

Maunsell JH (2004) Neuronal representations of cognitive state: reward or attention? Trends Cogn Sci 8:261–265.

Mazur JE (1987) An adjusting delay procedure for studying delayed reinforcement. In: The effect of delay and of intervening events on reinforcement value (Commons M, Mazur J, Nevin J, Rachlin H, eds), pp 55–73. Hillsdale, NJ: Lawrence Erlbaum.

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

Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. Nature 341:52–54.

Oristaglio J, Schneider DM, Balan PF, Gottlieb J (2006) Integration of visuospatial and effector information during symbolically cued limb movements in monkey lateral intraparietal area. J Neurosci 26:8310–8319.

Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. Nature 441:223–226.

Padoa-Schioppa C, Assad JA (2008) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. Nat Neurosci 11:95–102.

Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. Nature 400:233–238.

Posner MI (1980) Orienting of attention. Q J Exp Psychol 32:3–25.

Rachlin H, Raineri A, Cross D (1991) Subjective-probability and delay. J Exp Anal Behav 55:233–244.

Roesch MR, Olson CR (2005) Neuronal activity in primate orbitofrontal cortex reflects the value of time. J Neurophysiol 94:2457–2471.

Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci 22:9475–9489.

Samuelson PA (1947) Foundations of economic analysis. Cambridge: Harvard UP.

Savage LJ (1954) The foundations of statistics. New York: Wiley.

Schultz W (2004) Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology. Curr Opin Neurobiol 14:139–147.

Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275:1593–1599.

Sereno AB, Maunsell JH (1998) Shape selectivity in primate lateral intraparietal cortex. Nature 395:500–503.

Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J Neurophysiol 86:1916–1936.

Snyder LH, Batista AP, Andersen RA (1997) Coding of intention in the posterior parietal cortex. Nature 386:167–170.

Stoet G, Snyder LH (2004) Single neurons in posterior parietal cortex of monkeys encode cognitive set. Neuron 42:1003–1012.

Sugrue LP, Corrado GS, Newsome WT (2004) Matching behavior and the representation of value in the parietal cortex. Science 304:1782–1787.

Sugrue LP, Corrado GS, Newsome WT (2005) Choosing the greater of two goods: neural currencies for valuation and decision making. Nat Rev Neurosci 6:363–375.

Toth LJ, Assad JA (2002) Dynamic coding of behaviourally relevant stimuli in parietal cortex. Nature 415:165–168.

Von Neumann J, Morgenstern O (1944) Theory of games and economic behavior. Princeton: Princeton UP.

Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. Nature 447:1075–1080.

SUPPLEMENTAL DATA

Stability of intertemporal choice behavior and discount functions

In the main text, we quantified individual monkey temporal discounting behavior by examining intertemporal choice behavior averaged across the entire experiment, deriving a discount factor k for each monkey. Though previous experiments in animals and humans suggest that the discount function within an individual is generally stable (Richards et al., 1997; Simpson and Vuchinich, 2000), discounting behavior in our animals could in theory have been nonstationary (Schuck-Paim et al., 2004), either on a short time-scale (e.g. satiation over a day's session) or a long time-scale (e.g. behavioral changes over the course of the entire experiment).

To examine the short time-scale stability of discounting behavior, we divided behavioral choice data according to whether the experimental blocks occurred in the first half (Session-early) or second half (Session-late) of a given session. If the monkeys' valuation of delayed rewards changed substantially over the course of a session, one would expect differences in the observed choice behavior for the two groups of data. Figure S1a shows, individually by monkey, choice data and the best fit preference curves and discount functions for Session-early and Session-late data. As seen in the lower panels, the discount functions for early and late data do not significantly differ (Monkey W: $k_{early} = 0.174$ s^{-1}, $k_{late} = 0.151$ s^{-1}, $p = 0.31$; Monkey D: $k_{early} = 0.042$ s^{-1}, $k_{late} = 0.039$ s^{-1}, $p = 0.46$; permutation test, $n = 2000$).

To examine long time-scale stability, we divided behavioral choice data for each monkey according to whether they were recorded early (Experiment-early) or late (Experiment-late) in the overall course of the experiment. As seen in Figure S1b, discounting behavior did not change significantly between early and late sessions in the experiment (Monkey W: $k_{early} = 0.164$ s^{-1}, $k_{late} = 0.146$ s^{-1}, $p = 0.42$; Monkey D: $k_{early} = 0.037$ s^{-1}, $k_{late} = 0.044$ s^{-1}, $p = 0.06$; permutation test, $n = 2000$). These results suggest that, for each monkey, the average discount function derived from all behavioral data was an accurate estimate of delayed reward value for any given trial.

Characterization of the behavioral discount function

Different functional forms have been used to characterize the delay-dependent decline in subjective value observed during intertemporal choice behavior, notably the normative exponential function and the hyperbolic function initially described by Mazur. In the main text, both behavioral and neural discount functions were modeled as hyperbolic; we examine here how accurately these alternative functions characterize monkey intertemporal choice data. We fit each monkey's choice data to the single parameter equations describing exponential (1) or hyperbolic (2) discount functions:

$$SV = Ae^{-kD} \tag{S.1}$$

$$SV = \frac{A}{1+kD} \tag{S.2}$$

where A is the amount, D is the delay to reinforcement, and k is a parameter governing the rate of decrease in subjective value. The four preference curves and the discount function were simultaneously fit with a two-parameter binary logit model by maximum likelihood, with one parameter describing the discount function and one describing the logit function (Hosmer et al., 2000). To determine which discount function better describes the choice data, we compared the two models using Akaike's Information Criteria (AIC; Akaike, 1974). The AIC estimates the goodness of fit of a statistical model while accounting for model complexity; lower AIC values indicate a better model approximation to the underlying process. The AIC measures here show that the hyperbolic model is a better description than the exponential model in both monkeys (Monkey D: $\mathrm{AIC}_{\mathrm{hyp}} = 6688$, $\mathrm{AIC}_{\mathrm{exp}} = 6738$; Monkey W: $\mathrm{AIC}_{\mathrm{hyp}} = 4721$, $\mathrm{AIC}_{\mathrm{exp}} = 4811$); because both models have the same number of parameters, the lower hyperbolic AIC values strictly indicate a higher likelihood and a better characterization of the data. These results are consistent with the large majority of behavioral studies that report a better fit to discounting data using the hyperbolic function.

Fitting the behavioral discount function and four preference curves simultaneously with a two-parameter binary logit model provides two primary advantages: first, fitting the discount function to the choice data directly via MLE allows a direct comparison of different theoretical discount models (exponential versus hyperbolic) via the likelihood quantities, and second, the use of a single noise parameter β as opposed to four individual noise parameters represents a more parsimonious underlying choice model. However, an alternative stepwise method to fit the behavioral data exists, where the individual preference curves are first fit separately to identify indifference points, and the discount curve is subsequently fit to the identified indifference points. To ensure that the exact method of fitting the behavioral data does not affect our results, we also examined the behavioral choice data with this stepwise fit. Fit in this manner, the indifference points in both monkeys are well-described by hyperbolic discount functions (Monkey D: $k = 0.039\ \mathrm{s}^{-1}$, $R^2 = 0.98$; Monkey W: $k = 0.120\ \mathrm{s}^{-1}$, $R^2 = 0.99$). The discount rate parameter is statistically indistinguishable from that fit with the simultaneous binary logit model in one monkey (Monkey D: $k_{simultaneous} = 0.040\ \mathrm{s}^{-1}$, $k_{stepwise} = 0.039\ \mathrm{s}^{-1}$) and slightly smaller in the other animal (Monkey W: $k_{simultaneous} = 0.158\ \mathrm{s}^{-1}$, $k_{stepwise} = 0.120\ \mathrm{s}^{-1}$). Importantly, the temporal discount functions fit with this alternative method remain significantly different between the two animals ($p < 1.10\mathrm{x}10^{-16}$, permutation test) and still match the respective neural discount functions, falling well within the neural 95% confidence intervals (Monkey D: $k_{neural} = 0.038\ \mathrm{s}^{-1}$, 95% C.I. [0.010, 0.069]; Monkey W: $k_{neural} = 0.157\ \mathrm{s}^{-1}$, 95% C.I. [0.100, 0.236]). Furthermore, using the stepwise discounting fits does not change the primary conclusion of the paper: LIP firing rates are significantly modulated by subjective value ($R^2_{SV} = 0.146$, $p < 10^{-10}$, F-test) and not choice probability ($R^2_{CP} = 0.005$, $p = 0.23$, F-test).

Specificity of neurometric-psychometric match

Neural discount functions were obtained by fitting the hyperbolic discount model (S.2) to mean normalized firing rate data across the LIP population. As reported in the main text, the neural discount functions in both monkeys did not significantly differ from the corresponding behavioral discount functions during forced and free choice trials, suggesting a match between the behavioral and neural representations of subjective value. Here we examine the specificity of the neural discount functions to the individual monkeys. To compare the neural discount functions, we performed a permutation test using the combined neural data of both monkeys as a reference dataset. Two datasets corresponding to the size of the original datasets ($n = 48$, $n = 23$) were created by sampling without replacement and neural discount factors were calculated. This process was repeated 2000 times to generate a distribution of the differences in k values under the null hypothesis that the neural data from the two monkeys were generated by a single process. Compared to the distribution under the null hypothesis, the experimentally observed difference in neural discount functions is significant for both forced choice ($p = 0.0005$) and free choice trials ($p = 0.047$), suggesting that LIP activity represents individual-specific (in other words, idiosyncratic) delay-discounted value.

Temporal dependence of the neural discount function

In the main text, initial analysis of LIP activity focused on the interval 0-200 ms after peripheral target onset, as this corresponds to the peak transient visual response in the LIP population. The sliding window regression analysis in the main text (Fig. 6) indicates that LIP activity in later time periods is also dependent on subjective value, and thus should display similar neural activity (i.e. a corresponding neural discount function). To examine the robustness of the neurometric-psychometric correspondence seen in the main text (Fig. 3 and 4), we repeated the neural discount function analysis at multiple time intervals. Figure S2a and S2b displays the results of analyses on data from the window 150-350 ms after target onset (these are the same analyses shown in Figures 3 and 4 for 0-200 ms). This analysis reveals that there is nothing privileged about the 0-200 ms interval, as both forced choice and free choice neural discount functions in the 150-350 ms period also match their respective behavioral discount functions (Monkey D: $k_{behav} = 0.040\ s^{-1}$, $k_{neural\text{-}forced} = 0.040\ s^{-1}$, 95% C.I. [0.017, 0.075], $k_{neural\text{-}free} = 0.054\ s^{-1}$, 95% C.I. [0.023, 0.120]; Monkey W: $k_{behav} = 0.158\ s^{-1}$, $k_{neural\text{-}forced} = 0.125\ s^{-1}$, 95% C.I. [0.090, 0.201], $k_{neural\text{-}free} = 0.151\ s^{-1}$, 95% C.I. [0.077, 0.449]). We also analyzed the 150-350 ms data as in Figure 4 and find that, as for the 0-200 ms window, LIP activity is a much better explained by subjective value ($R^2 = 0.88$) rather than delay ($R^2 = 0.59$, Fig. S2b). Additionally, we repeated the neural discount function analysis for multiple non-overlapping 200 ms wide windows during early and middle target presentation (0-1000 ms). Plotted in Figure S2c are the free choice neural discount factors (k), separated by monkey, as a function of trial time. These data show that neural discount rates are generally robust to the exact temporal window used for analysis. Note that the time when both neural discount rates are closest to the respective behavioral discount rates (between 300-400 ms) corresponds to the peak of subjective value influence as determined by regression analysis in Figure 6.

Modulation of LIP activity by subjective value versus choice probability

To compare the relative influence of subjective value or choice probability on LIP activity, we employed univariate and multivariate linear regression to analyze block-averaged firing rate data. As reported in the main text, we found that LIP activity over much of the trial and particularly after target presentation was explained by delay-dependent subjective value. Importantly, when we compared the effect of subjective value and choice probability, this early-trial LIP modulation was independent of varying choice probability.

The results in the main text suggest that subjective value modulation is a general effect across the recorded LIP population. One potential concern is that the effects are driven by a subset of the data; the two individual monkeys demonstrated different behavioral discount functions (Monkey D: $k = 0.040\ s^{-1}$; Monkey W: $k = 0.158\ s^{-1}$), raising the possibility that the observed subjective value effects were driven by data from one monkey. To examine the effects by individual animal, we segregated the neuronal data by monkey and repeated the regression analyses reported in the main text. As shown in Figure S3, subjective value and not choice probability explains a significant portion of LIP modulation over much of the duration of intertemporal choice trials in both Monkey D and Monkey W, as indicated by R^2 values for univariate regressions on subjective value or choice probability performed in 100 ms windows across the duration of the trial. As in the main text, we examined activity in the 0-1000 ms interval following target presentation in further depth (Fig. S3, bottom panels). Univariate regression slopes were significant for subjective value in both monkeys (Monkey D: $b_{SV} = 0.657$, $p = 0.003$; Monkey W: $b_{SV} = 0.808$, $p = 1.9\text{x}10^{-7}$) and significant for choice probability in neither monkey (Monkey D: $b_{CP} = -0.019$, $p = 0.768$; Monkey W: $b_{CP} = 0.209$, $p = 0.074$). These results were confirmed with multiple regression analysis including both subjective value and choice probability as regressors (regression coefficients; Monkey D: $b_{SV} = 0.859$, 95% C.I. [0.399-1.320], $b_{CP} = -0.117$, 95% C.I. [-0.251-0.018]; Monkey W: $b_{SV} = 1.038$, 95% C.I. [0.680-1.396], $b_{CP} = -0.228$, 95% C.I. [-0.478-0.022]). Together, these results suggest that subjective value and not choice probability drives LIP modulation early in decision-making, regardless of the individual delay discounting functions.

In the main text, we analyzed LIP data grouped across all delayed reward magnitude conditions under which neural data was collected (0.143, 0.163, 0.260 ml). Here we examine subjective value modulation of LIP activity when the different magnitude conditions are examined separately. Figure S4 shows block-averaged normalized firing rates in the 0-1000 ms period after target onset, as a function of either subjective value (left) or choice probability (right). As demonstrated by the univariate regression lines (red), subjective value modulates firing rate to a much greater extent than choice probability. Subjective value explains a higher proportion of the variance across the population in all three magnitude conditions (mag. 0.143 ml: $R^2_{SV} = 0.202$, $R^2_{CP} = 0.035$; mag. 0.163 ml: $R^2_{SV} = 0.054$, $R^2_{CP} = 0.0002$; mag. 0.260 ml: $R^2_{SV} = 0.136$, $R^2_{CP} = 0.021$), and two of the three subjective value regression models are significant (F-test, $p < 0.05$), compared to none of the choice probability models.

Finally, because the primary analyses relied on linear regression analyses, we tested the validity of these models by residuals analysis, emphasizing the early visual period (0-1000 ms after target onset). Figure S5 shows residual data from the subjective value and choice probability univariate regressions. As indicated by the histogram of residuals (top) and normal probability plots (bottom), the residuals follow an approximately normal distribution. Both the residual distributions indicated are not significantly different from a normal distribution (SV: $p = 0.338$; CP: $p = 0.205$; Lilliefors test). Additional analyses indicate that the distributions of residuals are homoscedastic and unrelated to either the independent variables (subjective value, choice probability) or the dependent variable (normalized firing rate).

Together, these results indicate that LIP activity – in the early and middle stages of the decision process - represents subjective value and not choice probability. While it might be expected that choice probability and value would be strongly correlated, these findings suggest that the stochastic choice behavior responsible for the observed choice probability can be quite independent of value (for example, at the extremes of the preference curves). A key reason for this independence is that choice behavior is necessarily a function of the values of all options under consideration in a decision: in a given block, a monkey's choice probability depends on the values of both the delayed and immediate choices. Thus the subjective value of a single option (the RF target) can vary widely without a (significant) change in the associated choice probability (because it is at a maximum or a minimum), if the value of the other option remains much lower (or higher). This relationship is explicitly defined in the logit function (see Materials and Methods), where choice probability is acknowledged to be a joint function of both option values. Thus, in the context of a decision process, *values* are inputs and *choice* (and thus on average, choice probability) represents an output; the nonlinear nature of this relationship, as well as its dependence on multiple option values, allows for a perhaps surprising degree of separation between subjective value and choice probability under these conditions.

Effect of saccadic tuning strength on subjective value and choice probability coding

We selected LIP neurons for neurophysiological recording based solely on spatially selective responses to visual stimulus presentation. However, LIP neurons are known to exhibit varying degrees of spatially modulated peri-saccadic activity, which could affect the level of choice probability modulation observable early in the trial. In particular, the predominance of value and not choice-related modulation observed in the main text could be explained if the population of recorded LIP neurons lack strong spatial selectivity during choice.

To quantify saccadic spatial selectivity for each neuron, we compared peri-saccadic activity (from 200 ms before to 200 ms after saccade initiation) between trials with saccades into the RF (FR_{in}) to those with saccades away from the RF (FR_{out}), using a spatial selectivity index (SI):

$$SI = \frac{FR_{in} - FR_{out}}{FR_{in} + FR_{out}} \tag{S.3}$$

Possible SI values range from -1 to 1, with positive values indicating higher activity for saccades into the response field and values near zero indicating no directional difference in firing rates and thus no spatial selectivity. As shown in Figure S6a, we found significantly positive saccadic spatial tuning in both animals (Monkey D: mean $SI = 0.14$, $p = 2.97\text{x}10^{-5}$; Monkey W: mean $SI = 0.17$, $p = 4.14\text{x}10^{-6}$; t-test), with the large majority of neurons showing positive SI values (59/71 neurons, 83.1%); there was no difference in the saccadic tuning index between animals ($p = 0.52$, t-test). Thus, LIP neurons in this task were in fact spatially selective at the time of saccade, suggesting that a lack of spatially-selective intentional coding does not explain the failure to observe choice probability signals early in the decision process.

To further ensure that the absence of significant choice probability coding in our primary results (main text and Figure 6) was not a result of low spatial selectivity, we repeated the subjective value versus choice probability analyses on a subset of LIP neurons with strong saccadic spatial tuning. We identified LIP neurons with SI values larger than the median SI value (0.125), and examined the modulation of this subpopulation of tuned neurons using simple linear regression. As in Figure 6 of the main text, Figure S6b plots average firing rate in the 0-1000 ms interval immediately following target presentation as a function of subjective value (left) or choice probability (right). Even in this subpopulation of LIP neurons with high spatially selective saccade activity, firing rates are still modulated by subjective value and not choice probability ($R^2_{SV} = 0.175$, $p = 3.8\text{x}10^{-7}$; $R^2_{CP} = 0.001$, $p = 0.68$; F-test).

As an additional measure of saccadic spatial selectivity, we conducted a univariate regression across all neural data at the level of individual trials to examine the influence of saccade direction on LIP firing rate (normalized). This regression was conducted separately for nonoverlapping 100 ms intervals across the duration of the trial. Figure S7 plots the coefficient of determination (R^2) values of these regressions for saccade choice (black), as well as those for univariate regressions of LIP activity on either subjective value (blue) or choice probability (red); because the primary purpose of the supplemental analysis was to examine the relative effect of choice versus choice probability in LIP activity, significance levels were not included in the supplemental figure. There is a strong influence of saccade choice on LIP activity that peaks at the time of saccade initiation, confirming significant spatial selectivity in the activity of the recorded LIP population. Note that the addition of trial-to-trial variability in firing rate to the regression, while subjective value and choice probability remain constant across all trials in a block, necessarily reduces the R^2 values compared to the block-wise regression in Figure 6. Presumably, time-varying differences in this additional variance in firing rate leads to the shift in the peak of the subjective value effect, relative to the block-wise regression; for example, lower trial-to-trial variability early in the trial would decrease the early R^2 values to a lesser extent, shifting the peak to earlier in the trial. Subjective value remains the primary determinant of LIP activity early in the trial. Additionally, note that at no time in the trial does choice probability provide more explanatory power than the actual saccade choice. Together, these data suggest that LIP transitions from a coding of subjective value into a coding of saccade direction, without an intervening, separate representation of average choice probability.

Intertemporal choice reaction times

In the main text, we quantified behavioral discount functions which demonstrate that subjective value is a consistent, predictable function of delay. However, it is possible that other processes – such as motivation or attention – also vary with delay, and that it is these processes that directly modulate LIP activity (Goldberg et al., 2002; Bisley and Goldberg, 2003). To address this possibility, we examined reaction times in the intertemporal choice task (Posner, 1980). We quantified reaction time in each trial as the duration between the offset of the central fixation target, which signaled the earliest time the monkey was permitted to initiate a saccade, and the actual initiation of saccade. In particular, we examined forced choice trials requiring a saccade towards the delayed reward target; in forced choice trials, the randomized presentation of instruction cues prevents the monkey from deciding a saccade direction early in the trial, increasing the ability to detect any potential reaction time effect. We focus on trials with instructed saccades towards the RF target to determine if there is a delay-dependent allocation of attention. If diminishing attentional allocation to the delayed reward target with increasing delay were responsible for the observed LIP neural activity, reaction times should increase consistently as a function of delay. As shown in Figure S8, reaction times in both monkeys are non-monotonic (inverted U-shaped) functions of delay, with the longest values generally occurring at intermediate delays.

References

Akaike H (1974) New look at statistical-model identification. *Ieee Transactions on Automatic Control* AC19:716-723.

Bisley JW, Goldberg ME (2003) Neuronal activity in the lateral intraparietal area and spatial attention. Science 299:81-86.

Goldberg ME, Bisley J, Powell KD, Gottlieb J, Kusunoki M (2002) The role of the lateral intraparietal area of the monkey in the generation of saccades and visuospatial attention. Ann N Y Acad Sci 956:205-215.

Hosmer DW, Lemeshow S, Cook ED (2000) Applied logistic regression. New York: Wiley.

Posner MI (1980) Orienting of attention. Q J Exp Psychol 32:3-25.

Richards JB, Mitchell SH, deWit H, Seiden LS (1997) Determination of discount functions in rats with an adjusting-amount procedure. J Exp Anal Behav 67:353-366.

Schuck-Paim C, Pompilio L, Kacelnik A (2004) State-dependent decisions cause apparent violations of rationality in animal choice. PLoS Biol 2:e402.

Simpson CA, Vuchinich RE (2000) Reliability of a measure of temporal discounting. Psychol Rec 50:3 16.

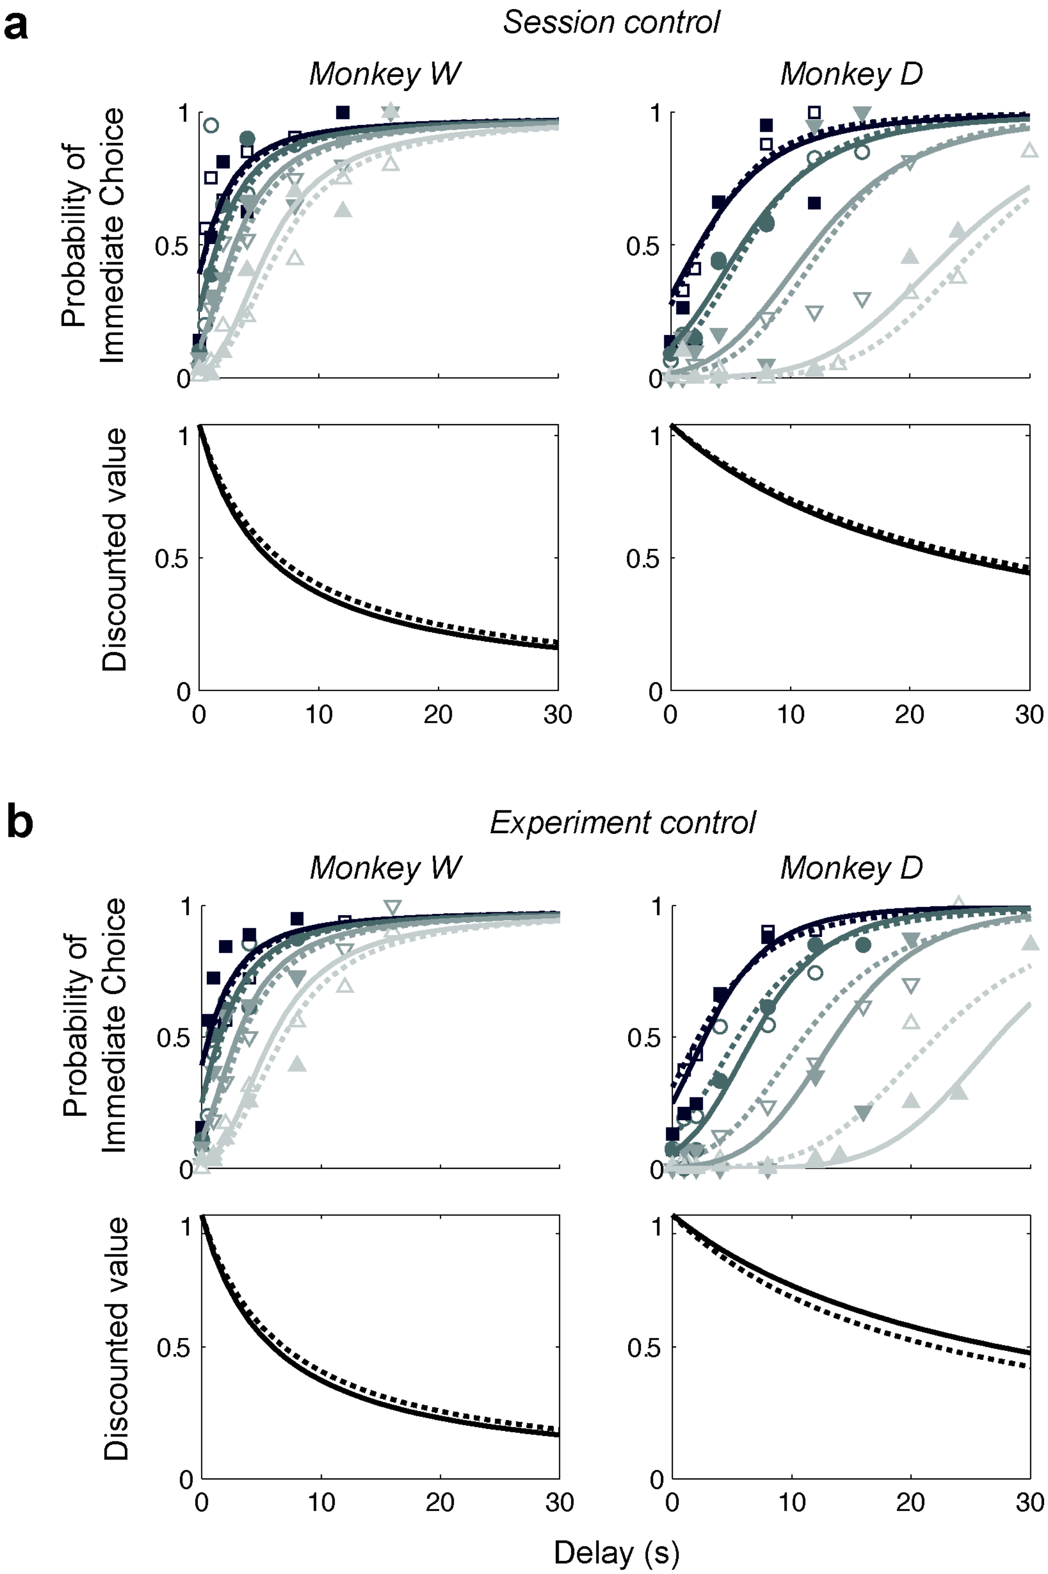

Figure S1. Stability of intertemporal choice and behavioral discount functions. (a) Choice data and best fit choice curves (top) and discount functions (bottom), segregated by time of occurence in a session (solid symbols and lines, early; open symbols and dashed lines, late). Color and symbols as used in main text and Figure 1. Early and late discount curves did not significantly differ in both monkeys (*p* = 0.31, *p* = 0.46, permutation test), suggesting stationary discounting behavior within recording sessions. **(b)** Choice data, choice curves, and discount functions segregated by time of occurrence across the duration of the experiment. Early and late discount curves did not significantly differ (*p* = 0.42, *p* = 0.06, permutation test), suggesting that the discount function was also stable between sessions and across the experiment.

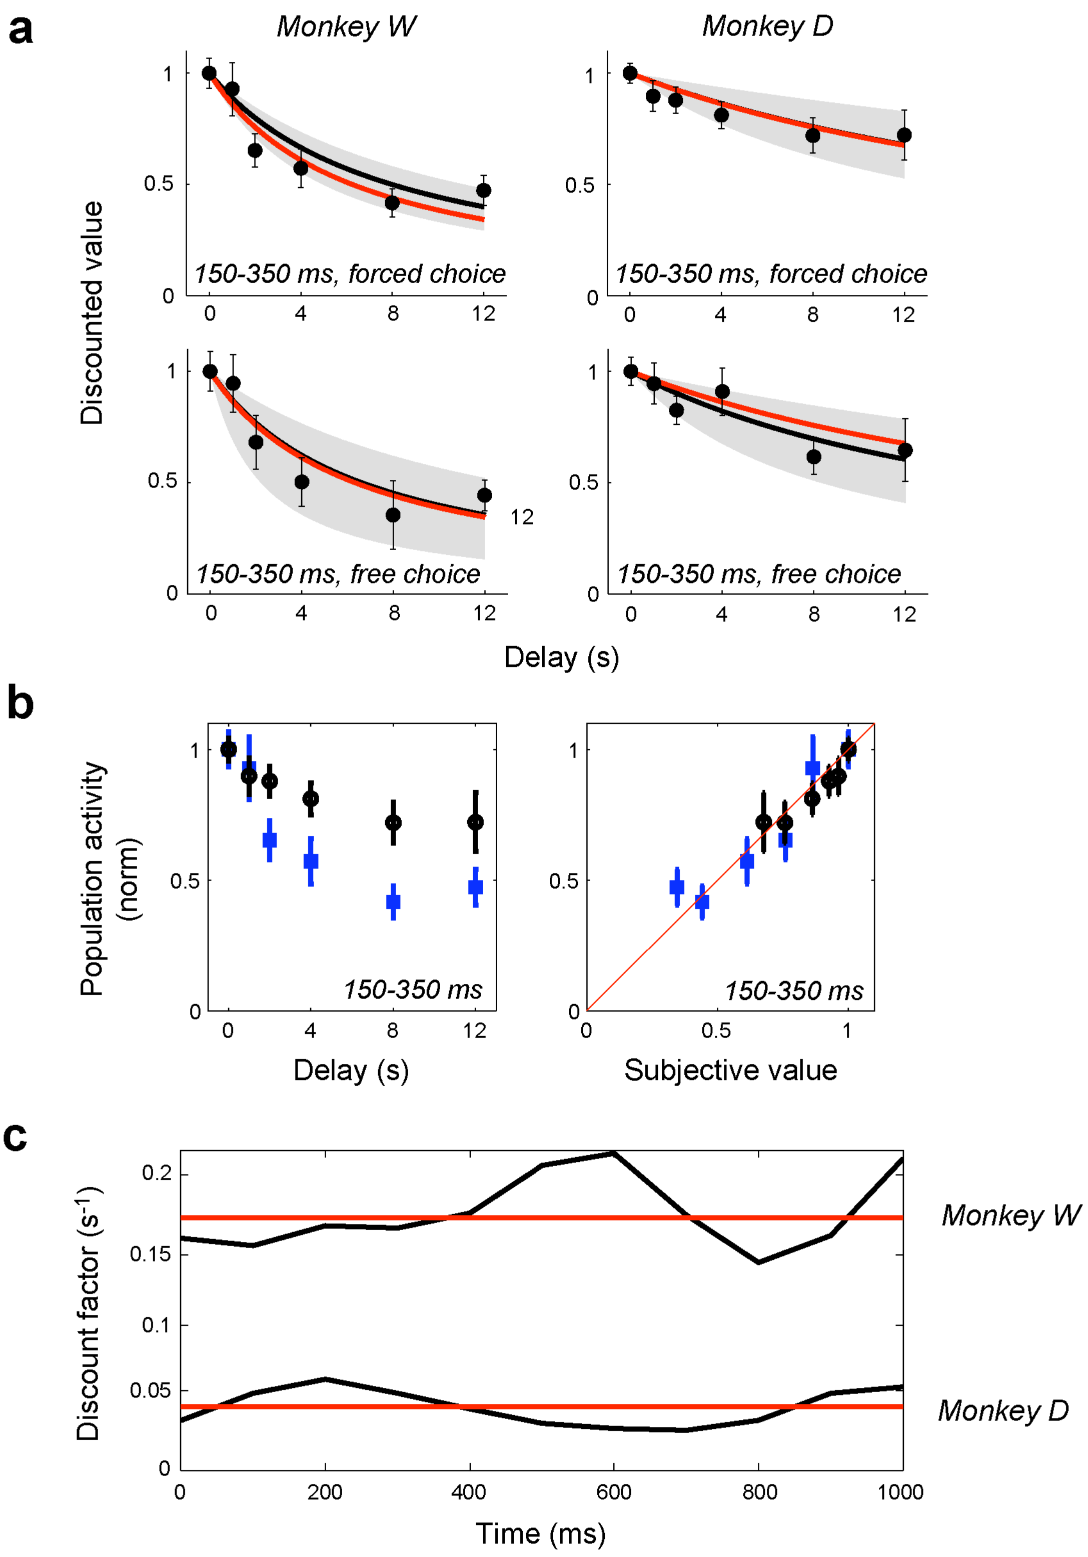

Figure S2. Temporal dependence of neural discount function. (a) Forced choice (top) and free choice (bottom) neural discount functions for the interval 150-350 ms after peripheral target onset. Mean (±s.e.m.) normalized LIP activity is shown as a function of delay, relative to activity at zero delay. As in Figure 3, behavioral discount functions are shown in red and neural discount functions in black, with 95% bootstrap confidence intervals in gray. **(b)** Normalized population activity in the interval 150-350 ms after target onset, as a function of delay (left) and subjective value (right). **(c)** Free choice discount factors as a function of time in trial, in non-overlapping 200 ms windows. Behavioral discount rates shown in red. These data show that the correspondence between LIP ativity and the behavioral discount function is not limited to the 0-200 ms window used in the main text.

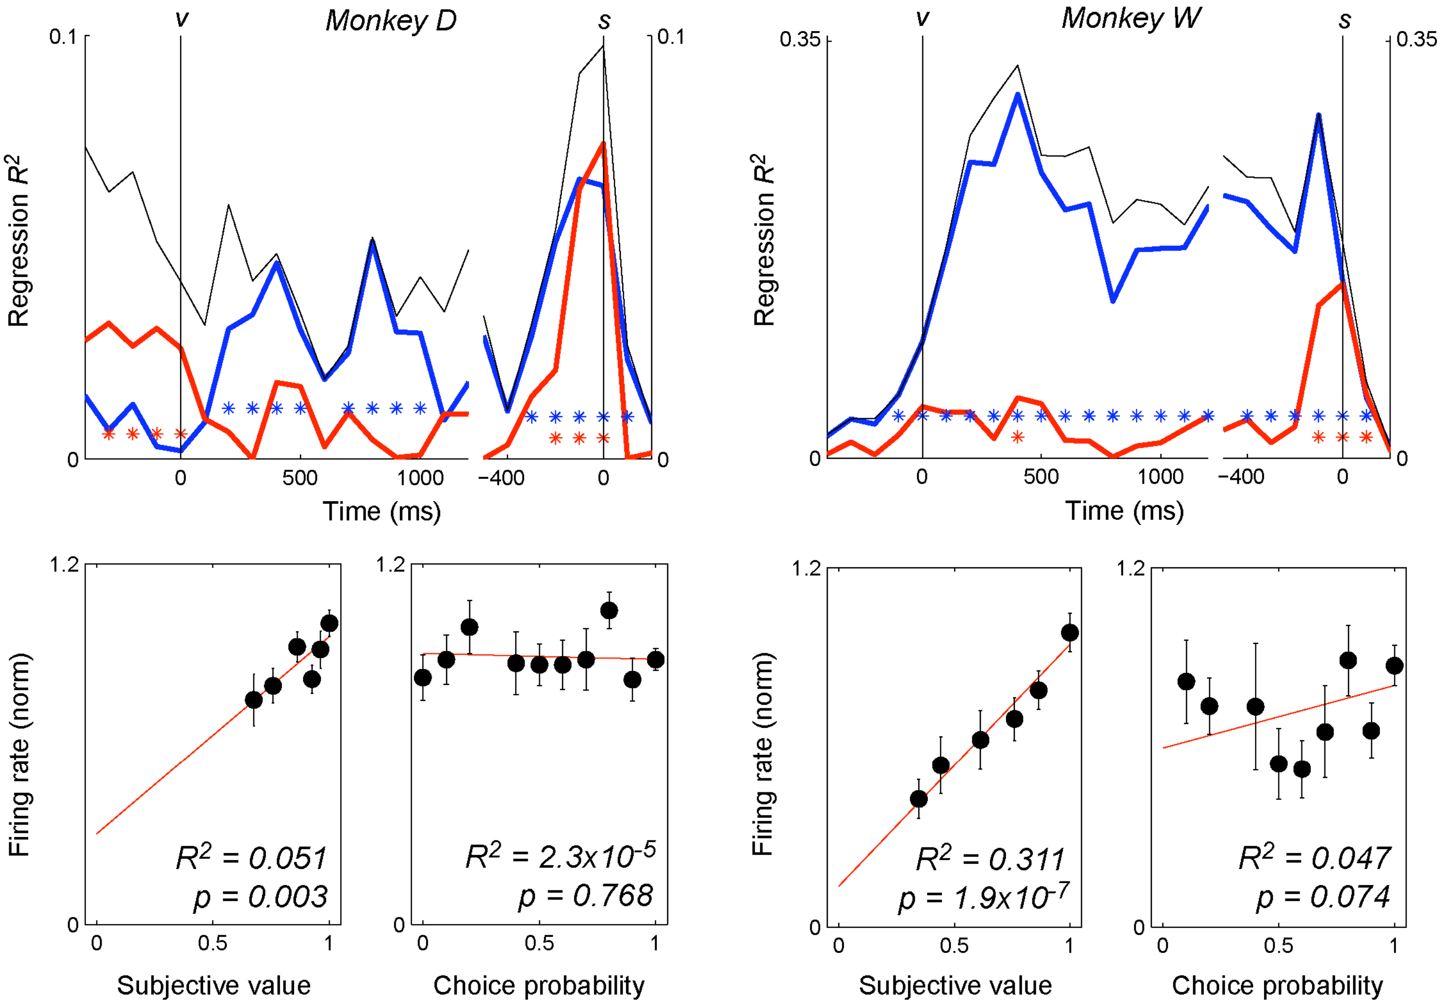

Figure S3. Subjective value versus choice probability in individual animals. Top panels, coefficient of variation (R^2) values for univariate regressions of block neural activity on subjective value (blue) or choice probability (red), separated by individual monkey. For comparision, R^2 values for multiple regression analyses including both subjective value and choice probability as covariates are also shown (black). Regressions were performed in nonoverlapping 100 ms windows across the length of the trial, with activity either aligned to target onset (*v*) or initiation of saccade (*s*). Asterisks indicate significant R^2 values for the corresponding univariate regressions ($p < 0.05$, F test). Bottom panels, neural activity during target presentation (0-1000 ms after target onset) as a function of either subjective value or choice probability, separated by individual monkey. LIP activity is significantly related to subjective and not choice probability, as evaluated by univariate regression slopes (red; Monkey D: $b_{SV} = 0.657$, $p = 0.003$; $b_{CP} = -0.019$, $p = 0.768$; Monkey W: $b_{SV} = 0.808$, $p = 1.9\text{x}10^{-7}$; $b_{CP} = 0.209$, $p = 0.074$). Regression analyses were performed on all relevant block data points and binned for illustration purposes.

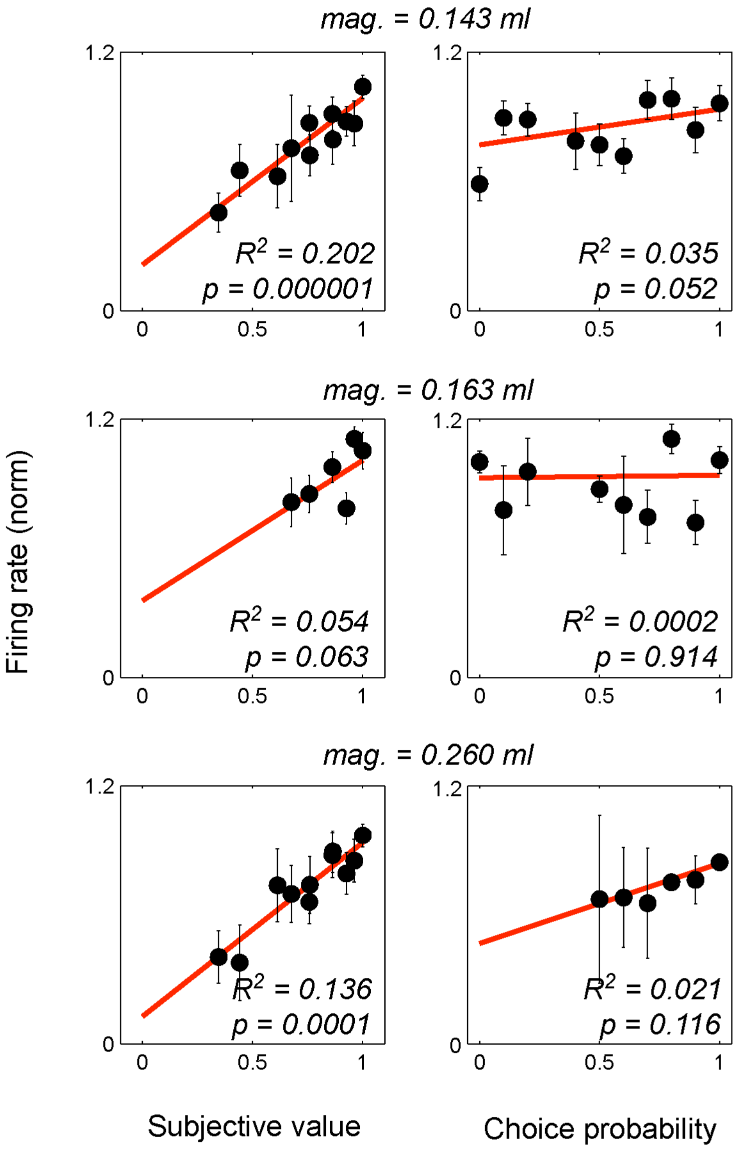

Figure S4. Subjective value effects segregated by reward magnitude condition. LIP activity plotted as a function of either subjective value (left) or choice probability (right) for three different delayed reward magnitudes (0.143 ml, 0.163 ml, 0.260 ml). Neural data corresponds to the 0-1000 ms window after target onset, averaged over each block of neural data. Regression slopes (red) and *p*-values for univariate linear regressions are indicated for each magnitude condition. Regression analyses were performed on all relevant block data points and only binned for illustration purposes.

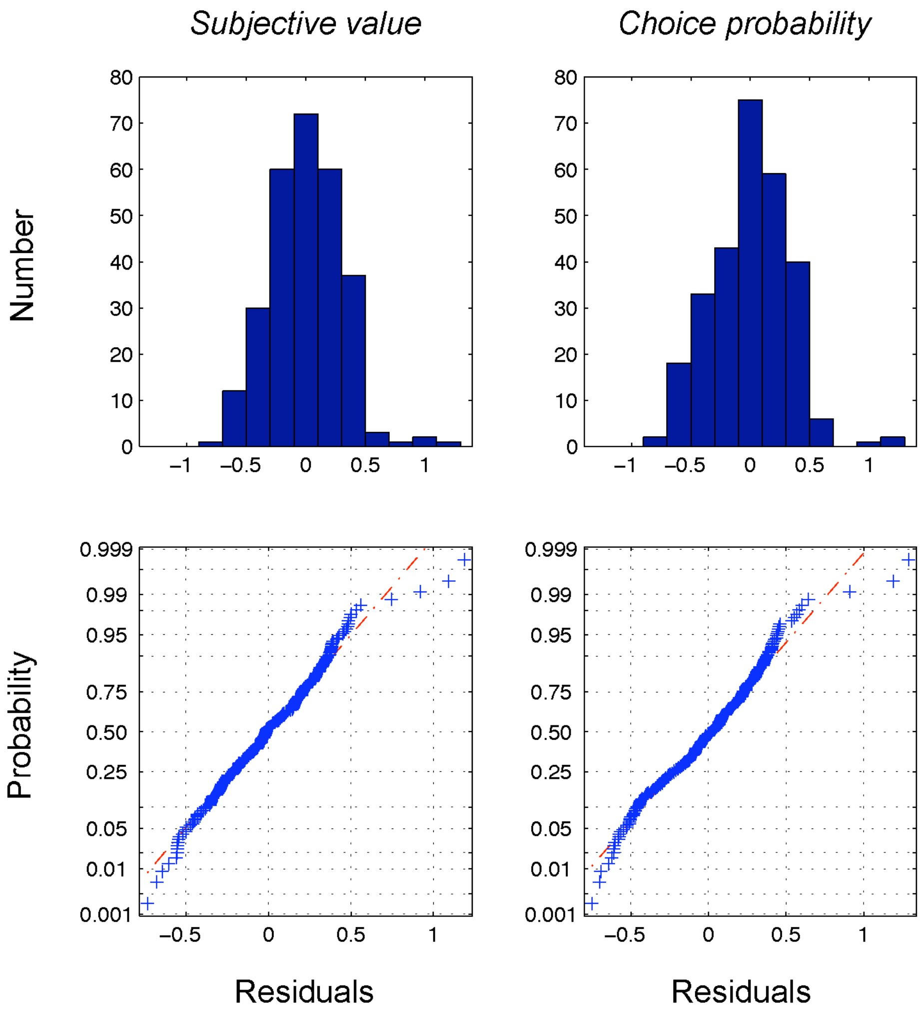

Figure S5. Residuals analysis: validity of subjective value and choice probability regression analyses. Univariate regression residuals for subjective value (left) and choice probability (right). Top row plots the histogram of observed residuals in the individual regressions. Bottom row plots residuals as normal probability plots. The distribution of residuals in both cases is not signficantly different from a normal distribution (subjective value: $p = 0.338$; choice probability, $p = 0.205$; Lilliefors test).

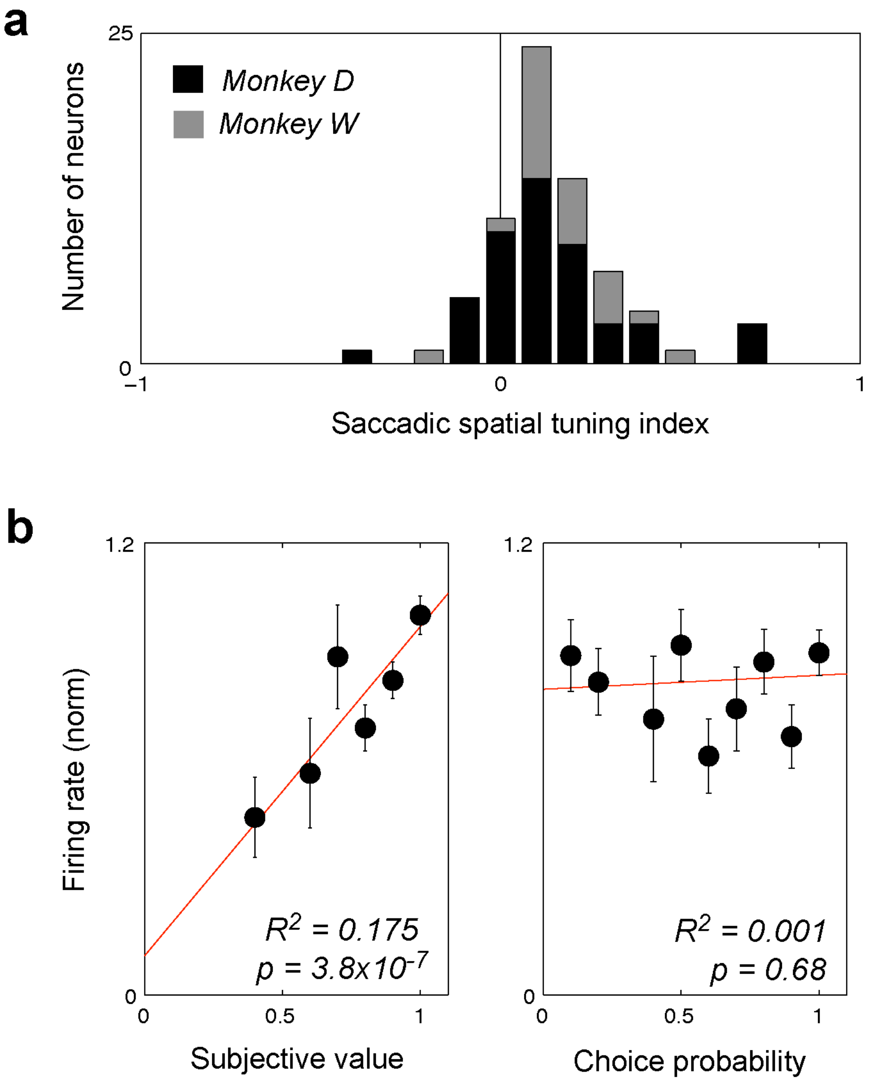

Figure S6. Effect of saccadic spatial selectivity on value and choice probability coding. (a) Distribution of saccadic spatial selectivity indices (SI) across the LIP population. The index is a measure of peri-saccadic spatial tuning, quantified as the difference in firing rate between saccades into and away from the response field, divided by the sum of the firing rates. Positive values of the index indicate activity at the time of saccade that is higher in response to an eye movement towards the visual-stimulus mapped response field; index values near zero indicate no spatial selectivity. SI values are significantly positive for each monkey individually as well as the entire population of neurons. **(b)** Spatially selective neurons are modulated by value and not choice probability. Mean LIP activity as a function of subjective value (left) or choice probability (right) for neurons with SI values greater than the median index value (0.125). Analyses are analogous to those in Figure 6 of the main text. Regression lines (red) are significant for subjective value (R^2 = 0.175, p = 3.8x10^{-7}, F-test) but not choice probability (R^2 = 0.001, p = 0.68, F-test).

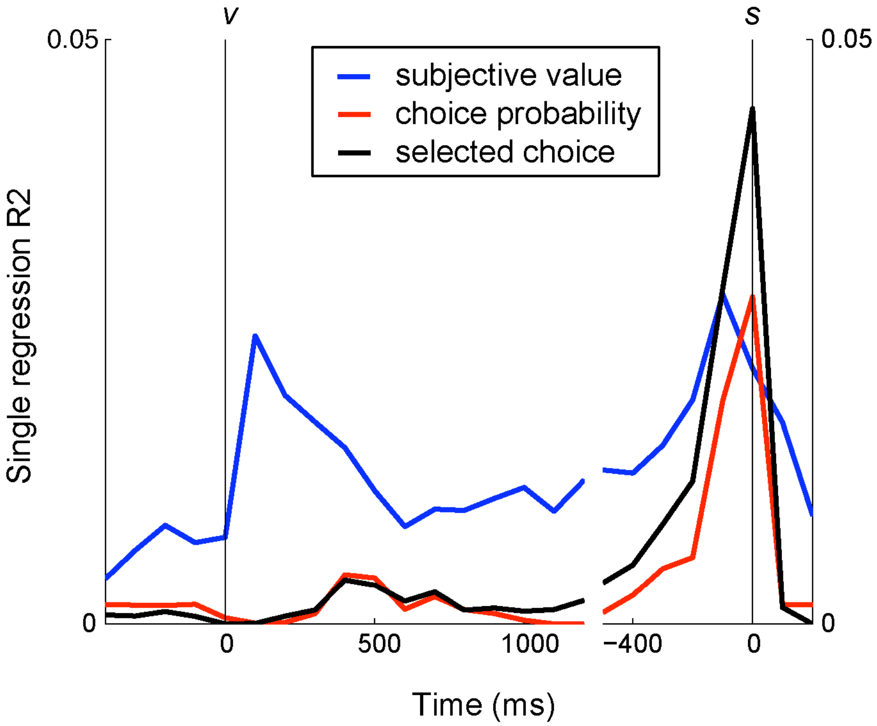

Figure S7. Relative strength of subjective value, choice probability and selected choice coding. Lines plot coefficient of variation (R^2) for individual univariate regressions of LIP activity on either subjective value (blue), choice probability (red), or the selected saccade (black) in free choice trials. Normalized firing rate data from every individual trial was included in the regression analyses, which were performed in nonoverlapping 100 ms windows across the length of the trial and aligned to either target onset (*v*) or saccade (*s*). Note that choice probability provides no additional explanatory power beyond that provided by the saccadic choice.

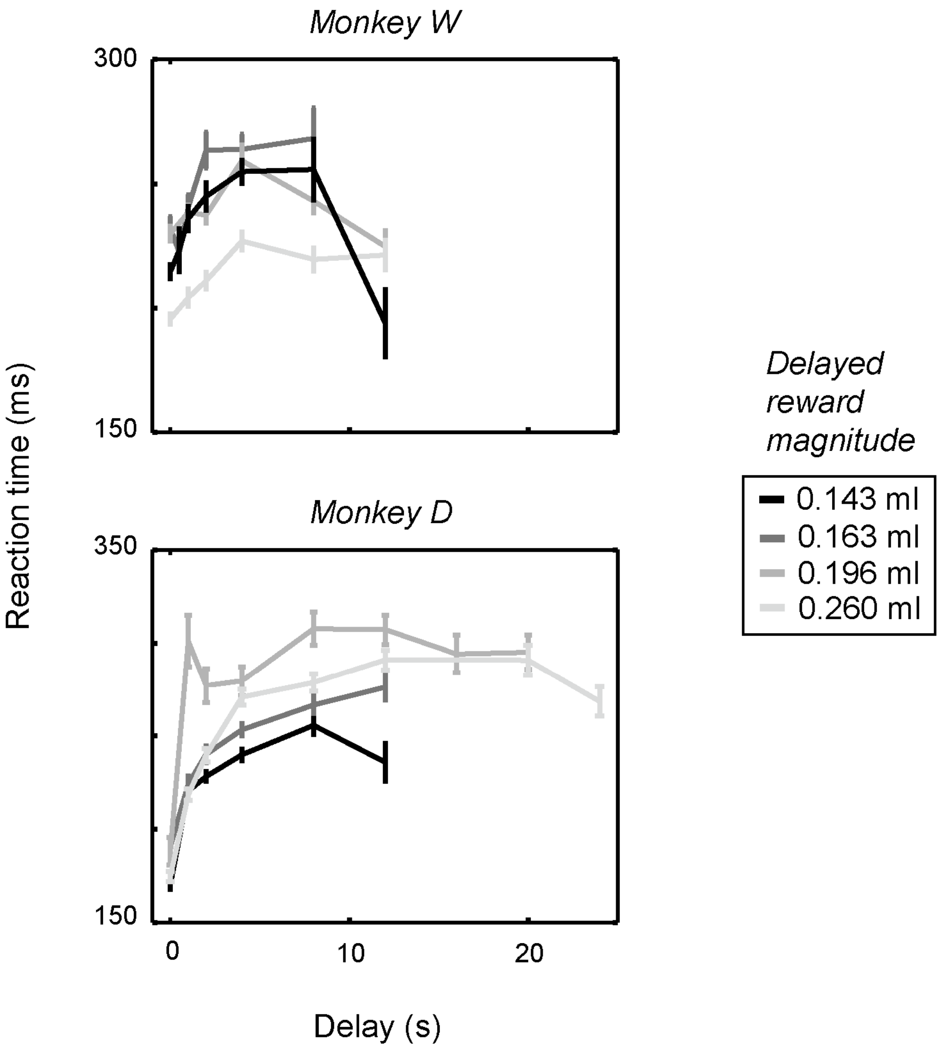

Figure S8. Reaction time is not a monotonic function of delay. Mean ± s.e.m. reaction time as a function of delay to reward for each individual monkey. Data shown corresponds to forced choice trials where the monkey was instructed to complete a saccade towards the delayed reward target. Individual lines correspond to different delayed reward magnitudes.