

Bounded Integration in Parietal Cortex Underlies Decisions Even When Viewing Duration Is Dictated by the Environment

Roозbeh Kiani, Timothy D. Hanks, and Michael N. Shadlen

Howard Hughes Medical Institute, National Primate Research Center, and Department of Physiology and Biophysics, University of Washington, Seattle, Washington 98195

Decisions about sensory stimuli are often based on an accumulation of evidence in time. When subjects control stimulus duration, the decision terminates when the accumulated evidence reaches a criterion level. Under many natural circumstances and in many laboratory settings, the environment, rather than the subject, controls the stimulus duration. In these settings, it is generally assumed that subjects commit to a choice at the end of the stimulus stream. Indeed, failure to benefit from the full stream of information is interpreted as a sign of imperfect accumulation or memory leak. Contrary to these assumptions, we show that monkeys performing a direction discrimination task commit to a choice when the accumulated evidence reaches a threshold level (or bound), sometimes long before the end of stimulus. This bounded accumulation of evidence is reflected in the activity of neurons in the lateral intraparietal cortex. Thus, the readout of visual cortex embraces a termination rule to limit processing even when potentially useful information is available.

Key words: lateral intraparietal area; LIP; perceptual decision; bounded accumulation; leaky accumulation; direction discrimination; electrophysiology; psychophysics; vision

Introduction

A hallmark of cognition is the ability to choose among alternatives based on evidence bearing on the state of the world and the expected consequences of any particular choice. A neural correlate of simple perceptual decisions, involving a choice between two actions based on the stimulus identity (Fechner, 1860; Wundt, 1865; Green and Swets, 1966; Link, 1992; Gold and Shadlen, 2007), has been demonstrated in sensory-motor association areas that are concerned with the preparation of the actions (Platt and Glimcher, 1997; Schall and Thompson, 1999; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Glimcher, 2003; Romo and Salinas, 2003; de Lafuente and Romo, 2005; Gold and Shadlen, 2007). For example, in monkeys trained to judge the motion direction of dynamic random dots and to report their decision with a saccade, neurons in several visuomotor integration areas, including lateral intraparietal cortex (LIP), represent the transformation of the motion information into a decision variable that directs the saccade (Kim and Shadlen, 1999; Gold and Shadlen, 2000, 2003; Horwitz and Newsome, 2001; Shadlen and Newsome, 2001).

In many instances, a decision between alternatives must also incorporate a rule for terminating the decision process, in effect stopping the deliberation and committing to a choice. A large body of theoretical and experimental work supports the idea that a single mechanism accounts for both the termination of the decision process and the ensuing choice. In a reaction-time (RT) version of the motion direction discrimination task, subjects are both faster and more accurate for stronger motion stimuli (Link, 1992; Roitman and Shadlen, 2002; Ratcliff and Smith, 2004). An accumulation-to-bound mechanism, in which momentary sensory evidence (furnished by direction-selective neurons) is accumulated over time toward a criterion level or bound, explains the choice and RT functions (Mazurek et al., 2003; Lo and Wang, 2006; Shadlen et al., 2006). Consistent with the accumulation-to-bound mechanism, LIP neurons increase or decrease their firing rates as evidence mounts for or against the direction associated with the choice target in the response field (RF) of each cell. Furthermore, just before the saccade, LIP neurons attain the same level of activity, independent of motion strength. Hence, in the RT task, where the stimulus viewing duration is dictated by the subject, it is hypothesized that LIP marks both the beginning and end of the decision process (Roitman and Shadlen, 2002; Gold and Shadlen, 2007).

In other instances, however, subjects must base their decisions on limited evidence, the availability of which is controlled by the environment; or, they may be required to postpone their behavioral response until it is solicited (see Fig. 1). Indeed, outside RT studies, such cases include most of the behavioral tasks that have been investigated by psychophysicists and neurophysiologists

Received Oct. 20, 2007; revised Jan. 4, 2008; accepted Jan. 26, 2008.

This work was supported by the Howard Hughes Medical Institute (HHMI), National Eye Institute Grant EY11378, and National Center for Research Resources Grant RR00166. T.D.H. is also supported by an HHMI predoctoral fellowship. We thank A. Churchland and T. Yang for comments on this manuscript and useful suggestions, and M. Mihali, V. K. Skypceck, and A. D. Boulet for technical assistance.

Correspondence should be addressed to Michael N. Shadlen, Howard Hughes Medical Institute, National Primate Research Center, and Department of Physiology and Biophysics, University of Washington, Seattle, WA 98195. E-mail: shadlen@u.washington.edu.

DOI:10.1523/JNEUROSCI.4761-07.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/283017-13\$15.00/0

(Rieke et al., 1997; Parker and Newsome, 1998). The division between the RT and fixed duration tasks is convenient for experimental purposes, and it can be exploited to ask an important question in sensory physiology: what is the termination rule for perceptual decisions in general? An idea that arose from RT studies is that there is an activity threshold, which terminates the decision process. In this paper, we consider the hypothesis that the beginning and end of the decision process are governed by the same neural mechanisms whether the animal or the environment determines the viewing time.

This idea poses two critical challenges. First, the representation of accumulated evidence in LIP is delayed by at least 100 ms compared with the activity of direction-selective neurons in the middle temporal area (MT) (Britten et al., 1996; Bair et al., 2002; Roitman and Shadlen, 2002; Osborne et al., 2004; Huk and Shadlen, 2005). When stimuli are brief, it is not known whether this delayed representation of accumulated evidence can contribute to the formation of decisions. A critical step for establishing a role for LIP in such decisions is showing that decisions follow the course of neural activity in LIP, respecting the delayed neural representation. Second, when a behavioral response is solicited by an exogenous cue or instruction, there can be a discrepancy between the time of response and commitment to a choice. It is not known when and how subjects commit to a choice in such a situation.

In this study, we show that both of these challenges are met in a simple perceptual decision about motion direction. Hence, the accumulation-to-bound mechanism generalizes beyond simple RT tasks. Additionally, we show that termination of a decision process, while the stimulus remains visible, delimits the improvement in performance that might otherwise accrue with prolonged stimulus exposures (Burr and Santoro, 2001). The finding overturns a long-held assumption that such failure is a sign of inefficient (leaky) integration (Grice, 1972; McClelland, 1979, 1993; Busemeyer and Townsend, 1993; Diederich, 1995, 1997; Kim and Myung, 1995; Smith, 1995; Usher and McClelland, 2001; Ludwig et al., 2005; Busemeyer et al., 2006). Indeed, many ideal observer analyses (e.g., signal detection theory, information theory), which have been applied to perception and neurophysiology, presume that observers use all available stimulus information. We show instead that the decision process may terminate when the accumulated information reaches a critical level, even when subjects are required to view the whole stimulus.

Materials and Methods

The data set consists of 51 neurons in the LIP of two rhesus monkeys (both male; 9 and 11.5 kg) trained on a motion direction discrimination task with variable stimulus durations. All training, surgery, and recording procedures conformed to the National Institutes of Health *Guide for the Care and Use of Laboratory Animals* and were approved by the University of Washington Animal Care Committee.

Data acquisition and neuron selection. Each monkey was implanted with a headpost, a scleral eye coil for eye monitoring (Fuchs and Robinson, 1966; Judge et al., 1980), and a recording cylinder (Crist Instruments, Damascus, MD) suitable for magnetic resonance imaging (MRI). The recording cylinder was positioned above the intraparietal sulcus in the left hemisphere of one monkey and the right hemisphere of the other. A plastic grid (1 mm spacing; Crist Instruments) was used for precise placement of the recording microelectrodes ($\sim 1\text{ M}\Omega$ impedance at 1 kHz; Alpha Omega, Nazareth, Israel). The electrodes were advanced through a sterile stainless steel guide tube, which penetrated the dura and was held by the grid. Postoperative MRI (1.5 T scanner) was used to direct the placement of the electrodes within the recording chamber. Signals from the electrode were amplified and filtered before sorting.

Action potentials of individual neurons were sorted either by using a dual voltage-time window discriminator (Bak Electronics, Germantown, MD) or by template matching (Alpha Omega). The time of action potentials, trial events, and horizontal/vertical eye positions were stored at 1 kHz for analysis.

We targeted neurons in the ventral division of area LIP (LIPv) (Lewis and Van Essen, 2000) by registering our MRI scans with a high-resolution “flattened” scan distributed with the CARET software package (Van Essen, 2002). Neurons in LIPv tend to exhibit sustained activity in delayed and memory guided eye movement tasks (Bracewell et al., 1996). We screened 59 neurons with this property (memory-delay period, 600–1500 ms). Of these, all but eight exhibited similar patterns of response during the motion task. We excluded these neurons from the analysis, because they did not exhibit a clear stereotyped dip after onset of random dot motion, or they were uninformative during decision making. During data collection, we also used a visually guided delayed saccade task to targets in and out the RF of the neuron. These trials were identical to the memory-saccade task, except the target remained visible until the end of the trial (a saccade to the target).

Motion direction discrimination task. Each trial began with the appearance of a fixation point (FP) (0.3° diameter) at the center of the display monitor. The monkey was required to maintain its gaze to within $\pm 1.5^\circ$ of the FP so long as it was visible on the computer monitor. After 200–600 ms of stable fixation, two red targets appeared on opposite sides of the screen. During recording sessions, one of the targets was located in the response field of the recorded LIP cell. After a random delay (250–600 ms), the dynamic random dot stimulus appeared within a 5° circular aperture centered on the FP. The motion stimulus is described in detail previously (Roitman and Shadlen, 2002; Gold and Shadlen, 2003). The percentage of coherently moving dots in each trial was chosen randomly from the following values: 0, 6.4, 12.8, 25.6, and 51.2%. The net direction of motion was toward one of the two targets. The stimulus duration was sampled randomly from a truncated exponential distribution (range, 80–1500 ms; mean, 330 ms). The flat hazard rate associated with this distribution of random durations minimizes the monkey's ability to anticipate the end of the trial (Ghose and Maunsell, 2002; Janssen and Shadlen, 2005).

After the termination of the motion stimulus, the FP was extinguished. This event (the Go signal) instructed the monkey to initiate an eye movement to one of the choice targets to indicate its decision about the direction of motion. In half of the trials, the Go signal coincided with offset of the motion stimulus. In the remaining trials, a random delay (500–1000 ms; approximately exponential) period intervened between motion offset and the Go signal. All trial types were randomly interleaved. The saccadic eye movement had to be initiated in <500 ms from the Go signal. The monkey received a liquid reward for each correct response and in half of the 0% coherence trials. Errors were followed by 1 s extension of the 1.5 s intertrial interval. The monkeys chose one of the two targets in almost all of the trials (99.1%). Analyses are based on 42,647 (physiology and psychophysics) plus 60,921 (psychophysics) trials on this task.

In the motion pulse experiment, the motion coherence was increased or decreased by 3.2% for 200 ms in every trial. For 0% coherence trials, this change of coherence resulted in a weak net motion toward one of the two targets during the pulse. To accommodate one pulse on every trial, the minimum stimulus duration was increased to 280 ms (truncated exponential distribution; range, 280–1500 ms; mean, 650 ms). The pulse began at a random time, with uniform probability density, between the beginning of the stimulus and 200 ms before the end of the stimulus. The pulse strength and duration were chosen to minimize the number of conditions and to increase the power of the analysis. We used a weak pulse strength, well within the variability of the ongoing motion, to avoid a conspicuous event in the trial, which might promote a change in decision strategy. Because the motion pulses had no bearing on the reward, a detectable motion pulse could encourage the monkey to ignore it. We collected 18,017 trials from one monkey for this experiment.

Motion energy analysis. The motion energy in each trial was calculated by using two pairs of spatiotemporal filters (see Fig. 4A). Each pair was either selective for the direction of coherent motion (i.e., toward T_{in}) or

its opposite direction (the $\pm x$ dimension below). Each directional filter is formed as the sum of two space-time separable filters (Adelson and Bergen, 1985; Shadlen and Carney, 1986). The spatial filters are even and odd symmetric fourth-order Cauchy functions as follows:

$$\begin{aligned} f_1(x, y) &= \cos^4(\alpha) \cos(4\alpha) \exp\left(-\frac{y^2}{2\sigma_g^2}\right) \\ f_2(x, y) &= \cos^4(\alpha) \sin(4\alpha) \exp\left(-\frac{y^2}{2\sigma_g^2}\right), \end{aligned} \quad (1)$$

where $\alpha = \tan^{-1}(x/\sigma_c)$. Along the x -dimension, the filters resemble Gaussian-weighted sinusoids. The envelope and the period of the carrier sinusoids are controlled by the order (fourth) and σ_c (0.35°). They are windowed along the orthogonal (y) dimension using a Gaussian envelope (σ_g°). The two temporal impulse responses were based on a linear filter of the following form:

$$\begin{aligned} g_1(t) &= (60t)^3 \exp(-60t) \left[\frac{1}{3!} - \frac{(60t)^2}{(3+2)!} \right] \\ g_2(t) &= (60t)^5 \exp(-60t) \left[\frac{1}{5!} - \frac{(60t)^2}{(5+2)!} \right] \end{aligned} \quad (2)$$

The constants in Equations 1 and 2 confer a spatiotemporal frequency passband consistent with that of MT neurons (Movshon et al., 1988) and centered on the speed of coherent motion in our display.

A linear motion filter is constructed by adding appropriate combinations of the f_{ij} . For example, $f_1g_1 + f_2g_2$ and $f_2g_1 - f_1g_2$ pass motion in the $+x$ direction. These two linear filters are in space-time quadrature. They were convolved with the three-dimensional spatiotemporal pattern of the dynamic random dot stimulus. The results of the two convolutions were squared and summed together to extract the local motion energies at each point in the image as a function of time. These local motion energies were summated across space to calculate the motion energy over the stimulus as a function of time. Finally, motion energies for opponent directions were subtracted from each other to measure the net motion energy in one direction (see Fig. 4B,D).

To generate the predicted patterns of motion energy in Figure 4C, we applied a leaky- and a bounded-accumulator model to motion stimuli matching those analyzed from the experiment (other panels in Fig. 4). The decision mechanisms were as follows: (1) choose the larger of two leaky accumulators when the motion stimulus terminates, or (2) choose the perfect accumulator that reaches its bound first, even if the motion has not terminated. Each model has two free parameters: (1) a coefficient (k) that converts motion energy to the momentary evidence, and (2) a bound height (B) or a leak time constant (τ) for the bounded- or leaky-accumulation models, respectively. These free parameters were adjusted to fit the behavioral performance of the monkey (see Fig. 3) (bounded-accumulation, $k = 0.46$, $B = 19.9$; leaky-accumulation, $k = 0.45$, $\tau = 298.5$ ms). Both models achieve good fits to the data in Figure 3, although a model comparison favors bounded accumulation (Bayes information criterion, 59.5; Bayes factor, 8.35×10^{12}). Importantly, the free parameters of the models were not chosen to fit the motion energy profiles depicted in Figure 4D. Thus, the fitting ensures that the predictions are based on mechanisms that are compatible with the behavioral observations.

Analysis of behavioral data. The probability correct values shown in Figure 3 were calculated by dividing the trials into 20 quantiles (bins) based on stimulus duration. The probability of a correct response was calculated for each coherence level in each bin. Therefore, the data points in Figure 3 represent equal number of trials (1/20 of the total). The number of stimulus duration bins was not crucial for the results, and similar results were obtained for different bin numbers. The data from each bin were fit by a psychometric function of coherence (Quick, 1974) as follows:

$$P(C)_{t_1 \leq t \leq t_2} = 0.5 + 0.5 \left(1 - \exp\left(-\left[\frac{C}{\alpha_r}\right]^{\beta_r}\right) \right), \quad (3)$$

where α_r and β_r are the fit parameters, and C is the coherence level. The value of α_r is the discrimination threshold and corresponds to the coher-

ence level that elicits 81.6% correct responses. β_r determines the steepness of the psychometric function for a particular threshold. The fits to Equation 3 were obtained using maximum likelihood. The discrimination thresholds were similar for delay and no-delay trials. The data for the two trial types were combined for the calculation of discrimination thresholds (α_r) and for the motion energy analysis (see Figs. 3, 4). Similar results were obtained for each trial type individually.

To find the stimulus durations for which the accumulation process deviated from perfect, we fit a bilinear function to the 20 discrimination thresholds. The first line, which was constrained to a slope = -0.5 , was fit to the first n data points (see Fig. 3B, red line). The second line was fit to data points n to 20 and could take any slope. The maximum likelihood fit places the transition point at $n = 15$, which corresponds to stimulus durations 403–443 ms (mean, 420 ms).

In the motion pulse experiment, we measured the choice bias caused by the pulses using the following logistic regression:

$$P_{\text{right}} = \frac{1}{1 + e^{-\beta(C + \alpha_1 + \alpha_2 I)}}, I = \begin{cases} 1 & \text{rightward pulse} \\ -1 & \text{leftward pulse} \end{cases} \quad (4)$$

where, β and α_i are fit parameters. C is the coherence level that takes positive and negative values for rightward and leftward motion, respectively. Therefore, α_2 represents the shift of the psychometric function, in units of motion coherence, caused by the pulses. The null hypothesis is that motion pulses do not cause significant bias in the behavior (H_0 : $\alpha_2 = 0$).

To measure the effect of pulse time on choice bias for the long-duration trials, we extended Equation 4 to the following:

$$P_{\text{right}} = \frac{1}{1 + e^{-\beta(C + \alpha_1 + \alpha_2 I + \alpha_3 T_p + \alpha_4 T_p I)}} \quad (5)$$

where T_p is a group indicator for the pulse time. The trials were binned into quartiles based on the start time of the pulse, and T_p took values 1–4 corresponding to the quartiles. α_4 indicates how the shift of the psychometric function changes as the pulse time increases (H_0 : $\alpha_4 = 0$). This value was significantly negative ($\alpha_4 = -0.40 \pm 0.19$) indicating a diminishing effect for later pulses. Also, α_2 is significantly larger than 0 (1.51 ± 0.52 ; $p = 0.002$), compatible with the results of Equation 4. α_3 was indistinguishable from 0 (-0.12 ± 0.19 ; $p = 0.26$), indicating no inherent effect of pulse time independent of pulse direction.

The fits to Equations 4 and 5 were obtained using maximum likelihood. The fits for Figure 5C–F were obtained from Equation 4, but the slopes of psychometric function (β) and α_1 were constrained to be similar for all quartiles. Similar results were achieved without this constraint.

Analysis of neural data. Average spike rates (Figs. 2, 6, left panels) were calculated by aligning action potentials to the onset of the random dot stimulus. Data from each trial were truncated 50 ms before the saccadic eye movement, except for the analysis of bound crossing times (see below). For each trial, the spike train was converted into a rate function by convolution with $[1 - \exp(-t/\tau_r)] \exp(-t/\tau_d)$, where τ_r , the rise time constant, equals 1 ms, and τ_d , the decay time constant, equals 25 ms. The peristimulus time histogram (PSTH) for each cell was calculated by averaging the rate functions across the selected trials. The population PSTH was calculated by averaging the PSTHs across the recorded cells. For the analysis in Figure 6, short- and long-RT groups were formed separately for each neuron and for each coherence level, thereby ensuring that an equal number of trials contributed to the two groups. The traces shown in Figure 6 were formed by averaging across the five coherence values for a neuron and then combining these averages across the neurons. The strategy permits each neuron to contribute equally to the averaged rates shown. Other weighting schemes achieve similar results.

In addition to inspection of the firing rate functions, we estimated the start of decision-related activity in LIP by analyzing the variance of the firing rate across trials as a function of time from motion onset. We measured the spike counts in 100 ms epochs from each trial. By shifting this 100 ms window, we estimated a time-dependent mean and variance of the counts across the trials, using all motion strengths and directions. It is useful to consider the measured spike count variance as the sum of an

input variance, attributable to variation in the spike rate (across trials), plus the variance of the spiking process. Under the assumption that spiking is a point process, the input variance may be estimated by subtracting $\phi \times \text{mean spike count}$ from the measured variance. ϕ is an idealized Fano factor that would determine the spike count variance in the absence of any variability in the input. We set $\phi = 1$ for Figure 2B, but the results of analysis are not critically dependent on the exact value of ϕ . For an accumulation of noisy evidence (e.g., random walk), the input variance is expected to increase linearly as a function of time.

To determine whether the buildup rate of neural activity was different for short- and long-RT trials (see Fig. 6), we divided the T_{in} trials for each cell and coherence level into equally sized short- and long-RT groups. The neural activity for each RT group was then averaged across coherences and cells. A linear function was fit to the firing rates measured in a 150 ms window (6 nonoverlapping 25 ms periods to allow more or less independent samples of firing rate) beginning 207 ms after motion onset.

$$y = \beta_1 + \beta_2 I_T + \beta_3 t + \beta_4 I_{RT}$$

$$I_T = \begin{cases} 1 & \text{short-RT group} \\ 0 & \text{long-RT group} \end{cases}, \quad (6)$$

where y is the average population firing rate at time t in the specified window, and β_i is the set of fitted coefficient estimated by least square method. β_4 represents the difference in buildup rates for the short- and long-RT groups. The null hypothesis is that the buildup rate is equal for the two RT groups ($H_0: \beta_4 = 0$). β_2 represents the difference in firing rate at the beginning of the accumulation process for the two RT groups. It was not significantly different for the three panels in Figure 6 ($p > 0.1$). For the analysis of the effect of baseline firing rate on response time, see Equation 8.

For the analysis in Figure 7, we calculated the expected bound crossing time of the firing rates on single trials resulting in T_{in} choices. The bound height for each neuron was estimated using the trials that contained a delay period. The sustained activity was measured in a 500 ms window beginning 400 ms from the onset of the delay period, averaged for all T_{in} choices. This sustained level of activity was similar for short-duration stimuli (<200 ms; firing rate, 34.9 ± 3.1 spikes/s) and long-duration stimuli (>700 ms; firing rate, 35.3 ± 3.2 spikes/s) and for different coherence levels. A three-way ANOVA with coherence and duration as the fixed effects and the neurons as the random effect showed no significant effect for coherence ($F_{(4,270)} = 1.07$; $p = 0.37$) or stimulus duration ($F_{(1,58)} = 0.36$; $p = 0.55$). Each trial was assigned a bound-crossing time, T_{bc} , by finding the first time when the smoothed firing rate exceeded the bound height for at least 30 ms. The 30 ms window was used to balance the need to detect bound crossings reliably while preserving enough temporal precision to calculate the correlations. The procedure detected bound crossings for 81.2% of individual trials for the short-duration, no-delay stimuli, 95.2% of short-duration delay trials, and 94.4% of trials for the long-duration stimuli. To ensure the positive correlations are not attributable to the censorship of the neural responses before the saccade, the search for the bound-crossing time was extended to 200 ms after the saccade. Similar results were obtained without this extension. We also used Monte Carlo methods to validate this analysis (Fig. S3, available at www.jneurosci.org as supplemental material).

To test whether the correlation coefficients between T_{bc} and RT are larger for short-duration trials with and without delay periods, we used the following regression analysis:




Figure 1. Direction discrimination task. After the monkey fixated, two choice targets appeared in the periphery. One of the targets (T_{in}) was within the RF of the neuron, indicated by the gray shading. The other target (T_{out}) was in the opposite hemifield. After a 250–600 ms delay period, dynamic random dots appeared within a 5° circular aperture centered on the fixation spot. The stimulus remained on for 80–1500 ms. Motion strength in each trial was selected randomly from a predefined set; the net direction was either toward T_{in} or T_{out} . The disappearance of the fixation spot (Go signal) instructed the monkey to execute a saccadic eye movement to one of the choice targets. The Go signal either coincided with termination of the motion stimulus (no-delay trials) or followed a 500–1000 ms delay period. A liquid reward was administered for choosing the target along the direction of motion and in half of the trials with no coherent motion. All trial types were randomly intermixed.

$$z = \beta_1 + \beta_2 C + \beta_3 I + \beta_4 CI \quad (7)$$

where z is the transformed value of the correlation coefficients (Fisher z), I is an indicator variable that is one for the no-delay trials and zero for the delay trials, and C is the motion coherence. The z -transformed correlation coefficients have normal distribution with known SEs, thus permitting weighted regression. The null hypothesis is no difference in the z values for the two trial types ($\beta_3 = 0$). No other terms were significant in the analysis ($p \geq 0.1$), indicating that motion strength did not affect the correlation.

We tested whether variation in baseline firing rate could explain the correlation between T_{bc} and saccadic reaction times. We evaluated the multiple regression model as follows:

$$RT = \beta_1 + \beta_2 T_{bc} + \beta_3 \bar{b}, \quad (8)$$

where RT is the response time from the Go signal in a trial, and $\bar{b}$ is the standardized baseline activity in the trial in a 100 ms window before the motion onset (i.e., the measured value expressed as a difference from the mean for the neuron in units of sample SD). For all coherence levels, β_2 remained significantly different from zero ($p < 0.01$; $H_0: \beta_2 = 0$). Moreover, baseline firing rates did not affect RT ($p > 0.17$; $H_0: \beta_3 = 0$) for any coherence level.

Only the trials in which the monkeys completed the direction discrimination task were used for the analyses. Except for the analysis of response variance and the analysis of error trials, all graphs and analyses of neural responses used only trials in which the monkey responded correctly. For this purpose, all trials using the 0% coherent motion are considered correct, although the reward was administered randomly in half of the trials. The correlation between T_{bc} and RT, and also the difference in buildup rate of activity varied continuously as a function of motion duration. We chose to focus on the short (<200 ms) and longer (>700 ms) durations, that is the two ends of the continuum, to simplify the




Figure 2. Time course of neural and behavioral responses. **A**, Average activity from 51 LIP neurons. The neural responses begin to convey information about motion direction and choice after ~ 200 ms. **B**, The rise in neural response variance corresponds with accumulation of noisy evidence. The graph shows the input variance, the portion of variance that is not explained by the change in mean firing rate (see Materials and Methods). The arrow indicates when the input variance begins to increase reliably. **C**, Distribution of the monkey's RT from the motion onset for very short, no-delay trials (motion durations, 80–150 ms; mean, 100 ms). **D**, LIP responses to trials in **C**. The brown horizontal bar shows the range of motion durations. The box-and-whisker at lower right depicts the range of RTs for these trials. The central box shows the median and interquartile range.

presentation of data. The results do not rest critically on this arbitrary choice.

Results

Two rhesus monkeys were trained to perform a direction discrimination task (Fig. 1). In each trial, a dynamic random-dot motion stimulus was presented in one of two directions. The direction, strength (coherence), and duration of motion were varied randomly in each trial. The net direction of motion could be toward one of the two targets, which were shown on the opposite sides of the visual field. One of the targets was in the RF of the recorded neuron. After the motion stimulus, the fixation point turned off immediately for half of the trials (no-delay trials) or after a 500–1000 ms delay period for the other half (delay trials). The disappearance of the fixation point was the Go signal for the monkeys. They reported their decision about motion direction by making a saccadic eye movement to the target corresponding to the direction of motion. The stimulus duration was controlled by the experimenter and varied randomly from 80 to 1500 ms according to an exponential distribution.

We first describe the time course of decision-related neural activity in area LIP and then turn to an analysis of the monkey's behavioral performance. We show that bounded accumulation of evidence explains the performance on this task. We then return to the neurophysiology to describe a neural correlate of bounded accumulation predicted from the behavioral analysis.

A challenge to the role of LIP in the decision process

We recorded from 51 LIP neurons while the monkeys performed the direction discrimination task. All neurons had spatially selective persistent activity on simple delayed eye movement tasks to peripheral targets (see Materials and Methods). While the monkey performed the motion task, one of the choice targets (T_{in}) was in the response field of the neuron; the random dot motion dis-

play and the other choice target (T_{out}) were outside the RF. As shown in Figure 2A, the firing rate of these neurons reflects both the choice the monkey makes at the end of the trial (T_{in} or T_{out}) and quality of evidence (the motion strength) in support of these choices (indicated by line color, Fig. 2A), consistent with previous investigations (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Huk and Shadlen, 2005). Thus, it has been suggested that LIP represents a decision variable that mediates the transformation of accumulated evidence to a binary choice on this task (Gold and Shadlen, 2007). It is the time course of this activity that concerns us here.

Close inspection of Figure 2A reveals that the firing rates associated with T_{in} and T_{out} choices begin to diverge >200 ms after motion onset (Kim and Shadlen, 1999; Roitman and Shadlen, 2002; Huk and Shadlen, 2005). This divergence marks the beginning of the formation of the decision variable represented in LIP. This time can be estimated with greater sensitivity by measuring the variance of the firing rate across trials as a function of time. The variance increases when the neural activity diverges for different choices and stimulus strengths. If LIP neurons reflect an accumulation of noisy evidence from earlier visual areas, then a portion of the measured variance, termed the “input variance” (see Materials and Methods), should increase linearly as a function of time (Fig. 2B). This increase in input variance was first evident 207 ± 13 ms after the motion onset. It is the earliest point in time that LIP can be said to represent information bearing on the decision. Indeed, additional time must elapse before the evidence can inform a decision. For example, the differences in mean responses of individual neurons do not achieve a common statistical criterion until 257 ± 9 ms after motion onset (t test; $p < 0.05$).

The delayed representation of the accumulated evidence in LIP poses a challenge to the role of LIP neurons in the formation of decisions on the motion task. The latency of decision-related activity is at least 100 ms longer than the emergence of directional responses from neurons in areas MT and V1 (Britten et al., 1996; Bair et al., 2002; Osborne et al., 2004), and it exceeds simple saccadic latencies (Carpenter, 1988). It seems possible that the brain could acquire enough information to reach a decision about direction before LIP neurons even begin to exhibit decision-related changes in their firing rate.

We therefore challenged the hypothesis that LIP plays a role in decision making on this task by examining decisions based on brief exposures to visual motion. In many trials, the monkey received the Go signal long before the dissociation point in LIP activity was reached. We measured the latency between onset of motion and the monkey's eye movement response on these short-duration trials (ranging from 80 to 150 ms; 23% of trials). Although the motion was brief, the eye movements were initiated only after a relatively long latency. The distribution shown in Figure 2C has its central tendency almost 400 ms after stimulus onset (mean, 393 ± 0.7 ms; interquartile range, 366–413 ms). This is 150–200 ms after the emergence of the decision build up in LIP. In fact, the entire range of response times for these short

trials relative to the neural activity is compatible with a process that awaits the dissociation point in LIP activity before forming a decision (Fig. 2*D*).

A possible concern is that the monkeys simply procrastinated in these trials. We do not think this is likely for two reasons. First, there was a clear incentive to respond quickly to increase the reward rate. This is a potent incentive in our experience (Janssen and Shadlen, 2005). Second, the same set of stimuli, matched for coherence and exposure duration, gave rise to much faster responses on trials incorporating a delay period after the stimulus. For the short-duration stimuli, the mean RT from the Go signal was 272.9 ± 0.6 ms. Insertion of a delay period reduced this latency by >50 ms (mean RT, 219.1 ± 0.6 ms; *t* test, $p < 10^{-8}$), which is similar to the latencies of simple visually guided saccades in our screening trials (220.7 ± 0.8 ms).

Bounded integration of evidence

The build up of activity in LIP represents an accumulation of evidence in favor of one or the other direction alternatives. The stochastic motion stimulus used in these experiments produces a stream of evidence samples that are statistically independent over the time course of the trial. Were the accumulation perfect, it would lead to an improvement in sensitivity (i.e., accuracy) governed by the square root of viewing duration. The data, however, are only partly consistent with this prediction. When the duration of the stimulus increased, the probability of a correct response improved for all coherence levels (Fig. 3*A*), but this improvement saturated for moderate stimulus durations. We quantified this effect by extracting a threshold from the psychometric functions obtained at each of the 20 time bins depicted in Figure 3*A*. Figure 3*B* shows a clear decrease in discrimination threshold (i.e., an increase in sensitivity) when the monkeys were provided longer viewing durations. A square root of time improvement in the accuracy would result in a slope of -0.5 on the log-log graph in Figure 3*B*. Although the data approximately conform to this prediction for the shorter-duration stimuli (<420 ms), there is a clear deviation from this relation at longer durations. It is as if for the shorter-duration stimuli the monkey used all the available information, but for the longer durations, the information was used only partially.

There are two possible explanations for this departure from perfect integration. The first possibility is that subjects accumulate the information throughout the trial, but part of the accumulated information leaks away simultaneously as time goes by. This “leaky accumulator” hypothesis has been suggested previously (Smith, 1995; Usher and McClelland, 2001; Bussemeyer et al., 2006) and is widely assumed to account for psychophysical performance, although it has not been tested directly, to our knowledge. An important prediction of this hypothesis is that the information acquired later in a trial should exert a larger influence on the subject’s choice than the information acquired earlier, which has “leaked” away by the time the decision is rendered. Thus leaky integration explains the diminishing improvement on long duration trials, because it posits that information presented at the beginning of the trial is lost, at least partially.

The second possibility is that subjects accumulate information only up to a point when the decision process terminates, which could occur through an explicit bound (Link, 1992; Mazurek et al., 2003; Palmer et al., 2005) or an implicit termination, such as one mediated by attractor dynamics (Wang, 2002; Wong and Wang, 2006). Thereafter, information is simply ignored. The idea of a bounded accumulation is motivated by a confluence of behavioral, neurophysiologic, and modeling studies in the con-




Figure 3. Limited improvement in perceptual accuracy with longer motion viewing duration. *A*, Probability of correct response as a function of stimulus duration for the five motion strengths. The trials are divided into 20 quantiles based on stimulus duration ($n \approx 1024$ trials per data point). *B*, Behavioral discrimination thresholds deviate from perfect accumulation for longer viewing duration. The discrimination thresholds were estimated by fitting cumulative Weibull function to the 20 columns of points in *A*. The threshold is motion strength supporting 81.6% correct choices. The red line shows the expected change of threshold for perfect accumulation of evidence. Error bars represent SE.

text of reaction-time tasks, where subjects control viewing duration (Luce, 1986; Link, 1992; Carpenter and Williams, 1995; Roitman and Shadlen, 2002; Mazurek et al., 2003; Ratcliff and Smith, 2004; Lo and Wang, 2006). It is possible that a bound might terminate the integration process even when viewing duration is controlled by the environment (here, the computer). One of the important predictions of this hypothesis for our task is that the information provided later in the trial should exert a smaller influence on choice. Thus, bounded accumulation explains the diminishing improvement in long-duration trials, because it posits that information is simply ignored once the decision bound is attained.

These two hypotheses can be differentiated by analyzing the time course of the motion evidence leading to the monkey’s choices (Fig. 4). The analysis is especially informative on trials using the neutral motion stimulus (0% coherence) shown for long durations (e.g., >700 ms). In each trial, a particular random




Figure 4. The effect of motion information on decisions diminishes at longer viewing times. **A**, Spatiotemporal filters used in the motion energy calculation. The filters in the left and right columns are selective for opposite directions of motion. The two filters in each column form a quadrature pair. Application of the filters permits extraction of the motion energy as a function of time in each trial. **B**, Average motion energy for the 0% coherence trials with durations longer than 700 ms ($n = 1811$). Positive and negative values indicate rightward and leftward motion, respectively. The shaded region indicates SEM. **C**, Expected separation of motion energy profiles for rightward (red) and leftward (blue) choices for a simulated bounded accumulator (top) and a simulated leaky accumulator (bottom). In trials with long durations, leaky accumulation would render information at the beginning of the trial irrelevant, whereas bounded accumulation would render information at the end of the trial irrelevant. **D**, Separation of motion energy profiles for the monkeys' rightward and leftward choices (same trials as in **B**).

dot stimulus gives rise to a noisy stream of motion information that tends to fluctuate in magnitude and direction. For the 0% coherence, the net motion energy (Adelson and Bergen, 1985; Watson and Ahumada, 1985; Britten et al., 1993) is zero, on average, as shown in Figure 4B. However, at any moment, in any one trial, the stimulus contains motion energy toward or away from the targets. In principle, these fluctuations ought to influence the monkey's choices. We show that they do, and moreover the time course of their influence is better explained by bounded accumulation.

Figure 4C illustrates the predictions of the leaky and bounded accumulator models. Both models were optimized to match the accuracy results for the entire data set (see Materials and Methods). According to the leaky accumulator hypothesis, the motion energy profiles associated with the two choices should differ toward the end of the long trials. Early fluctuations in motion energy should have little bearing on the decision, because they fail to accumulate (they leak away). According to the termination hypothesis, the motion energy profiles should differ toward the beginning of the longer trials. Late fluctuations in motion energy should have little bearing on the decision, because they often occur after the decision has terminated. The results depicted in Figure 4D clearly favor the termination hypothesis. The random fluctuations in motion energy that occur early during stimulus viewing affect the choices, whereas later fluctuations are less effective. Here, we focused on the longer-duration trials merely because they are long enough to have distinct early and late por-

tions. Similar patterns are apparent in the motion energy profiles of shorter-duration trials and in the motion energy profiles of trials with >0% coherence (data not shown).

To verify the motion energy analysis and to provide a direct test for the differential effect of early and late evidence in our task, we performed an additional experiment on one of the monkeys. In this experiment, a 200 ms pulse of 3.2% coherent motion was introduced at a random time on each trial. Such weak motion pulses have been shown previously to influence the monkey's choice on the reaction time version of the direction discrimination task (Huk and Shadlen, 2005; Wong et al., 2007). We measured the bias of the monkey's choice for the rightward and leftward pulses. To allow inclusion of pulses on every trial, we extended the minimum stimulus duration to 280 ms. The motion pulses significantly biased the monkey's choice toward the pulse direction (Fig. 5A) (psychometric function shift across all trials, $1.3 \pm 0.1\%$ coherence; $p < 10^{-8}$) (Eq. 4). Importantly, these pulses were more effective when they were presented earlier in the trial. To demonstrate this, we examined all the trials lasting 700 ms or longer. As shown in Figure 5C–F, the early pulses caused significantly larger biases than the late ones (Fig. 5B) ($p = 0.018$) (Eq. 5). In fact, the bias caused by the late pulses was indistinguishable from 0 (third quartile: shift, $0.2 \pm 0.4\%$, $p = 0.26$; fourth quartile: shift, $-0.1 \pm 0.4\%$, $p = 0.63$). This direct experimental manipulation confirms the conclusion drawn from the analysis of motion energy: the brain ignores the later informa-

tion, presumably because it has already terminated the decision process.

Neural correlate of bounded integration in short-duration trials

The preceding analysis suggests that the mechanism for decision making in long-duration trials resembles the type of bounded accumulation of evidence that is thought to explain the speed and accuracy of decisions in choice-reaction time tasks (Ratcliff, 1980; Smith, 1988; Link, 1992; Gold and Shadlen, 2007). We therefore hypothesized that a similar process, reflected in the activity of LIP neurons, might underlie the monkey's choices in this experiment. Although the present experiment is not a reaction time experiment, there is variation, nonetheless, in the response times relative to the Go signal. We have demonstrated already that decision-related activity in LIP precedes the behavioral response and is predictive of the monkey's choice. We next asked whether the neural activity is also correlated with the monkey's RT in a way that is expected from an accumulation-to-bound mechanism.

We performed two analyses to test this idea. In the first, we asked whether the rate of buildup of activity in area LIP affects the time of saccade initiation. We first consider all T_{in} trials with stimulus durations of 200 ms or less in which the Go signal coincided with offset of the motion stimulus (i.e., no-delay trials). In these trials, the Go signal is given before the accumulation begins in area LIP. According to our hypothesis, the variation in RT in these trials should reflect the time it takes to reach a decision, hence the rate of evidence accumulation. We therefore divided the trials into two groups based on the monkey's RT relative to the Go signal, termed short- and long-RT groups (Fig. 6A). Each neuron contributed half of its trials at each motion strength to the short- and long-RT groups.

The firing rates for the two groups, averaged across coherence levels, are shown in Figure 6A. The buildup rate of neural activity was significantly larger for the trials in which the monkey responded faster (Fig. 6A) (slope change, 48.7 ± 21.3 spikes/s²; $p = 0.02$). This is consistent with our hypothesis that the LIP firing rates reflect formation of a decision, which terminates earlier when the buildup is faster. If this is correct, then we would expect the differences in the buildup rates for short- and long-RT groups to be absent when the decision is made before the Go instruction.

This prediction is supported by the graphs in Figure 6, B and C. For short-duration trials (<200 ms) in which a 500–1000 ms delay period was introduced before the Go signal (i.e., delay trials), the firing rate of the short- and long-RT trials essentially overlapped from the beginning of the accumulation process (Fig. 6B) (slope change, -17.4 ± 18.2 spikes/s²; $p = 0.18$). Similarly, the firing rates were overlapping for trials with stimulus durations of 700 ms or more (Fig. 6C) (slope change, -6.6 ± 19.0




Figure 5. Late motion pulses do not bias the monkey's behavior, whereas early pulses do. In each trial, a weak motion pulse (3.2% coherence to the right or left) was introduced at a random time. The range of stimulus durations was 280–1500 ms to accommodate these pulses. **A**, The probability of choosing the rightward target is plotted as a function of motion strength for all trials with rightward (black) and leftward (gray) motion pulses. Positive and negative coherences correspond to rightward and leftward motions, respectively. The motion pulses significantly biased the monkey's choice toward their corresponding target. This bias is quantified by the horizontal separation of the two functions, here equivalent to $1.3 \pm 0.1\%$ coherent motion. **B–F**, Effect of pulse time on pulse effectiveness. Trials with long stimulus duration (>700 ms) were divided into four groups of equal size (quartiles) based on the pulse time. For each group, a pair of psychometric functions was constructed for the two pulse directions (**C–F**). The horizontal separation is plotted as a function of pulse time in **B** (error bars are SE of the shifts). The motion pulses caused a significant shift in the psychometric functions for the earlier two quartiles (**C, D**), but not for the later quartiles (**E, F**). The plots in **C–F** focus on the middle of the coherence range to allow better visualization of the effect of the weak motion pulses.

spikes/s²; $p = 0.37$). The absence of a correlation between buildup rates and response times in Figure 6, B and C, can be explained if the decision process has terminated before the Go instruction on these trials. For the trials comprising all three panels of Figure 6, decisions probably arose faster when LIP activity built up faster. However, faster decisions only affect the measured RT in Figure 6A. For the other panels, the variation in RT after the Go signal does not reflect the speed of decision formation. The traces in Figure 6, B and C, are overlapping, because they are formed from a random collection of faster and slower decisions. Indeed, they remain overlapping for the duration of the trial.

In a second analysis, we measured a correlation between the monkey's RT and an estimate of the end of the decision on single trials. The estimate of the decision time was obtained in two steps: (1) the spike discharge was smoothed to obtain a continuous estimate of spike rate, and (2) this rate estimate was compared with a decision criterion or bound. The bound height for each neuron was estimated using the sustained level of firing rate measured during the waiting period between offset of the stimulus




Figure 6. Time course of LIP activity for T_{in} choices in short- and long-RT trials. For each cell and coherence level, the trials were divided into two groups based on the median RT. The neural activity was averaged across coherences and cells. **A**, No-delay trials with short motion duration (<200 ms). **B**, Delay trials with short motion duration (<200 ms). **C**, No-delay trials with longer motion duration (>700 ms).

and Go instruction on the delay trials. Of course, we do not know the actual level of the bound, but according to our hypothesis, it should be near the level of firing rate that a neuron attains at the end of the long-duration stimuli or during the delay period after the stimulus in the delay trials (Mazurek et al., 2003). Guided by this assumption, we extracted the bound-crossing times, which we take to be a noisy estimate of decision time, by finding the first time when the firing rate on the trial exceeded the bound height for at least 30 ms. We then asked whether this estimate bears any detectable relationship to the measured variation in reaction times, relative to the Go instruction.

A weak positive correlation is evident on short-duration stimuli. For 0% coherence trials with no delay and short stimulus durations (<200 ms), the monkey's RTs relative to the Go signal were positively correlated to the decision times estimated from the neural data (Fig. 7A) ($r = 0.16$; $p = 1.4 \times 10^{-4}$, Fisher z). The correlation remained significant for all coherence levels ($p < 0.01$). The time from the estimated neural decision to saccade initiation was 94.4 ± 2.0 ms (mean $\pm$ SEM for all trials represented in Fig. 7A,B). This value depends on the choice of bound and the method for smoothing the spike trains, but the sign and significance of the

correlation did not depend critically on these assumptions. We also ensured that the correlation was not explained by variations in baseline firing rates at the beginning of the trial (see Materials and Methods, Eq. 8). The analysis indicates that whatever effect the baseline activity may have on the response time is exerted through changing the bound-crossing time.

Importantly, inserting a delay period before the Go signal reduced the correlation between the bound-crossing time and the response time (Fig. 7C,D) ($p = 0.0005$) (Eq. 7). In fact, the correlation coefficients were not significantly different from 0. Here, our comparison is focused on the delay and no-delay trials with short stimulus durations because of the similarly large number of trials. The lack of correlation is expected, because, according to the bounded accumulation hypothesis, the decision has completed long before the Go signal in these trials. The variable RT in these trials is presumably explained by whatever mechanism accounts for variable motor latencies in the absence of a decision among alternatives (Carpenter, 1988; Hanes and Schall, 1996). The lack of correlation in these trials also rules out potentially trivial explanations for the correlation such as increased vigilance, arousal, attentiveness, or level of preparedness. Such causes would have induced similar correlations for the delay and no-delay trials. Together, these analyses (Figs. 4–7) support the hypothesis that a bounded accumulation of evidence, represented in LIP activity, underlies the monkey's choices even when the stimulus duration is controlled by the environment.

Errors are governed by the same mechanism as the correct responses

In error trials, the LIP neurons reflected the monkey's choice rather than the direction of random dot motion (Fig. 8A): the firing rate rose to almost the same level for the T_{in} trials independent of whether the monkey's choice was correct or incorrect. Moreover, the rate of buildup for fast and slow errors greatly resembled that of the pattern observed with fast and slow correct choices. To appreciate this, we pooled error trials using the lower range of motion strengths (6.4, 12.8, and 25.6%; errors were too infrequent above this range). For short stimulus duration trials (<200 ms) without a delay period, the buildup rate of neural activity was significantly larger in the trials in which the monkey responded faster (slope change, 59.4 ± 31.2 spike/s²; $p = 0.05$). Similar to the correct trials, the difference in the buildup rate was not significant for delay trials with short stimulus durations (slope change, -11.6 ± 29.6 spike/s²; $p = 0.35$). [The number of error trials for the long-durations stimuli (>700 ms) was too few to support a similar analysis.]

Finally, the analysis of neural bound-crossing times, estimated from individual short-duration trials, indicated significant correlation between the bound-crossing time and the monkey's re-

sponse time for the error trials (Fig. 8*B*). These results suggest that the monkey commits to an error, because the accumulated evidence reaches a bound (in this case, the wrong bound) perhaps because of noise in the stimulus or noise in the neural responses in sensory areas.

Discussion

The present study extends previous work on the neural basis of decision making in two ways. It shows that the putative neural correlate of a decision variable in LIP is represented in a time frame that is compatible with the timing of decisions. Also, it shows that the decision process carries its own rule for termination even when the period in which evidence is available and the time for action are under the control of the environment.

When provided with more information, the monkeys' accuracy improved, but by less than expected for perfect accumulation of evidence (Fig. 3). By examining the time course of motion information in the display and by inserting brief motion pulses at various times during motion viewing, we determined that improvement was curtailed, because the monkeys ignored information late in the trial. They did not appear to lose, or leak away, information acquired at the beginning of the trial (Figs. 4, 5). We hypothesize that this diminishing contribution of information to the decision (Luna et al., 2005) is a result of application of a termination rule before the Go signal, and that the termination rule is a bound on the accumulated information. According to this hypothesis, the monkey commits to one of the choices when the accumulated evidence reaches a criterion level (Link, 1992; Mazurek et al., 2003). A leaky accumulation process, although theoretically capable of describing the observed curtailed improvement of accuracy (Usher and McClelland, 2001; Ratcliff and Smith, 2004; Busemeyer et al., 2006), is not supported in this task.

By ruling out the leaky accumulator, we do not disregard the presence of all types of leak in the brain. All processes in the brain are ultimately implemented by leaky devices: neurons. However, it is possible to arrange leaky neurons in a recurrent network to implement nearly perfect accumulators (Wang, 2002; Major and Tank, 2004; Machens et al., 2005; Cisek, 2006; Ma et al., 2006; Wong and Wang, 2006; Bogacz, 2007).

We focused on leaky accumulation as the main alternative to bounded accumulation, but our findings rule out a much larger class of alternative models. One possibility is that the accumulation is simply limited in time. According to this alternative, the subject commits to a choice at some point in time and ignores the information arriving later. Another related possibility is that information is simply weighted less as time elapses. These models can explain the early separation of motion energy profiles, similar to Figure 4*D*. However, they fail to explain the LIP neural responses in our experiment, because they predict no correlation between the neural threshold crossing time and RT. Moreover, a more careful analysis of the accuracy data augurs poorly for the




Figure 7. Correlation between saccadic response times and the neural bound crossing time (T_{bc}). T_{bc} for each trial is the time when the firing rate first exceeded a threshold specified by the sustained level of activity of the cell in the delay period (see Materials and Methods). **A**, Relationship between RT and T_{bc} for 0% coherence trials with short durations (<200 ms) and no delay period. Each point represents one trial. The confidence ellipse is stretched $1.5 \times$ SD along the principle components of the data points. **B**, Pearson's correlation coefficient of RT and T_{bc} for the five motion strengths. Error bars represent the 95% confidence interval. **C**, **D**, Same as **A** and **B** for trials with a 500–1000 ms delay between the end of stimulus and the Go signal (same range of stimulus durations as in **A** and **B**).

plausibility of these alternatives. First, the time-limited accumulation fails to account for the changes of discrimination thresholds (Fig. S1, available at www.jneurosci.org as supplemental material). Second, because they dismiss the concept of a bound on accumulated evidence, all the alternatives to bounded accumulation assert a testable prediction: a pulse delivered at a particular time should affect the decision process identically for all motion strengths. In contrast, the bounded accumulation model predicts that the effectiveness of a pulse depends on both the time it is delivered and the strength of the motion that it perturbed. This is because for weaker motion strengths, the accumulated evidence takes longer to approach the bound. Therefore, the pulse should remain effective at later times for weaker motion than for stronger motion. This is exactly what we observed (Fig. S2, available at www.jneurosci.org as supplemental material).

We also found that, for each motion strength, the rate of rise and the bound-crossing times of LIP neural activity were correlated with the monkey's RT for short-duration trials but not for longer durations or delay trials (Figs. 6, 7). This pattern is also consistent with an accumulation-to-bound mechanism (Mazurek et al., 2003). It suggests that the accumulated evidence, represented in LIP, tends to plateau or saturate at the longer viewing durations used in previous experiments (Shadlen and Newsome, 2001). LIP appears to play a similar role whether the time course of evidence bearing on the decision is controlled by the subject or by the experimenter. Notably, despite the long latency for the representation of accumulated evidence in LIP, activity in this structure preceded the monkey's behavioral re-




Figure 8. Error responses are governed by the same mechanism as the correct responses. **A**, Average activity of the LIP neurons in correct and error trials. Neural responses were averaged across all stimulus durations and nonzero motion strengths. **B**, Correlation between RT and T_{bc} in error trials with short stimulus durations and no delay period. Error bars represent the 95% confidence interval. The number of errors was small at the higher motion strengths (e.g., $n = 52$ for 25.6% coherence), hence the larger confidence interval.

sponses even for extremely short stimuli. The finding is consistent with a recent microstimulation study, which supports a causal role for LIP in the decision process on this task (Hanks et al., 2006).

It is remarkable that the decision times estimated from the spike discharge on single trials should have any detectable relationship to the response times measured behaviorally. After all, the behavior is undoubtedly controlled by many neurons, and the estimates of bound-crossing times from a single neuron in each trial are highly variable. Thus, it is not surprising that the slopes of regression between the bound-crossing time and the monkey's response time are <1 , the value that might be expected for a mechanism that controls the decision. Additionally, for trials with fast neural bound crossings, the monkey's response time cannot decrease proportionally, because it is limited by the motor preparation delay and the time needed to detect the Go signal. Like other weak trial-by-trial correlations between neural activity and decisions, the observed correlations are probably explained by the tendency of neurons with similar response characteristics (in this case, RFs) to covary in their variable discharge (Zohary et al., 1994; Shadlen et al., 1996; Parker and Newsome, 1998; Pezaris et al., 1999).

The distribution of stimulus durations in our task helped us

identify the termination rule. We used exponentially distributed durations, which prevented monkeys from expecting the end of stimulus and encouraged them to start the decision process as the motion stimulus appeared (Janssen and Shadlen, 2005). Such temporal alignment of the decision process to an experimentally controlled reference time is necessary to explore when and how the decision process terminates. In contrast to this approach, a stimulus duration that is fixed or predictable (Britten et al., 1993; Shadlen and Newsome, 2001; Parker et al., 2002; DeAngelis and Newsome, 2004; Krug, 2004; Krug et al., 2004; Uka and DeAngelis, 2004) would allow subjects to start the decision process at times other than the beginning of stimulus. Such temporal smear across different trials could lead to the incorrect conclusion that subjects use the information throughout the whole stimulus presentation. The important point is that when the experimenter controls the duration of a stimulus, the brain may or may not exploit the extra time to its advantage. The policy that governs decision time is presumably the setting of the bound on the evidence. This bound establishes the trade off between accuracy and speed of decisions. The latter is measured when the decision is coupled to an immediate action, as in RT tasks. However, when the experimenter (or the environment) controls the stimulus duration, the same considerations appear to be in play, although the decision time is covert.

Such flexibility is essential for adjustment to the complexities of the real world. The stochastic function underlying the stimulus in our task was stationary, not biasing the monkey to any particular period in time. In contrast, a nonstationary function, more generally the nonstationarities of the real world, result in periods of time with more reliable evidence than others. If the source of evidence changes in a predictable way, for example, later evidence might be more informative. We predict that in these situations, the brain can flexibly change strategies to limit the evidence accumulation to periods with more relevant information, not necessarily the beginning of the stimulus. This possibility is an interesting direction for future research, especially in the light of our current finding that subjective stopping rules generalize beyond the setting of reaction time tasks.

Thresholds and accumulation to thresholds seem to be widely used in neural mechanisms, from spiking of individual neurons (Gerstein and Mandelbrot, 1964; Ricciardi, 1977; Shadlen and Newsome, 1994) to higher brain functions. Thresholds are involved in decision making (Laming, 1968; Link and Heath, 1975; Smith, 1988; Ratcliff et al., 1999; Reddi and Carpenter, 2000; Shadlen and Gold, 2004), target selection, and saccade initiation (Hanes and Schall, 1996; Schall and Thompson, 1999). They may be useful for the perception of temporal intervals (Leon and Shadlen, 2003; Janssen and Shadlen, 2005) and contribute to learning. They are even hypothesized to play a role in memory retrieval (Ratcliff, 1978; Wagner et al., 2005). The use of a threshold would necessitate mechanisms for detection of threshold crossing. In our task, a potential candidate is an interconnected network that includes LIP and the basal ganglia (Lo and Wang, 2006). This circuitry can detect the bound crossing in LIP activity and is capable of implementing a flexible bound, which is necessary for adjustable behavior in a complex environment (Link, 1992; Mazurek et al., 2003; Ratcliff and Smith, 2004; Palmer et al., 2005).

In fact, such a threshold-crossing mechanism together with a recurrent accumulator network (Cisek, 2006; Ma et al., 2006; Wong and Wang, 2006; Bogacz, 2007; Wong et al., 2007) could explain many of the phenomena touched on here. Some of these recurrent networks (Wang, 2002; Wong and Wang, 2006) have

attractor states that implement a decision termination without an explicit bound, although they too need a mechanism to detect absorption to the attractor state to initiate an action (Kiani et al., 2006; Lo and Wang, 2006). Such a mechanism is indispensable, particularly for reaction time tasks and for the short-duration trials in the task described in the current paper.

The presence of a termination rule invites us to revisit the contribution of sensory neurons to decisions. For stimulus durations of 2 s in our task, Britten et al. (1992) have shown that individual monkey MT neurons can perform as well as the monkey, provided that the activity of the neuron is averaged over the full stimulus presentation period. However, as demonstrated here, monkeys base their decisions on accumulation of evidence for much shorter durations. Thus, we predict that if the neural activity was averaged over the relevant period on each trial, individual neurons would perform substantially worse than they do from 2 s worth of information and therefore substantially worse than the monkey (Cook and Maunsell, 2002; Mazurek et al., 2003). Similarly, the relationship between trial-to-trial variability in the firing rate of sensory neurons and behavioral choices [choice probability (Britten et al., 1996; Parker and Newsome, 1998)] varies for relevant and irrelevant periods of neural responses. Calculation of choice probability over long stimulus presentations mixes the relevant and irrelevant periods and leads to underestimation. The similarity of the monkeys' performance on RT and fixed duration tasks (Roitman and Shadlen, 2002) adds further support to our contention that the brain simply ignores additional evidence that is available in the stimulus and in the firing of its sensory neurons. The use of a termination rule is not a consequence of training; we observed a similar diminishing effect of viewing duration in monkeys that were never trained on an RT task.

When the environment establishes a clear and predictable epoch for decision making, why should subjects make decisions based on a fraction of the available information? Why forego the increased reward that would accrue by improving accuracy? One answer may be the important, but often overlooked, concept of the sampling cost. Collecting evidence for reaching a decision consumes limited mental resources and retards engagement in other decisions or cognitive functions. When sampling of information is costly, the brain must strike a balance between the amount of collected information and the improvement of accuracy. Economists study this topic under the banner of satisficing (Simon, 1959). A striking prediction of such a balance is that even when subjects are not engaged in reaction time tasks, they will limit the sampling of information. We show that subjects behave in this way in our experiment, and we suggest a mechanism, bounded accumulation of evidence, based on behavioral and neural data. Our hypothesis unifies the decision mechanisms used in RT and non-RT tasks.

More generally, termination rules allow the brain to cascade states with complex dependencies on each other (Fine et al., 1998; Caspi et al., 2004). Unlike simple behaviors, which are more or less automatic, cognitive processes may be dynamically switched on and off at will. Knowledge of when to start the next process presupposes control over the termination of preceding states. Thus, the cost in accuracy associated with ignoring available information may be offset by capacities that underlie the very flexibility of thought.

References

- Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. *J Opt Soc Am A* 2:284–299.
- Bair W, Cavanaugh JR, Smith MA, Movshon JA (2002) The timing of response onset and offset in macaque visual neurons. *J Neurosci* 22:3189–3205.
- Bogacz R (2007) Optimal decision network with distributed representation. *Neural Netw* 20:564–576.
- Bracewell RM, Mazzoni P, Barash S, Andersen RA (1996) Motor intention activity in the macaque's lateral intraparietal area. II. Changes of motor plan. *J Neurophysiol* 76:1457–1464.
- 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.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10:1157–1169.
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100.
- Burr DC, Santoro L (2001) Temporal integration of optic flow, measured by contrast and coherence thresholds. *Vision Res* 41:1891–1899.
- Busmeyer JR, Townsend JT (1993) Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol Rev* 100:432–459.
- Busmeyer JR, Jessup RK, Johnson JG, Townsend JT (2006) Building bridges between neural models and complex decision making behaviour. *Neural Netw* 19:1047–1058.
- Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62.
- Carpenter RHS (1988) *Movements of the eyes*, Ed 2. London: Pion.
- Caspi A, Beutter BR, Eckstein MP (2004) The time course of visual information accrual guiding eye movement decisions. *Proc Natl Acad Sci USA* 101:13086–13090.
- Cisek P (2006) Integrated neural processes for defining potential actions and deciding between them: a computational model. *J Neurosci* 26:9761–9770.
- Cook EP, Maunsell JH (2002) Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat Neurosci* 5:985–994.
- de Lafuente V, Romo R (2005) Neuronal correlates of subjective sensory experience. *Nat Neurosci* 8:1698–1703.
- DeAngelis GC, Newsome WT (2004) Perceptual “read-out” of conjoined direction and disparity maps in extrastriate area MT. *PLoS Biol* 2:E77.
- Diederich A (1995) Intersensory facilitation of reaction time: evaluation of counter and diffusion coactivation models. *J Math Psychol* 39:197–215.
- Diederich A (1997) Dynamic stochastic models for decision making under time constraints. *J Math Psychol* 41:260–274.
- Fechner GT (1860) *Elements of psychophysics*. Leipzig, Germany: Breitkopf and Härtel.
- Fine S, Singer Y, Tishby N (1998) The hierarchical hidden markov model: analysis and applications. *Machine Learn* 32:41–62.
- Fuchs AF, Robinson DA (1966) A method for measuring horizontal and vertical eye movement chronically in the monkey. *J Appl Physiol* 21:1068–1070.
- Gerstein GL, Mandelbrot B (1964) Random walk models for the spike activity of a single neuron. *Biophys J* 4:41–68.
- Ghose GM, Maunsell JH (2002) Attentional modulation in visual cortex depends on task timing. *Nature* 419:616–620.
- Glimcher PW (2003) The neurobiology of visual-saccadic decision making. *Annu Rev Neurosci* 26:133–179.
- Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–394.
- Gold JI, Shadlen MN (2003) The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J Neurosci* 23:632–651.
- Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Green DM, Swets JA (1966) *Signal detection theory and psychophysics*. New York: Wiley.
- Grice GR (1972) Application of a variable criterion model on auditory reaction time as a function of the type of catch trial. *Percept Psychophys* 12:103–107.
- Hanes DP, Schall JD (1996) Neural control of voluntary movement initiation. *Science* 274:427–430.
- Hanks TD, Ditterich J, Shadlen MN (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat Neurosci* 9:682–689.
- Horowitz GD, Newsome WT (2001) Target selection for saccadic eye move-

- ments: direction-selective visual responses in the superior colliculus. *J Neurophysiol* 86:2527–2542.
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25:10420–10436.
- Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 8:234–241.
- Judge SJ, Richmond BJ, Chu FC (1980) Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res* 20:535–538.
- Kiani R, Hanks TD, Shadlen MN (2006) When is enough enough? *Nat Neurosci* 9:861–863.
- Kim C, Myung IJ (1995) Incorporating real-time random effects in neural networks: a temporal summation mechanism. In: *Proceedings of the Seventeenth Annual Meeting of the Cognitive Science Society* (Moore JD, Lehman JF, eds), pp 472–477. Hillsdale, NJ: Erlbaum.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Krug K (2004) A common neuronal code for perceptual processes in visual cortex? Comparing choice and attentional correlates in V5/MT. *Philos Trans R Soc Lond B Biol Sci* 359:929–941.
- Krug K, Cumming BG, Parker AJ (2004) Comparing perceptual signals of single V5/MT neurons in two binocular depth tasks. *J Neurophysiol* 92:1586–1596.
- Laming DRJ (1968) *Information theory of choice reaction time*. New York: Wiley.
- Leon MI, Shadlen MN (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* 38:317–327.
- Lewis JW, Van Essen DC (2000) Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *J Comp Neurol* 428:112–137.
- Link SW (1992) *The wave theory of difference and similarity*. Hillsdale, NJ: Erlbaum.
- Link SW, Heath RA (1975) A sequential theory of psychological discrimination. *Psychometrika* 40:77–105.
- Lo CC, Wang XJ (2006) Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat Neurosci* 9:956–963.
- Luce RD (1986) *Response times: their role in inferring elementary mental organization*. Belfast, UK: Oxford UP.
- Ludwig CJ, Gilchrist ID, McSorley E, Baddeley RJ (2005) The temporal impulse response underlying saccadic decisions. *J Neurosci* 25:9907–9912.
- Luna R, Hernandez A, Brody CD, Romo R (2005) Neural codes for perceptual discrimination in primary somatosensory cortex. *Nat Neurosci* 8:1210–1219.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Machens CK, Romo R, Brody CD (2005) Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* 307:1121–1124.
- Major G, Tank D (2004) Persistent neural activity: prevalence and mechanisms. *Curr Opin Neurobiol* 14:675–684.
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cereb Cortex* 13:1257–1269.
- McClelland JL (1979) On the time relations of mental processes: an examination of systems of processes in cascade. *Psychol Rev* 86:287–330.
- McClelland JL (1993) Toward a theory of information processing in graded, random, interactive networks. In: *Attention and performance XIV: synergies in experimental psychology, artificial intelligence and cognitive neuroscience* (Meyer DE, Kornblum S, eds). Cambridge, MA: MIT.
- Movshon JA, Newsome WT, Gizzi MS, Levitt JB (1988) Spatio-temporal tuning and speed sensitivity in macaque visual cortical neurons. *Invest Ophthalmol Vis Sci* 29:327.
- Osborne LC, Bialek W, Lisberger SG (2004) Time course of information about motion direction in visual area MT of macaque monkeys. *J Neurosci* 24:3210–3222.
- Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J Vis* 5:376–404.
- Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21:227–277.
- Parker AJ, Krug K, Cumming BG (2002) Neuronal activity and its links with the perception of multi-stable figures. *Philos Trans R Soc Lond B Biol Sci* 357:1053–1062.
- Pezaris JS, Sahani M, Andersen RA (1999) Response-locked changes in auto- and cross-covariations in parietal cortex. *Neurocomputing* 26–27:471–476.
- Platt ML, Glimcher PW (1997) Responses of intraparietal neurons to saccadic targets and visual distractors. *J Neurophysiol* 78:1574–1589.
- Quick Jr RF (1974) A vector-magnitude model of contrast detection. *Kybernetik* 16:65–67.
- Ratcliff R (1978) A theory of memory retrieval. *Psychol Rev* 85:59–108.
- Ratcliff R (1980) A note on modeling accumulation of information when the rate of accumulation changes over time. *J Math Psychol* 21:178–184.
- Ratcliff R, Smith PL (2004) A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111:333–367.
- Ratcliff R, Van Zandt T, McKoon G (1999) Connectionist and diffusion models of reaction time. *Psychol Rev* 106:261–300.
- Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. *Nat Neurosci* 3:827–830.
- Ricciardi L (1977) *Diffusion processes and related topics in biology*. New York: Springer.
- Rieke FM, Warland D, de Ruyter van Steveninck R, Bialek W (1997) *Spikes: exploring the neural code*. Cambridge, MA: MIT.
- 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.
- Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4:203–218.
- Schall JD, Thompson KG (1999) Neural selection and control of visually guided eye movements. *Annu Rev Neurosci* 22:241–259.
- Shadlen M, Carney T (1986) Mechanisms of human motion perception revealed by a new cyclopean illusion. *Science* 232:95–97.
- Shadlen MN, Gold JI (2004) The neurophysiology of decision-making as a window on cognition. In: *The cognitive neurosciences* (Gazzaniga MS, ed). Cambridge, MA: MIT.
- Shadlen MN, Newsome WT (1994) Noise, neural codes and cortical organization. *Curr Opin Neurobiol* 4:569–579.
- 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.
- Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16:1486–1510.
- Shadlen MN, Hanks TD, Churchland AK, Kiani R, Yang T (2006) The speed and accuracy of a simple perceptual decision: a mathematical primer. In: *Bayesian brain: probabilistic approaches to neural coding* (Doya K, Ishii S, Pouget A, Rao RPN, eds). Cambridge, MA: MIT.
- Simon HA (1959) Theories of decision-making in economics and behavioral science. *Am Econ Rev* 49:253–283.
- Smith PL (1988) The accumulator model of two-choice discrimination. *J Math Psychol* 32:135–168.
- Smith PL (1995) Psychophysically principled models of visual simple reaction time. *Psychol Rev* 102:567–593.
- Uka T, DeAngelis GC (2004) Contribution of area MT to stereoscopic depth perception: choice-related response modulations reflect task strategy. *Neuron* 42:297–310.
- Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev* 108:550–592.
- Van Essen DC (2002) Windows on the brain: the emerging role of atlases and databases in neuroscience. *Curr Opin Neurobiol* 12:574–579.
- Wagner AD, Shannon BJ, Kahn I, Buckner RL (2005) Parietal lobe contributions to episodic memory retrieval. *Trends Cogn Sci* 9:445–453.
- Wang XJ (2002) Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36:955–968.
- Watson AB, Ahumada Jr AJ (1985) Model of human visual-motion sensing. *J Opt Soc Am A* 2:322–341.
- Wong KF, Wang XJ (2006) A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 26:1314–1328.
- Wong KF, Huk AC, Shadlen MN, Wang XJ (2007) Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision-making. *Front Comp Neurosci* 1:1–11.
- Wundt W (1865) *Lehrbuch der Physiologie des Menschen*. Stuttgart: Ferdinand Enke.
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143.




Figure S1.
Kiani, Hanks & Shadlen

Comparison of bounded and time-limited accumulation models. The black dots show the monkey's discrimination thresholds as a function of stimulus duration (same as Fig. 3). The red curves show the best fit of the bounded accumulation model (**a**) and the time-limited accumulation model (**b**).

In the bounded accumulation model, the noisy momentary sensory evidence, which is proportional to the motion coherence, is accumulated until a decision bound is reached. This model has two free parameters: the proportionality constant between coherence and average momentary evidence, and the height of the bound. In the time-limited integration model the momentary sensory evidence is accumulated for a fixed duration and the decision is made based on the final accumulated evidence. This model has two free parameters: the proportionality constant and the integration period. The figure shows the changes of discrimination thresholds as a function of stimulus duration in the best fit of the two models (red). The bounded accumulation fits the data much better. A model comparison between the time-limited integration and bounded accumulation models yields a Bayes factor of 8×10^{-15} , indicating that the time-limited integration is inferior — values smaller than 0.1 indicate strong evidence against a model, and values smaller than 0.01 are considered decisive (Jeffreys 1961, *The Theory of Probability*). Similarly, the Bayes information criterion indicates a very large difference between the two models: 64.9. The time limited integration requires the discrimination thresholds stay similar for all durations longer than the integration period, a feature that is absent in the data. It is also important to note that the best fit of the time limited accumulation model is achieved for an integration period of 484.0 ms, which is much larger than the mean stimulus duration, 330 ms, or the median, 269 ms.

The fit of the time-limited integration model to data improves by assuming that the stopping time is not fixed, but variable. It is known that the sense of time passage follows

Weber law. The time-limited integration model with variability at stopping time has, three free parameters: a Weber fraction and the two parameters of the model without variability. The best fit is achieved for an average stopping time of 602.1 ms. This model is still inferior to the bounded accumulation model with two free parameters (Bayes factor = 0.001; Bayes information criterion = 13.3).




Figure S2
Kiani, Hanks & Shadlen

The impact of a constant motion pulse (3.2% coherence, 200-ms long) on the monkey's choice depends on both the time it is delivered and the strengths of the motion that it perturbs. In the figure, the trials for each motion strength were divided into 10 bins based on the pulse time. Each data point shows the difference of probability correct for pulses toward and away from the rewarded target. At times later than 400 ms the pulse has little effect on the monkey's choice, as demonstrated in the paper (Fig. 5). At intermediate times (e.g., ~200-300 ms), pulses that perturb intermediate and weak motion strengths lose their efficacy at different times with the effects persisting later for weaker motion strengths, consistent with a bounded accumulation of evidence. To test the statistical significance of this effect we performed the following logistic regression on the trials in which the pulse was delivered sooner than 400 ms. (The data were not binned for this analysis.)

$$P_{right} = \frac{1}{1 + e^{-\left(\alpha_1 + \alpha_2 I + \alpha_3 T_p + \alpha_4 C + \alpha_5 T_p I + \alpha_6 |C| I + \alpha_7 T_p |C| I\right)}}, \quad I = \begin{cases} 1 & \text{rightward pulse} \\ -1 & \text{leftward pulse} \end{cases}$$

where T_p is the time of the pulse and C is the signed motion strength (positive for rightward and negative for leftward motion). The null hypothesis is that the effect of pulse as a function of pulse time is not influenced with motion strength ($H_{07}: \alpha_7 = 0$). However, α_7 was significantly negative (-8.48 ± 4.28 , $p=0.02$) indicating that the effect of the pulse diminishes faster for stronger motion.




Figure S3
Kiani, Hanks & Shadlen

The positive correlation coefficients between T_{bc} and RT are not caused by the attrition of neural responses around the time of saccade. The figure shows the distribution of the expected correlation coefficients for 0% coherence trials when attrition is imposed on a distribution that has no real correlation (see below). The observed correlation in the data (red line) is significantly larger than the correlation expected by the attrition ($p < 0.001$).

The bound crossing time was detected in $>80\%$ of the trials, and its correlation with RT is, therefore, expected to be a good representative of the underlying process. However, the positive correlation observed in the data might be due to attrition of neural responses around the time of saccade. Attrition, when imposed on an uncorrelated 2D distribution, can cut out part of the distribution and produce an artifactual correlation. We used Monte Carlo methods to estimate the bias that would be caused by the attrition of the neural responses under the null hypothesis: zero correlation. For this analysis the response times and the bound crossing times were both measured relative to the *Go* signal. In the first step of the analysis, we found the bivariate normal distribution with zero covariance that maximized the likelihood of the observed cloud of data points under the censorship criterion dictated by the attrition: $T_{bc} \leq RT + 170$. The 170 ms was added to the RT because the neural traces extended 200ms after the saccade and, for the reliable detection of the bound crossing, we required the neural activity to stay above the bound for 30ms. This first step recovers the joint distribution of RT and T_{bc} under the null hypothesis of zero correlation. In the second step of the analysis, we sampled randomly from the retrieved bivariate normal distribution, respecting to the censorship criterion, and measured the expected correlation coefficient. For all motion strength, the correlation measured in our experiment was significantly larger than the correlation expected by the attrition ($p < 0.005$).