

Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex

Roозbeh Kiani and Michael N. Shadlen

The degree of confidence in a decision provides a graded and probabilistic assessment of expected outcome. Although neural mechanisms of perceptual decisions have been studied extensively in primates, little is known about the mechanisms underlying choice certainty. We have shown that the same neurons that represent formation of a decision encode certainty about the decision. Rhesus monkeys made decisions about the direction of moving random dots, spanning a range of difficulties. They were rewarded for correct decisions. On some trials, after viewing the stimulus, the monkeys could opt out of the direction decision for a small but certain reward. Monkeys exercised this option in a manner that revealed their degree of certainty. Neurons in parietal cortex represented formation of the direction decision and the degree of certainty underlying the decision to opt out.

Choice certainty—the degree to which a decision-maker believes a choice is likely to be correct—affects a variety of cognitive functions: how we plan subsequent actions, how we react and learn from mistakes, and how we justify our choices to others. Choice certainty is pivotal for planning actions in a complex environment in which subsequent decisions depend on pending outcomes of previous decisions (1–3). For example, a decision to undergo a risky operation depends, among other factors, on the degree of certainty that the diagnosis is correct. Psychologists have long proposed that choice certainty serves as a link between the physical world and belief: It provides a graded scale that allows us to translate our convictions into suitable actions (4, 5).

Despite the importance of choice certainty, its neural mechanisms are poorly understood. It is well established that choice certainty is closely correlated with both decision accuracy and reaction time (6–11). This close relationship suggests that the same mechanism that underlies the decision-making process might underlie certainty judgments (1, 10, 12, 13). It has been suggested that neurons in orbitofrontal cortex and cingulate cortex, which are known to represent reward expectation or conflict, represent reward expectations associated with decision uncertainty (14–16). However, these neurons do not give rise to a representation of decision uncertainty but presumably receive this information from neurons that compute this quantity in the decision-making process.

The neural mechanism of decision-making has been investigated using simple perceptual

tasks in which a monkey makes a categorical choice between two or more discrete options based on a sensory stimulus (17). When the monkey is required to report the perceived direction of motion by a saccadic eye movement, neurons in lateral intraparietal cortex (LIP) represent the accumulation of evidence, termed a decision variable, that supports the target in their response fields (RFs) (17–19). Furthermore, these neurons signal the termination of the decision process when their firing rates reach a critical level or bound (19–21). Theoretical and experimental studies raise the possibility that the neural computations approximate a form of probabilistic reasoning about the alternatives (22–24). We hypothesize that the graded, time-dependent firing rates of LIP neurons also represent choice certainty. Our hypothesis, therefore, unifies the representation of the three components of decisions—choice, reaction time, and certainty—in a single neural population.

Two monkeys made perceptual decisions about the net direction of motion in a dynamic random-dot display (Fig. 1A) (25). Task difficulty was controlled by varying both the percentage of coherently moving dots and the viewing duration. After a delay period, the fixation point was extinguished, which instructed the monkey to indicate its direction choice by making an eye movement to one of the direction-choice targets. On a random half of the trials, the monkey was given the option to abort the direction discrimination and to choose instead a small but certain reward associated with a third saccade target. This “sure target” was shown during the delay period, at least 500 ms after the random-dot motion was extinguished. During motion viewing, the monkey did not know whether the surebet option would arise. The task design, a form of postdecision wagering (26–28), ensured that the monkey made a decision about motion di-

rection on each trial. We hoped that the monkey would choose the sure target when less certain of the high-stakes direction choice, allowing us to study neural responses associated with choice certainty.

We first describe behavioral observations, which demonstrate that the postdecision wager reflects choice certainty. We then demonstrate a neural correlate of this certainty in the LIP firing rate. Together, these observations support a mechanism in which the same decision variable, represented by LIP neurons, underlies both the choice and the degree of certainty in that choice.

The monkeys opted for the sure target when the chance of making a correct decision about motion direction was small. They exercised this option more frequently for the weaker motion strengths and for the shorter stimulus durations ($P < 10^{-8}$) [equation 1 (25) and Fig. 1B], that is, when the probability of making an error was higher ($P < 10^{-8}$) [equation 2 (25) and Fig. 1C]. More interestingly, when the monkeys waived this option, the choice accuracy was better than on the trials when they were not offered the postdecision wager ($P < 10^{-3}$) [equation 3 (25) and Fig. 1C]. This improvement was apparent at almost all motion strengths and stimulus durations. It implies that the monkeys did not choose the sure target on the basis of stimulus difficulty but instead based on a sense of uncertainty on each trial. This same pattern was observed on a subset of trials in which identical random-dot patterns were repeated ($P < 0.025$) (fig. S1), which suggests that the source of information about difficulty is not governed solely by properties of the stimulus but also by internal variability that renders the evidence more or less reliable to the decision-maker.

We recorded extracellularly from 70 LIP neurons while the monkeys performed this task. These neurons exhibited spatially selective persistent activity that predicted whether an eye movement was planned into the RF of the neuron on a memory-guided saccade task (29–31). For the main motion task, we placed one of the direction targets (T_{in}) in the RF of the recorded neuron, the other direction target (T_{opp}) on the opposite side of the screen, and the sure target (T_s) orthogonal to the axis that connected the two direction targets.

Figure 2A shows responses of an example neuron for trials without the sure target. The neural activity after motion onset underwent a brief dip and then diverged to indicate the monkey's decision for T_{in} or T_{opp} . The activity persisted through the delay period until the eye movement (18). For simplicity, the graph combines all motion strengths and stimulus durations (25), but, as shown previously (18–21), the buildup of firing rate reflected the stimulus strength ($P = 0.01$) [equation 10 (25) and fig. S2], compatible with the representation of accumulated evidence in favor of T_{in} . We observed a similar divergence and persistence of activity for

Howard Hughes Medical Institute, National Primate Research Center, and Department of Physiology and Biophysics, University of Washington, Seattle, WA 98195, USA. E-mail: roozbeh@u.washington.edu (R. K.); shadlen@u.washington.edu (M.N.S.)

T_{in} and T_{opp} choices on the trials in which T_s was presented but was waived by the monkey (Fig. 2B, solid traces).

In contrast, when the monkey chose T_s , the activity after the motion changed more gradually and achieved intermediate values compared with the T_{in} and T_{opp} choices. This pattern persisted into the delay period until T_s appeared ($P < 10^{-8}$, t test). Before T_s appeared, the monkey did not know whether the sure bet would be offered. After the onset of T_s , there was a dip in activity, followed by a return to the level of activity preceding the onset of T_s . When the monkey chose T_s , the response gradually converged to the T_{opp} level. The profile of activity suggests that even before the onset of T_s , the neuron was informative about whether the monkey would choose or waive this option should it be offered.

We observed a similar pattern of activity across the population of 70 LIP neurons (32). Intermediate firing rates during motion viewing and the early delay were associated with choosing the sure target later in the trial, as shown by the population average firing rates (Fig. 2D). To quantify this effect in single neurons, we compared activity in the 200-ms period before T_s

onset (Fig. 2D, hatched box) for trials in which the monkey selected or waived the sure-bet option. For motion toward T_{in} , the neural activity across the population was significantly smaller for T_s choices than for T_{in} choices [$P = 0.007$, analysis of variance (ANOVA)] (Fig. 2E). For motion toward T_{opp} , the activity was significantly larger for T_s choices than for T_{opp} choices ($P = 0.001$) (Fig. 2F).

These observations demonstrate that the monkey is more likely to opt for T_s when the LIP activity achieves an intermediate level of firing rate. However, a possible concern is that the intermediate level of activity represented by the mean firing rates from many trials is an unfair representation of the activity on single trials. According to this argument, the intermediate means might represent a mixture of the high and low firing rates that would have corresponded to T_{in} and T_{opp} choices, had the monkey indicated a direction choice on these trials. This “mixture of states” alternative makes a clear prediction, which is not supported by the data. If the intermediate means were solely mixtures of the responses associated with T_{in} and T_{opp} choices, then the variance should reflect the dispersion of

values associated with these extremes. This idea is rejected: The variance associated with T_s choices was significantly smaller than the variance associated with the mixtures of T_{in} and T_{opp} choices ($P = 4.7 \times 10^{-5}$, F-test). We conclude that these intermediate levels of activity are not artifactual but represent a low state of certainty.

This conclusion is supported further by comparing the activity from neurons on single trials with the monkey’s decision to choose or waive the T_s option (Fig. 3A). For each trial, from each neuron, we calculated the deviation of firing rate, in the epoch just preceding T_s onset, from an intermediate level. The magnitude of this deviation was inversely related to the probability that the monkey chose the sure-bet option ($P = 2.3 \times 10^{-5}$) [equation 11 (25) and Fig. 3A]. The influence of a single neuron on the probability of a postdecision wager is expected to be small because it is but one member of a large population of neurons that govern the behavior, presumably (33–35). Nonetheless, the significance of the effect is a strong indication that LIP responses represent the choice certainty.

This single-trial analysis addresses another possible concern. Because stimulus difficulty (i.e., motion coherence and duration) affects both LIP responses and confidence judgments, it seems possible that the correlation between LIP activity and the postdecision wager is merely accidental, that is, totally explained by the stimulus difficulty. Alternatively, if T_s choices are based on LIP activity, they should be influenced by both the stimulus and the noisy fluctuations of LIP firing rates. To address this, we performed a variant of the single-trial analysis described in the previous paragraph. We calculated the trial-to-trial fluctuation of LIP responses relative to the mean response dictated by each motion strength and direction. These residual fluctuations before the sure-target onset had significant leverage on the probability of choosing the sure target ($P = 4.0 \times 10^{-5}$) [equation 12 (25)]. This finding also held for the subset of trials in which we used identical random-dot motion stimuli ($P = 0.015$). Therefore, the linkage of neural responses with sure-target choices is not explained merely by their shared covariation with the stimulus. We conclude from these analyses that the variable discharge of LIP activity was related to the monkey’s choice certainty, whether these variations were caused by experimental manipulations (i.e., motion strength and duration) or random effects (i.e., neural noise).

The single-trial analyses have focused thus far on neural activity in the delay period, immediately preceding the onset of T_s . Is the LIP activity during decision formation also related to choice certainty? The evolution of neural activity accompanying motion viewing suggests an affirmative answer. The rate of change of LIP activity after motion onset, termed the buildup rate (36), was related to the probability of choosing

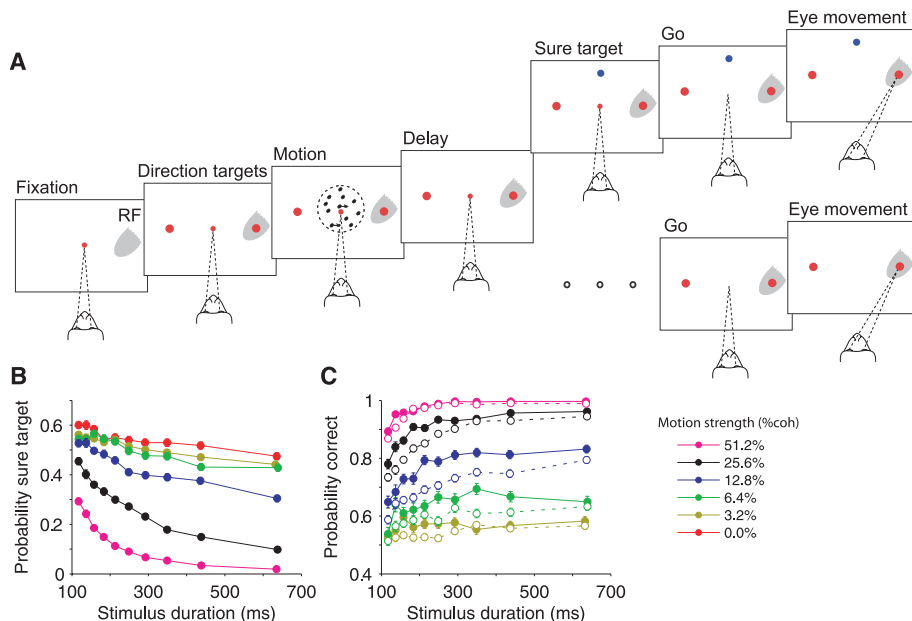


Fig. 1. Postdecision wagering behavior in monkeys is indicative of choice certainty. **(A)** The sequence of events in the task. After acquiring a central fixation point (small red circle), two direction targets (large red circles) appeared on the screen, one inside the neural RF (gray shading), the other on the opposite side of the screen. The motion stimulus appeared after a short delay, remained visible for 100 to 900 ms, and was followed by another delay (1200 to 1800 ms). On half of the trials (lower branch) the delay persisted until the fixation point was turned off, which served as a “go” signal that instructed the monkey to indicate the perceived direction of motion by a saccadic eye movement to one of the direction targets. A correct response led to a liquid reward; a wrong response led to no reward and a brief time-out. On the other half of the trials (upper branch), a third target was presented 500 to 750 ms after extinction of the motion. Choosing this sure target (T_s , blue circle) led to a smaller reward (~80% of correct reward). On these trials, the monkey could choose T_s or a direction choice. The two trial types were randomly interleaved. **(B)** The frequency of choosing T_s was greater when the motion strength (% coherence) was weak or the duration brief. The points are data, grouped in duration quantiles (deciles). Error bars (SE) are smaller than the symbols. **(C)** Decision accuracy when the T_s option was waived. The graph compares performance on trials in which T_s was not shown (open circles, dashed curves) with trials in which T_s was offered but waived (filled circles, solid curves).

T_s later in the trial. For stronger stimuli, the buildup was steeper ($P < 10^{-8}$) [equation 10 (25)], consistent with the accumulation of stronger


evidence, shorter decision times, and ultimately more accurate decisions (13, 17). According to our hypothesis, the buildup rates should tend

toward intermediate values when the monkey chose T_s . To test this, we performed a logistic regression analysis using buildup rates estimated from single trials. Deviation of the buildup rate from intermediate values was associated with a lower probability of choosing the sure target ($P = 0.017$) [equation 11 (25) and Fig. 3B]. Moreover, this link was not simply due to covariation of buildup rates and choice certainty with motion strength ($P = 0.0018$) [equation 12 (25)].

It is also interesting to note that, although the fluctuations in buildup rate and delay period activity were weakly correlated ($r = 0.10$, $P < 10^{-8}$), each exerted independent leverage on the likelihood that the monkey would opt for the T_s wager ($P < 0.03$) [equation 13 (25)]. In other words, both the evolution of decision-related activity and the sustained activity in the delay period carry information about choice certainty. Although both quantities reflect the state of evidence, variation in the buildup rate also affects the amount of time it takes to reach a decision (19–21, 37, 38), consistent with the long-held view that decision time contributes to choice certainty (8, 9, 12).

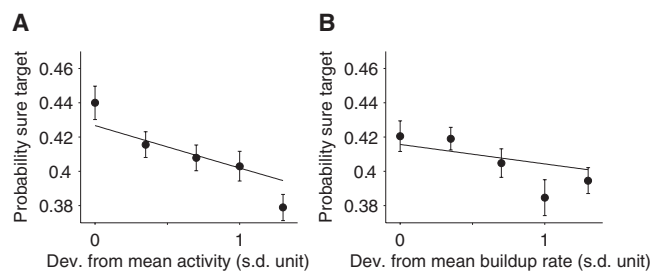
Indeed, a Bayesian framework that incorporates both evidence and decision time explains several aspects of the data. As previously shown, the left-right choices on this task are governed by the accumulation of evidence favoring one or the other option (17). This accumulation, which we call a decision variable, $v(t)$, is represented by the firing rates of LIP neurons. It begins at a neutral value and undergoes a random walk with drift (also termed drift diffusion) as evidence accumulates for and against the two direction alternatives. The decision terminates naturally when there is no more evidence (e.g., when the stimulus duration is short) or when v reaches a critical level or bound. In both cases, the choice is determined by the sign of v . As previously shown, this simple model explains the monkey's accuracy as a function of stimulus strength and viewing time. It explains both the diminishing returns associated with prolonged viewing in our experiment (fig. S3) (20) and the trade-off between speed and accuracy in reaction-time experiments (13, 39, 40). It also explains the saturating firing rate curves in Fig. 2.

A simple extension of this bounded evidence accumulation model also explains the postdecision wagering. The key insight is that both v and t convey information about certainty. Figure 4A shows the distribution of $v(t)$, combined across all stimulus strengths when the rewarded direction is, for example, rightward. Application of the decision rule described in the previous paragraph to $v(t)$ would lead to different proportions of correct and incorrect choices, depending on the motion strength. This transformation is shown in Fig. 4B, which replaces the probability distribution of $v(t)$ with the log odds of a correct decision. This is the log-posterior odds based solely on $v(t)$. For example, if a rightward stimulus is shown, the log odds of a



saccade initiation (right). (B) Responses from the same neuron on trials in which T_s was presented. The dashed lines show neural activity on trials in which T_s was chosen (black and gray, motion toward T_{in} and T_{opp} , respectively). The middle portion of the graph shows activity in the delay period, aligned to onset of T_s . (C and D) Population average responses of 70 LIP neurons from two monkeys. Same conventions as in (A) and (B). Firing rates from each neuron were normalized to the mean level before onset of the motion stimulus. (E) The activity before T_s presentation was smaller for T_s choices than for T_{in} choices. Each data point represents the mean activity of an LIP neuron in the 200 ms before T_s presentation [hatched rectangle in (D)]. Error bars, mean \pm SEM. Shading in the histogram shows significant cases ($P < 0.05$). The arrow shows the mean difference of normalized activity across the population (mean \pm SEM, -0.20 ± 0.03). (F) The activity before T_s presentation was larger for T_s choices than for T_{opp} choices. Same conventions as in (E) (mean difference = 0.18 ± 0.02).

Fig. 3. T_s choices were correlated with trial-to-trial variation of neural activity. Responses from single trials were represented as the absolute deviation, in units of standard deviation, from the mean value using all the trials from a neuron (z score). (A) The frequency of choosing T_s as a function of deviation from mean in the activity before T_s presentation. Curves are fits of equation 11 (25) to individual trials. The points illustrated on the graph were formed by grouping trials into five bins. (B) The frequency of choosing T_s as a function of deviations from the mean buildup rate of activity after motion onset. Same conventions as in (A).



The points illustrated on the graph were formed by grouping trials into five bins. (B) The frequency of choosing T_s as a function of deviations from the mean buildup rate of activity after motion onset. Same conventions as in (A).

correct choice is simply the log posterior odds that the stimulus is to the right

$$\underbrace{\text{Log} \frac{p(S_1|v(t))}{p(S_2|v(t))}}_{\text{log posterior odds}} = \underbrace{\text{Log} \frac{\sum_i p(v(t)|S_1, C_i) p(C_i|S_1)}{\sum_i p(v(t)|S_2, C_i) p(C_i|S_2)}}_{\text{log likelihood ratio}} + \underbrace{\text{Log} \frac{p(S_1)}{p(S_2)}}_{\text{log prior odds}}$$

where C is motion coherence, and S_1 and S_2 represent the rightward and leftward motion direction, respectively. The last term vanishes, because the prior probability that motion is left or right is equal. The summation terms implement marginalization over motion strength. The left side of the equation formalizes belief in the proposition $S = S_1$.

From the depiction in Fig. 4B, it is easy to imagine that opting out of the direction decision might happen when the expected chance of success based on $v(t)$ at decision time is less than a criterion level (Fig. 4C). This simple model explains the observed behavior and successfully predicts the amount of improvement in the probability of being correct for trials in which the monkey waives T_s . The model has only three free parameters (table S1), which were set by fitting the proportion of T_s choices and the probability of being correct for trials without T_s [dashed curves in Fig. 4D ($R^2 = 0.97$) and Fig. 4E ($R^2 = 0.98$)]. This establishes a prediction (not a fit) for the probability of being correct on the trials in which T_s was shown but waived (Fig. 4E, solid curves) ($R^2 = 0.95$). The agreement between this simple model and the data affirms the plausibility of the Bayesian sequential sampling framework (41).

Moreover, the evolution of $v(t)$ predicted by the model resembles qualitatively the responses of LIP neurons (Fig. 4F and fig. S4). For stronger motion, the decision variable associated with T_s choices follows less intermediate trajectories (note the separation of dashed curves), and the decision variable associated with direction choices rises (or falls) faster toward its plateau level. Both of these features are evident in the neural responses. The agreement is only approximate, presumably because neurons other than the ones we recorded contribute to the estimation of certainty (1, 22). These neurons might represent evidence for other directions of motion, but they are unlikely to represent the T_s choice option directly, as shown next.

To gain a better understanding of the representation of choice certainty across the population of LIP neurons, we recorded from 19 cells using the task configuration shown in Fig. 5. T_s was in the RF, whereas the direction-choice targets were not. Although T_s was not displayed until late in the delay period and only on half of the trials, its position was fixed throughout the course of the experiment. Nevertheless, these neurons

did not show a significant modulation of activity during the motion stimulus or in the ensuing delay period (42). Moreover, the weak activity that was present was uninformative about the choice to forego or choose the T_s option (Fig. 5B) ($P > 0.1$ for both directions of motion). Unlike the neurons with a direction choice target in the RF, the neurons that encode the location of T_s do not appear to represent choice certainty.

This observation argues against an alternative explanation of our finding based on allocation of attention to the T_s location. More generally, it provides additional evidence that

the monkey made a decision about the motion direction in the period preceding the onset of T_s , even on trials when it opted out of the direction task. There is no indication that the monkey approached the task as a choice between three alternatives, T_{in} , T_{opp} , and T_s . However, after the appearance of T_s , these neurons with a sure target in their RF became predictive of the postdecision wager. Although it is not obvious from the traces, the visual response in the first 200 ms was slightly larger when the monkey would choose T_s ($P < 0.01$, ANOVA), suggesting that T_s was more salient when there

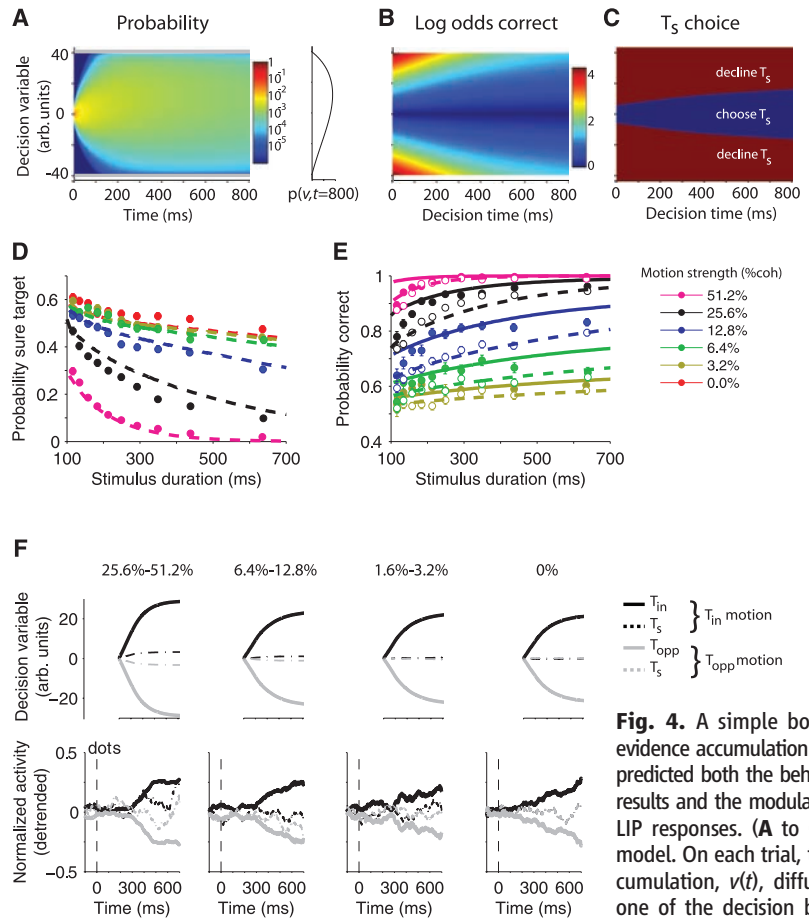


Fig. 4. A simple bounded evidence accumulation model predicted both the behavioral results and the modulation of LIP responses. (A to C) The model. On each trial, the accumulation, $v(t)$, diffuses to one of the decision bounds (gray lines). The process terminates when $v(t)$ reaches a bound or the stream of motion evidence ceases. (A) Representation of $v(t)$ as a propagating probability density, for all motion strengths, when the rewarded direction is rightward. Positive values for $v(t)$ represent accumulated evidence in favor of rightward. At time zero, the distribution is a delta function at $v = 0$. As time elapses, the range of $v(t)$ expands to fill the space between the two bounds, and there is a drift toward positive values, as shown by the probability density of v at $t = 800$ ms (inset to the right of color map). The distribution associated with leftward motion (not shown) is the mirror symmetric graph reflected about $v = 0$. (B) The log odds of a correct response based on the value of $v(t)$ at decision time. Correct responses are associated with larger v , but the relationship between v and probability correct changes with decision time. (C) T_s is chosen when the probability of a correct response is less than a criterion level. (D and E) Model fits and predictions. The three model parameters (table S1) were fit to the observed frequency of correct responses on trials in which T_s was not shown and the observed frequency of T_s choices on trials in which T_s was shown. These parameters predict the probability of a correct response on trials in which T_s was waived [solid curves in (E)]. (F) Comparison of model predictions and neural data. The average trajectory of $v(t)$ in the model was calculated for different coherence levels using the fit parameters. The calculation is based on the stimulus durations used in the experiment and assumes that v is fixed from termination of the accumulation process. The calculated trajectories (top) resemble the LIP responses (bottom). Neural responses were detrended by subtracting the mean response at each moment. Predictions are shifted by the neural latency (200 ms).

terminates when $v(t)$ reaches a bound or the stream of motion evidence ceases. (A) Representation of $v(t)$ as a propagating probability density, for all motion strengths, when the rewarded direction is rightward. Positive values for $v(t)$ represent accumulated evidence in favor of rightward. At time zero, the distribution is a delta function at $v = 0$. As time elapses, the range of $v(t)$ expands to fill the space between the two bounds, and there is a drift toward positive values, as shown by the probability density of v at $t = 800$ ms (inset to the right of color map). The distribution associated with leftward motion (not shown) is the mirror symmetric graph reflected about $v = 0$. (B) The log odds of a correct response based on the value of $v(t)$ at decision time. Correct responses are associated with larger v , but the relationship between v and probability correct changes with decision time. (C) T_s is chosen when the probability of a correct response is less than a criterion level. (D and E) Model fits and predictions. The three model parameters (table S1) were fit to the observed frequency of correct responses on trials in which T_s was not shown and the observed frequency of T_s choices on trials in which T_s was shown. These parameters predict the probability of a correct response on trials in which T_s was waived [solid curves in (E)]. (F) Comparison of model predictions and neural data. The average trajectory of $v(t)$ in the model was calculated for different coherence levels using the fit parameters. The calculation is based on the stimulus durations used in the experiment and assumes that v is fixed from termination of the accumulation process. The calculated trajectories (top) resemble the LIP responses (bottom). Neural responses were detrended by subtracting the mean response at each moment. Predictions are shifted by the neural latency (200 ms).

was greater choice uncertainty (43). The effect was weak (median difference = 7.7%), but as time elapsed during the remainder of the delay period the firing rates gave a clear indication of whether the monkey would choose the T_s .

Discussion. A connection between signal reliability, choice accuracy, and confidence has been proposed previously (1, 13, 14, 44, 45), but until now this connection has not been observed directly in the same neurons. Neurons in a variety of brain structures represent the size, preference, and probability of obtaining a reward (15, 46–54), but it is not known how these representations arise. The present results show that the same neurons that participate in decision formation (20, 55) carry the relevant signals for assigning the probability of obtaining a reward. It therefore seems likely that the computation of choice certainty is passed from LIP to brain structures that anticipate reward, and it is likely that feedback from these structures affects LIP in the epoch after the appearance of T_s to mediate the decision to choose or forgo the T_s option.

The mechanism underlying the representation of certainty in LIP is linked to the same evidence accumulation that underlies choice and decision time (17, 20). This accumulation is encoded in the firing rates of LIP neurons with RFs aligned to the choice targets representing the direction alternatives (18, 55–57). This is the decision variable, $v(t)$, that governs the choice of direction, having either attained a critical level—a decision termination bound—or by comparison to a criterion if the evidence stream ceases. This mechanism can be viewed as a merging of decision models based on sequential analysis (58–60) and signal detection

theory (61). The magnitude of this decision variable, combined with knowledge of elapsed time, maps directly to the probability of obtaining a reward. An associative-learning process based on LIP responses can therefore underlie the monkey's choice of T_s . The ability to explain the rich pattern of behavioral results and the qualitative agreement between model and physiology favors the simple conceptual model. It is probably also consistent with other models that exploit a broader population of LIP neurons to encode posterior probability (22, 41).

This simple mechanism brings certainty, which is commonly conceived as a subjective aspect of decision-making, under the same rubric as choice and reaction time (1, 62) and removes the need to resort to metacognitive explanations for certainty monitoring (45). Our findings support a low-level explanation of postdecision wagering in our task, but they do not preclude the possibility that an animal that experiences subjective awareness of degree of certainty might base such impressions on neural signals like the ones exposed here.

LIP neurons are hypothesized to encode the attentional salience or expected value of a visual saccade target (52, 53, 63), but these concepts cannot explain the pattern of LIP activity in our experiment. For example, a diversion of attention away from T_{in} to the potential location of T_s should have led to a reduction in firing rate for both T_{in} and T_{opp} directions during motion viewing and in the delay period before T_s appeared. Attention (or motor planning) might explain the activity just preceding saccades, but it does not explain the intermediate firing rates, particularly for T_{opp} directions, in the key epochs of interest.

A second alternative, expected reward, seems more plausible, at least to the extent that it mimics the belief that a choice will be correct. However, the expected value of T_{in} , in the objective sense (from economics), changes as a function of motion strength (psychometric function) (Fig. 1C), whereas the firing rate before T_s onset is minimally affected by motion strength when the monkey waives T_s (Fig. 4F and fig. S4). Even subjective expected value, which is synonymous with certainty, fails to capture fully the deeper insight our experiment reveals about mechanism: The evolution of decision-related activity that gives rise to a choice also underlies certainty and a wager based upon it.

A famous controversy in the history of probability theory concerned whether it is meaningful to embrace the truth of a hypothesis as a graded quantity expressed as a probability or whether, instead, hypotheses are simply true or false. The latter approach led frequentists to reject the Bayesian concept of degree of belief, relegating probability to the analysis of error rates in assertions of truth (64, 65). Our finding suggests that when the brain embraces a truth, it does so in a graded way so that even a binary choice leaves in its wake a quantity that represents degree of belief. From this perspective, our neural recordings support the idea of a “Bayesian brain” (66) and a neural mechanism of decision-making that does not flip into a fixed point or attractor state but instead approximates the formation of a probability distribution (41, 67). Accordingly, the intermediate levels of activity associated with less certain choices might be a sign of a more homogeneous level of activity across the population of neurons. Fundamentally, our results advance understanding of the neural mechanisms that underlie decision-making by coupling for the first time the mechanisms leading to decision formation and the establishment of a degree of confidence.

References and Notes

1. D. Vickers, *Decision Processes in Visual Perception* (Academic Press, New York, 1979).
2. N. D. Daw, Y. Niv, P. Dayan, *Nat. Neurosci.* **8**, 1704 (2005).
3. P. Dayan, N. D. Daw, *Cogn. Affect. Behav. Neurosci.* **8**, 429 (2008).
4. F. B. Sumner, *Psychol. Rev.* **5**, 616 (1898).
5. W. McDougall, *Psychol. Rev.* **28**, 315 (1921).
6. V. C. A. Henmon, *Psychol. Rev.* **18**, 186 (1911).
7. C. S. Pierce, J. Jastrow, *Proc. Natl. Acad. Sci. U.S.A.* **3**, 75 (1884).
8. J. Volkman, *Psychol. Bull.* **31**, 672 (1934).
9. D. M. Johnson, *Arch. Psychol.* **34**, 1 (1939).
10. W. M. Petrusic, J. V. Baranski, *Psychon. Bull. Rev.* **10**, 177 (2003).
11. D. Vickers, P. Smith, *Perception* **14**, 471 (1985).
12. R. J. Audley, *Br. Med. Bull.* **20**, 27 (1964).
13. S. W. Link, *The Wave Theory of Difference and Similarity*, Scientific Psychology Series (Erlbaum, Hillsdale, NJ, 1992).
14. A. Kepecs, N. Uchida, H. A. Zariwala, Z. F. Mainen, *Nature* **455**, 227 (2008).
15. C. Padoa-Schioppa, J. A. Assad, *Nature* **441**, 223 (2006).
16. B. Y. Hayden, A. C. Nair, A. N. McCoy, M. L. Platt, *Neuron* **60**, 19 (2008).
17. J. I. Gold, M. N. Shadlen, *Annu. Rev. Neurosci.* **30**, 535 (2007).

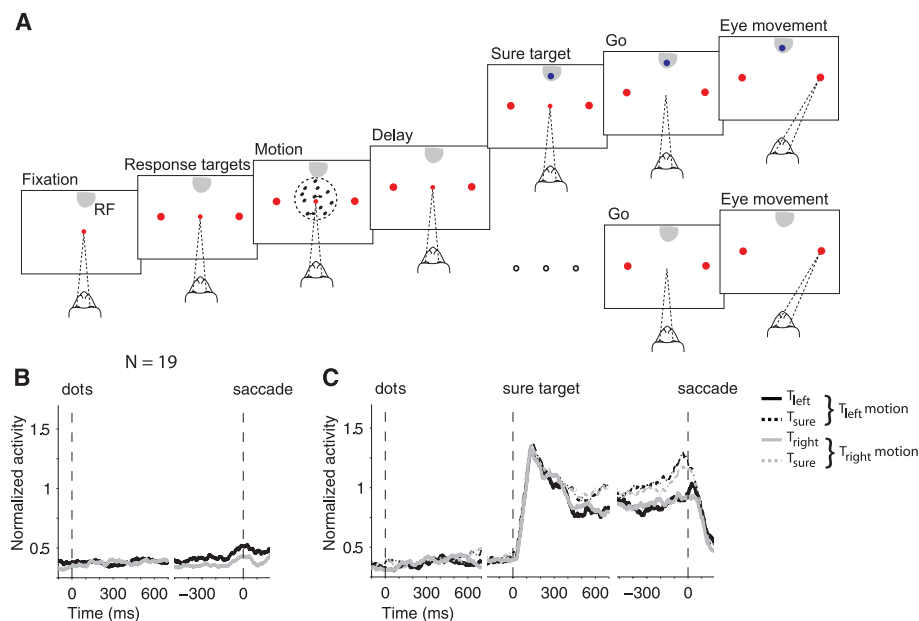


Fig. 5. Activity of LIP neurons when the location of T_s was in the RF. **(A)** Task design. For 19 neurons from two monkeys, we placed T_s in the RF. The high-stakes direction targets were outside the RF. Task sequence was otherwise unchanged. **(B)** Responses on trials in which T_s was not offered. **(C)** Responses on trials in which T_s was presented. Firing rates were normalized to the visual activity in the 300-ms epoch after onset of T_s .

18. M. N. Shadlen, W. T. Newsome, *J. Neurophysiol.* **86**, 1916 (2001).
19. J. D. Roitman, M. N. Shadlen, *J. Neurosci.* **22**, 9475 (2002).
20. R. Kiani, T. D. Hanks, M. N. Shadlen, *J. Neurosci.* **28**, 3017 (2008).
21. A. K. Churchland, R. Kiani, M. N. Shadlen, *Nat. Neurosci.* **11**, 693 (2008).
22. J. M. Beck *et al.*, *Neuron* **60**, 1142 (2008).
23. J. I. Gold, M. N. Shadlen, *Trends Cogn. Sci.* **5**, 10 (2001).
24. T. Yang, M. N. Shadlen, *Nature* **447**, 1075 (2007).
25. Materials and methods are available as supporting material on Science Online.
26. N. Persaud, P. McLeod, A. Cowey, *Nat. Neurosci.* **10**, 257 (2007).
27. W. E. Shields, J. D. Smith, D. A. Washburn, *J. Exp. Psychol. Gen.* **126**, 147 (1997).
28. N. Kornell, L. K. Son, H. S. Terrace, *Psychol. Sci.* **18**, 64 (2007).
29. C. L. Colby, M. E. Goldberg, *Annu. Rev. Neurosci.* **22**, 319 (1999).
30. J. W. Gnadt, R. A. Andersen, *Exp. Brain Res.* **70**, 216 (1988).
31. M. L. Platt, P. W. Glimcher, *J. Neurophysiol.* **78**, 1574 (1997).
32. The activity of each neuron was normalized to the average response in the 300 ms preceding the motion, that is, the period after the appearance of the direction choice targets in the RF.
33. A. J. Parker, W. T. Newsome, *Annu. Rev. Neurosci.* **21**, 227 (1998).
34. E. Zohary, M. N. Shadlen, W. T. Newsome, *Nature* **370**, 140 (1994).
35. M. N. Shadlen, K. H. Britten, W. T. Newsome, J. A. Movshon, *J. Neurosci.* **16**, 1486 (1996).
36. The buildup rate on each trial was calculated by fitting a line to the neural activity in a ~300-ms window starting at the dip of activity after motion onset (25).
37. E. P. Cook, J. H. Maunsell, *Nat. Neurosci.* **5**, 985 (2002).
38. D. P. Hanes, J. D. Schall, *Science* **274**, 427 (1996).
39. J. Palmer, A. C. Huk, M. N. Shadlen, *J. Vis.* **5**, 376 (2005).
40. D. Vickers, J. Packer, *Acta Psychol. (Amst.)* **50**, 179 (1982).
41. W. J. Ma, J. M. Beck, P. E. Latham, A. Pouget, *Nat. Neurosci.* **9**, 1432 (2006).
42. The activity of each neuron was normalized by its (visual) response in the 300 ms after T_r onset.
43. J. W. Bisley, M. E. Goldberg, *Science* **299**, 81 (2003).
44. J. N. Kim, M. N. Shadlen, *Nat. Neurosci.* **2**, 176 (1999).
45. J. D. Smith, M. J. Beran, J. J. Couchman, M. V. Coutinho, *Psychon. Bull. Rev.* **15**, 679 (2008).
46. M. Watanabe, *Nature* **382**, 629 (1996).
47. A. Izquierdo, R. K. Suda, E. A. Murray, *J. Neurosci.* **24**, 7540 (2004).
48. R. Kawagoe, Y. Takikawa, O. Hikosaka, *Nat. Neurosci.* **1**, 411 (1998).
49. M. I. Leon, M. N. Shadlen, *Neuron* **24**, 415 (1999).
50. L. Tremblay, W. Schultz, *Nature* **398**, 704 (1999).
51. J. D. Wallis, E. K. Miller, *Eur. J. Neurosci.* **18**, 2069 (2003).
52. L. P. Sugrue, G. S. Corrado, W. T. Newsome, *Science* **304**, 1782 (2004).
53. M. L. Platt, P. W. Glimcher, *Nature* **400**, 233 (1999).
54. M. A. Belova, J. J. Paton, C. D. Salzman, *J. Neurosci.* **28**, 10023 (2008).
55. T. D. Hanks, J. Ditterich, M. N. Shadlen, *Nat. Neurosci.* **9**, 682 (2006).
56. A. C. Huk, M. N. Shadlen, *J. Neurosci.* **25**, 10420 (2005).
57. C. T. Law, J. I. Gold, *Nat. Neurosci.* **11**, 505 (2008).
58. A. Wald, *Sequential Analysis* (Wiley, New York, 1947).
59. D. R. J. Laming, *Information Theory of Choice Reaction Time* (Wiley, New York, 1968).
60. R. D. Luce, *Response Times: Their Role in Inferring Elementary Mental Organization* (Oxford University Press, Belfast, 1986).
61. D. M. Green, J. A. Swets, *Signal Detection Theory and Psychophysics* (Wiley, New York, 1966).
62. P. L. Smith, *J. Math. Psychol.* **32**, 135 (1988).
63. J. Gottlieb, *Neuron* **53**, 9 (2007).
64. E. T. Jaynes, *Probability Theory: The Logic of Science*, G. L. Bretthorst, Ed. (Cambridge University Press, Cambridge, 2003).
65. D. Howie, *Interpreting Probability: Controversies and Developments in the Early Twentieth Century* (Cambridge University Press, Cambridge, 2007).
66. D. C. Knill, A. Pouget, *Trends Neurosci.* **27**, 712 (2004).
67. R. S. Zemel, P. Dayan, A. Pouget, *Neural Comput.* **10**, 403 (1998).
68. This work was supported by the Howard Hughes Medical Institute, National Eye Institute grant EY11378, and National Center for Research Resources grant RR00166. We thank T. Hanks, A. Pouget, A. Churchland, D. Lee, P. Phillips, J. Palmer, and C. D. Salzman for helpful discussions and comments and A. Boulet and K. Ahl for technical assistance.

Supporting Online Material

www.sciencemag.org/cgi/content/full/324/5928/759/DC1

Materials and Methods

Figs. S1 to S4

Table S1

References

5 December 2008; accepted 11 March 2009

10.1126/science.1169405

REPORTS

Characterization of Multipartite Entanglement for One Photon Shared Among Four Optical Modes

Scott B. Papp,^{1*} Kyung Soo Choi,^{1*} Hui Deng,² Pavel Lougovski,³ S. J. van Enk,³ H. J. Kimble^{1†}

Access to genuine multipartite entanglement of quantum states enables advances in quantum information science and also contributes to the understanding of strongly correlated quantum systems. We report the detection and characterization of heralded entanglement in a multipartite quantum state composed of four spatially distinct optical modes that share one photon, a so-called W state. By randomizing the relative phase between bipartite components of the W state, we observed the transitions from four- to three- to two-mode entanglement with increasing phase noise. These observations are possible for our system because our entanglement verification protocol makes use of quantum uncertainty relations to detect the entangled states that span the Hilbert space of interest.

Investigations of entanglement for two quantum systems have answered many fundamental questions in quantum physics (1, 2) and revealed powerful new capabilities of quantum mechanics within the field of quantum information science

(3–5). Many of these advances have used well-tested methods for the characterization of quantum entanglement in bipartite (two-component) systems (6, 7). Entangled states of more than two systems enhance our knowledge of quantum theory, because new classes of states are available (7–9). Beyond applications to conventional quantum computation (3), exotic multipartite states have emerged as crucial resources for new directions in quantum information processing such as measurement-based quantum computation (10, 11), quantum secret sharing (12), and quantum simulation (13). Despite

the extraordinary promise that they offer, unambiguously detecting multipartite entangled states is still a major challenge from both an experimental and a theoretical standpoint.

Genuine N -partite entanglement is realized only with the simultaneous participation of all N of the constituent systems. The exponential increase with N in the amount of information required to describe the overall quantum system, although exceedingly beneficial for large-scale quantum information protocols (3), makes the task of classifying (8, 9) and detecting such entangled states extremely difficult (7). Still, there are prescribed methods to detect entanglement in select classes of multipartite states that generally rely on reconstructing the density matrix $\hat{\rho}$. Linear entanglement witnesses supplemented by tomography of $\hat{\rho}$ have been used to detect entanglement in six (14) and eight (15) atomic ions, as well as for hyperentangled photons (16). A serious drawback of quantum-state tomography is the prohibitive number of measurements and their accuracies that are required with increasing N .

Our work focuses on a specific class of quantum states in which exactly one photon is coherently shared among N distinct optical modes in the form of

$$|W\rangle = \frac{1}{2} [(|1000\rangle + e^{i\phi_1} |0100\rangle) + e^{i\phi} (|0010\rangle + e^{i\phi_2} |0001\rangle)] \quad (1)$$

shown here for $N=4$ and with the relative phases ϕ , ϕ_1 , ϕ_2 of the modes. This is a so-called W state,

¹Norman Bridge Laboratory of Physics 12-33, California Institute of Technology, Pasadena, CA 91125, USA. ²Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA. ³Department of Physics, University of Oregon, Eugene, OR 97403, USA.

*These authors contributed equally to this work.

†To whom correspondence should be addressed. E-mail: hjkimble@caltech.edu

Supplementary Online Material

Materials and Methods

Behavioral task

We measured the choice certainty in a direction discrimination task using a post-decision wagering design (Fig. 1A). Two adult rhesus monkeys (*Macaca mulatta*) were trained to judge the net direction of a field of dynamic random dots ($S1$, $S2$). Task difficulty was controlled by varying the viewing duration (100-900 ms, truncated exponential distribution) and the percentage of coherently moving dots (motion strength) as described previously (20, $S3$). After extinction of the motion display, the monkeys held fixation through a delay period (1200-1800 ms) before extinction of the fixation point. On a random half of the trials, a third target (sure target, T_s) appeared at a random time during the delay period (500-750 ms after motion offset) and stayed on through the delay period. After extinction of the fixation point, the monkey indicated its choice by making a saccadic eye movement either to one of the two direction targets or to T_s , if present. The monkey received a liquid reward for a correct choice, and nothing plus a short timeout for an error. Choosing T_s , when present, always yielded a reward, but the reward size was smaller than that for a correct response. The reward ratio was adjusted to encourage the monkey to choose the sure target on nearly half of all trials in order to increase the statistical power of analyses on neural responses. The sure bet allowed the monkey to opt out of the high stakes decisions about motion direction.

We collected 150,558 trials (45,989 psychophysics and electrophysiology trials, and 104,569 psychophysics trials) from two monkeys in 200 experimental sessions. Across the sessions, the ratio of reward size for the sure bet and the high-stakes decisions was 0.79 ± 0.07 (median=0.81). On 114 sessions half of the trials for each motion strength shared a fixed seed for the random number generator. The fixed seed guaranteed a fixed sequence of motion by eliminating trial-to-trial variation of the stimulus frames. We collected 83,809 trials in these sessions (35,023 psychophysics and electrophysiology trials, and 48,786 psychophysics trials).

In electrophysiology sessions, we recorded the activity of single neurons in the ventral division of area LIP (LIPv) ($S4$) while the monkey performed the task. Standard electrophysiology techniques were used for these extracellular recordings (20, 21). LIPv was identified by three characteristic markers: (1) anatomical location, confirmed by stereotactical location and by registering the monkey's structural MRI scans with a high resolution "flattened" scan distributed with the CARET software package ($S5$); (2) transition of white and gray matter during recordings; and (3) sustained activity of neurons in delayed and memory guided saccade tasks ($S6$). We screened 81 neurons with sustained activity in the vicinity of LIPv. All but 10 neurons were included in the analysis. The excluded neurons did not exhibit clear selectivity during the decision making. All training and data collection 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.

Analysis of behavioral data

We used logistic regression to examine the effects of motion strength (c ; range 0 to 0.512) and duration (t ; 0.1 to 0.9 s) on the probability of choosing T_s :

$$P_{sure} = \left[1 + e^{-(\beta_0 + \beta_1 c + \beta_2 t)} \right]^{-1} \quad (1)$$

where β_i are fitted coefficients determined using maximum likelihood (binomial error).

The same equation describes the effect of the stimuli on the probability of a correct choice:

$$P_{cor} = \left[1 + e^{-(\beta_0 + \beta_1 c + \beta_2 t)} \right]^{-1} \quad (2)$$

To assess the improvement in accuracy when T_s was presented but waived, we extended Equation 2 as follows:

$$P_{cor} = \left[1 + e^{-(\beta_0 + \beta_1 c + \beta_2 t + \beta_3 I + \beta_4 c I + \beta_5 t I)} \right]^{-1} \quad I = \begin{cases} 1, & T_s \text{ shown} \\ 0, & T_s \text{ not shown} \end{cases} \quad (3)$$

where I is an indicator variable. The null hypothesis is that probability correct does not change in the presence of T_s ($H_0 : \beta_{3-5} = 0$). This analysis supports our claim that by waiving T_s , accuracy improved at all motion strengths and durations ($\beta_3 = 0.12 \pm 0.03$, $p = 0.00026$; $\beta_4 = 1.28 \pm 0.14$, $p < 10^{-8}$; $\beta_5 = 0.88 \pm 0.33$, $p = 0.008$). For the last term, the analysis was limited to viewing durations under 300 ms, which is the range in which t has the largest effect on the probability correct (Fig. 1C) (β_5 was not significantly different from zero when longer durations were included).

To make Figures 1B-C and S1, trials were divided into 10% quantiles (deciles) based on stimulus duration for each motion strength. The probability of choosing T_s was calculated for each decile. Also, we calculated the probability of choosing the correct direction target when the monkey did not choose the sure target. We pooled 1.6% and 3.2% motion coherence to avoid clutter at the lower half of the figures, and because 1.6% coherence was tested only on some sessions. Note that all the statistical analyses were performed on individual trials, not on the deciles.

Bounded accumulation model

In this model, the discrimination of rightward and leftward motion directions is based on integration of noisy sensory evidence by two accumulators, which correspond to the two direction choices. The accumulated evidence is termed decision variable. The process continues until one of the accumulators reaches a fixed bound, or until the stream of motion evidence terminates, whichever happens first. The duration of the accumulation process is termed decision time. On trials in which T_s is not presented, the choice is

dictated by the accumulator that hits the bound first, or has the larger decision variable. The mechanism thus resembles a race between two accumulators. We can simplify the process further by assuming that a piece of evidence in favor of one of the options is against the other option. By this assumption, the fluctuations of decision variable in the two accumulators become anti-correlated and the race model reduces to a one-dimensional diffusion process in which the decision variable fluctuates in the space between the two decision bounds (13, 39, S7).

The propagation of the probability density of decision variable over time can be calculated using the Fokker-Planck equation (S8):

$$\frac{\partial p(v,t)}{\partial t} = \left[-\frac{\partial}{\partial v} \mu(v,t) + \frac{\partial^2}{\partial v^2} D(v,t) \right] p(v,t) \quad (4)$$

where $D(v,t)$ is the diffusion coefficient ($D(v,t) > 0$) and $\mu(v,t)$ the drift or advective coefficient. In the context of bounded accumulation models, $p(v,t)$ represents the probability of decision variable (v) at time t , $\mu(v,t)$ the strength of momentary evidence and $D(v,t)$ the variance of the momentary evidence. Based on previous studies (S9, S10) we assume that both the advective coefficient and the diffusion coefficient are stationary:

$$\begin{aligned} \mu(v,t) &= kc \\ D(v,t) &= 1 \end{aligned} \quad (5)$$

where k is a constant. The boundary conditions of the Fokker-Planck equation are:

$$p(v, t_0) = \delta(v) \quad (6)$$

and

$$p(\pm B, t) = 0 \quad (7)$$

which enforce the constraints that the initial value of the decision variable is zero and that the accumulation terminates whenever the decision variable reaches one of the bounds ($\pm B$).

In this model, there is an implicit mapping between the probability of a correct response and the value of decision variable. If $v(t)$ denotes the decision variable at decision time, the log posterior odds is

$$\text{Log} \frac{p(S_1 | v(t))}{p(S_2 | v(t))} = \text{Log} \frac{\sum_i p(v(t) | S_1, C_i) p(C_i | S_1)}{\sum_i p(v(t) | S_2, C_i) p(C_i | S_2)} + \text{Log} \frac{p(S_1)}{p(S_2)} \quad (8)$$

where S_1 and S_2 are the two motion directions, and C_i are the set of coherences used in the experiment. The two motion directions happen with equal frequency, and motion coherence and direction vary independently. Equation 8, therefore, is reduced to:

$$\text{Log} \frac{p(S_1 | v(t))}{p(S_2 | v(t))} = \text{Log} \frac{\sum_i p(v(t) | S_1, C_i) p(C_i)}{\sum_i p(v(t) | S_2, C_i) p(C_i)} \quad (9)$$

where $p(C_i)$ for $C_i = 0$ is half of other coherence levels. On trials in which T_s is presented, the model recommends choosing T_s when the absolute value of log posterior odds is less than a criterion level (θ).

Overall, there are three free parameters in the model: k , B , and θ . k and B are both expressed in units of the diffusion coefficient, and thus have arbitrary units. We fit these free parameters by maximizing the likelihood of the observed frequency of correct responses on trials in which T_s was not shown, and the observed frequency of choosing T_s on trials in which T_s was shown. The parameters were then used to predict the frequency of correct responses on trials in which the monkey waived T_s (solid lines in Fig. 4E). The fits were performed on individual trials using binomial error terms. The R^2 values for the goodness of fits or predictions were computed for the data points (deciles) shown on the Figures 4D-E.

To create the predicted trajectories of the decision variable (Fig. 4F), we performed a Monte-Carlo simulation of the diffusion process. The distribution of stimulus duration for the simulated trials was similar to that used in the experiment. To average the trajectory of the decision variable across the simulated trials, we assumed that the value of the decision variable at decision time is retained after the termination of the process and during the delay period that separates the motion stimulus from T_s onset.

Analysis of neural data

Peri-stimulus time histograms (PSTHs) shown in Figures 2, 4, 5, S2 and S4 were smoothed by convolution of neural responses with a 100 ms box car filter. However, all statistical analyses were performed on spike counts without any smoothing. To create population PSTHs in Figures 2, 4, S2, and S4 we normalized the activity of each neuron to the average firing rate of all trials in the 300 ms period before dots onset, that is the period after onset of the direction choice targets. For Figure 4F, the PSTHs were detrended by subtracting from each PSTH the mean of the correct T_{in} and T_{opp} PSTHs. For the population PSTH in Figure 5 (T_s in RF), responses were normalized to the average firing rate in the 300 ms period after T_s onset.

Single trial analyses were based on two measurements. The first was the neural activity calculated in a 200 ms window ending at T_s onset. The second was the rate of buildup in neural activity following motion onset. To calculate the buildup rate we first identified a relevant window for each cell, and then fit a line to neural responses in that window on each trial. The window started at the nadir (dip) of activity following motion onset and ended when the average firing rate achieved a steady level. The time of dip for each neuron was calculated as the time of minimum activity in 100-500 ms after motion onset (mean=195.8±9.4 ms). The beginning of steady level of activity was calculated for each cell by fitting the following ramp-to-plateau function after the dip to the average response for correct T_{in} choices.

$$r = \gamma_0 + \gamma_1 \min\{t, t_p\}$$

where r is the spike rate, t is time relative to the dip, and t_p is the beginning of the sustained level of activity; γ_i and t_p are free parameters. t_p , which defines the width of the window, was 299.9 ± 7.6 ms across the population. Where necessary we trimmed the window to avoid overlap with the window used for the calculation of firing rate before T_s onset.

To test whether the buildup rate of activity is modulated by coherence we used the following regression

$$\dot{r} = \beta_0 + \beta_1 c \quad (10)$$

where \dot{r} is buildup rate of activity for individual trials with correct T_{in} choices, and c is motion coherence. The null hypothesis is that buildup rate is not modulated by motion coherence ($H_0 : \beta_1 = 0$). When applied to the population of cells, buildup rates of the trials belonging to individual cells were first standardized by calculating z-scores (subtracting the mean and dividing by standard deviation).

We quantified the change of activity associated with the post-decision wager by comparing firing rates in the 200 ms epoch before onset of T_s on trials in which the monkey chose or waived this option. We show means and standard errors in Figures 2E-F and apply 2-sided t-tests to illustrate significant cases ($p < 0.05$). Population summary statistics were based on two-way ANOVA with choice and cell identity as fixed and random effects, respectively.

We examined the variance of activity in the same 200 ms epoch before onset of T_s . If the mean firing rates associated with T_s choices arose from a mixture of values like those associated with the T_{in} and T_{opp} choices, then the ratio of the variance for all T_s choices is predicted to be equal to the variance of the pooled T_{in} and T_{opp} responses. Alternatively, if the mean represents the central tendency of an intermediate distribution, the variance should be smaller than this prediction. We performed an F-test to compare the variance of activity for T_s choices and the mixture of T_{in} and T_{opp} choices. For this analysis, neural responses on individual trials were standardized for each neuron and pooled across cells. Similar results were obtained by comparing the geometric mean of the variance ratios from each neuron to unity ($p = 0.0004$, t-test on log transformed ratios). As a further test, we also used Monte Carlo methods to recreate the average delay activity for T_s choices by sampling, with appropriate proportions, from the distribution of activity for T_{in} and T_{opp} choices (24). The simulated distributions had larger variance than the observed distribution both for trials with T_{in} motion direction ($p < 0.005$) and for trials with T_{opp} motion direction ($p < 10^{-4}$).

To test how the deviation of neural activity from intermediate values on individual trials influenced the probability of choosing the sure target we used the following logistic regression:

$$P_{sure} = \left[1 + e^{-(\beta_0 + \beta_1 |z|)} \right]^{-1} \quad (11)$$

where $|z|$ is the magnitude (absolute value) of the standardized firing rate measured in the 200 ms epoch preceding onset of T_s . The standardization was applied to each neuron

using the mean and standard deviation across all trials, regardless of outcome. The logistic regression was applied to the population of z-transformed single-trial responses. The fitted coefficient β_1 quantifies the degree that deviation from intermediate firing rates affects log odds of choosing T_s . Our hypothesis is that larger deviations are associated with a lower probability of T_s choices ($H_0 : \beta_1 \geq 0$). We used a similar logistic regression for the effect of buildup rate, using standardized values of buildup rate from the same trials.

To test whether the single-trial effects were explained by motion strength and duration, we expanded Equation 11 to incorporate these terms and their interaction:

$$P_{sure} = \left[1 + e^{-(\beta_0 + \beta_1|z| + \beta_2c + \beta_3t + \beta_4ct)} \right]^{-1} \quad (12)$$

The null hypothesis remains ($H_0 : \beta_1 = 0$).

To test if buildup rate during the decision making process conveyed information about the choice of sure target beyond what was conferred by the neural activity prior to sure target onset we used the following logistic regression.

$$P_{sure} = \left[1 + e^{-(\beta_0 + \beta_1|z_1| + \beta_2|z_2|)} \right]^{-1} \quad (13)$$

where z_1 and z_2 are standardized firing and buildup rates, as above. The null hypothesis is ($H_0 : \beta_2 = 0$).

Supplementary Figures

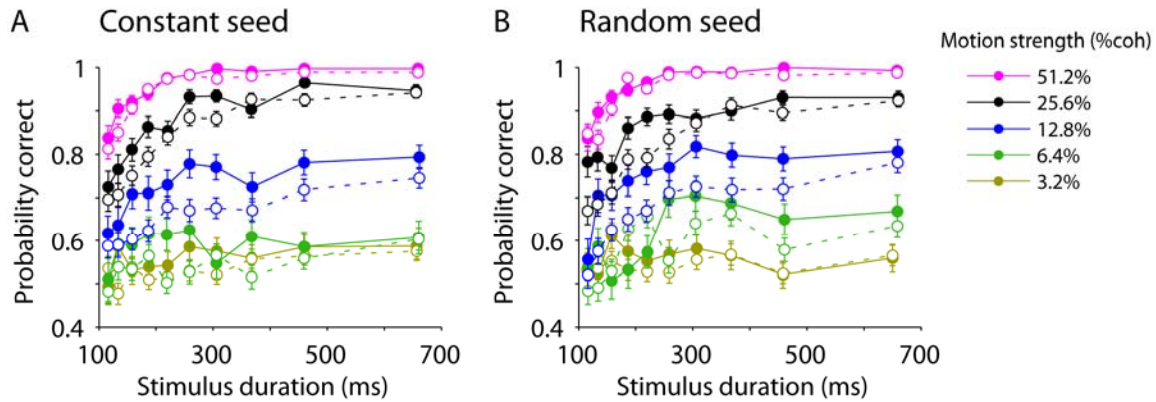


Figure S1.

Post-decision wagering reflects variability of neural signals as well as variation in stimulus difficulty. Here we show that the monkeys performed the task similarly on a subset of trials that used motion stimuli comprised of identical sequences of random dots. In some experimental sessions we used two trial types, which were randomly interleaved. For both trial types the probability of a correct response was larger when the monkey waived the sure target and chose the high stakes targets than when the monkey was not provided with the sure target. (A) Constant seed trials. An identical sequence of motion frames were shown for each motion strength. The trial-to-trial variability of stimulus was thus eliminated for these trials. (B) Random seed trials. The sequence of motion frames varied randomly from trial to trial.

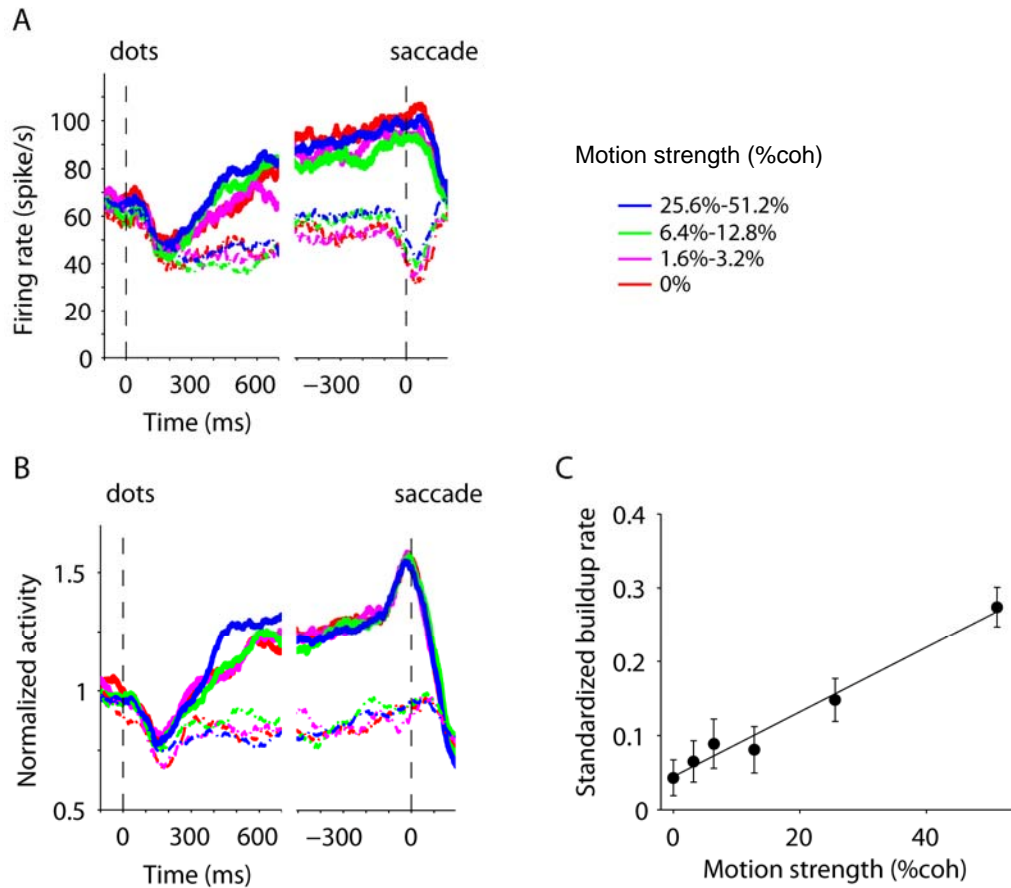


Figure S2.

The buildup of neural responses after motion onset was modulated by motion strength. (A) Responses of the example LIP neuron from Figure 2A-B. All conventions are similar to Figure 2A, except that trials are grouped by motion strength. After motion onset, the buildup of activity for T_{in} choices was larger for stronger motion stimuli. (B) Responses of the population of 70 LIP neurons. Same conventions as in A. The coherence-dependent increase in buildup rate is less evident because of the abundance of trials with short motion durations combined with variation in latency and magnitude of the buildup across neurons. These limitations are addressed in the next panel. (C) Standardized buildup rate across the population of neurons as a function of motion strength. Correct T_{in} choices, and all T_{in} choices for 0% motion strength, are included in this analysis. Buildup rates were calculated for single trials, as described in Methods, and then expressed as deviation from the mean in units of standard deviation. This procedure (standardization) was performed for each cell before pooling the data. The increase in buildup rate is well-described by a line fitted to the data. Error bars show SEM.

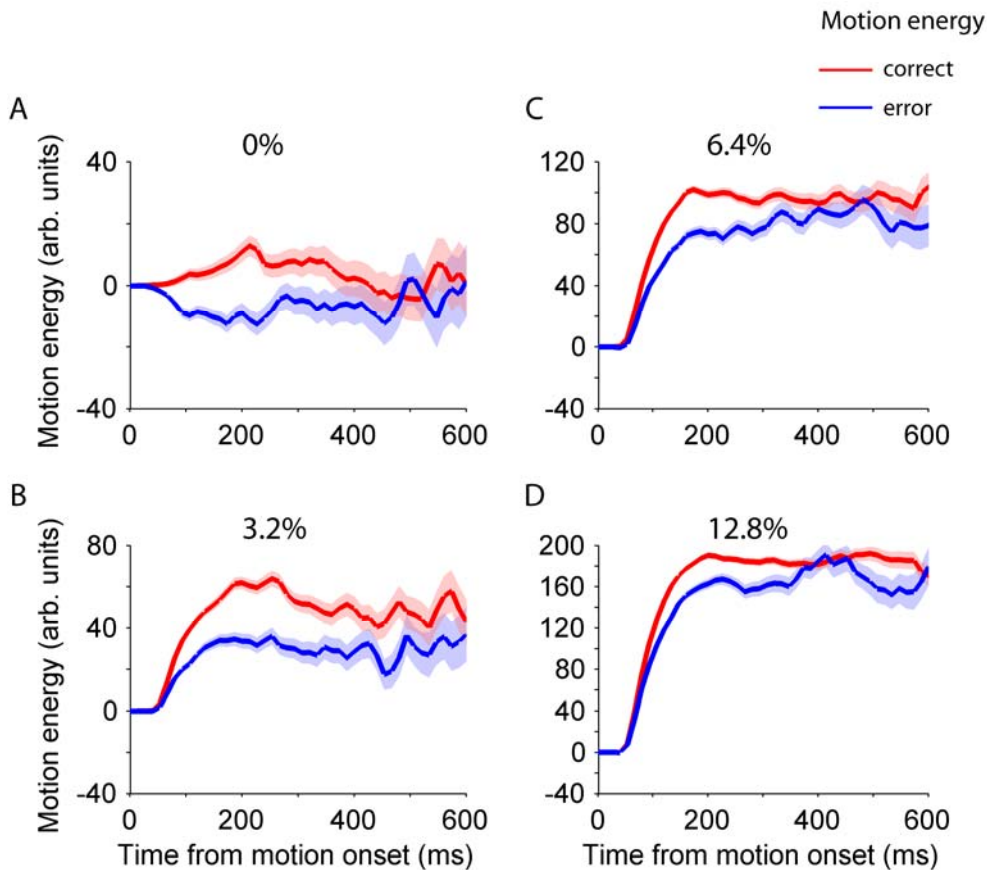


Figure S3.

Decision time is limited by a bounded evidence accumulation mechanism. The bounded accumulation model asserts a mechanism for terminating decisions. In reaction time tasks, this explains the relationship between decision speed and accuracy. In an experiment like ours, in which the stream of evidence is controlled by the experimenter, the mechanism asserts that some decisions terminate covertly, despite the availability of additional information from the stimulus. After the accumulated evidence reaches the bound, the incoming evidence would be largely ignored. Therefore, the initial part of the stimulus must have a larger influence on the monkey's choice. We performed a psychophysical reverse correlation analysis to test this assertion (*S11*, *S12*). On each trial, a particular random dot stimulus gave rise to a noisy stream of motion information, comprised of fluctuations in the magnitude and sign of evidence. These fluctuations were quantified by calculating the motion energy (*20*, *S13*). Positive motion energy indicates rightward motion for 0% coherence and motion in the 'correct' direction for non-zero coherences. The motion energy profiles for rightward and leftward choices (0% coherence, A) and correct and error choices (non-zero coherences, B-D) show a clear separation in support of the choice. The separation was clearest in the first ~400 ms, and overlapped at later times, consistent with a covert decision termination mechanism. Only trials with stimulus durations of 400 ms or larger contributed to this analysis.

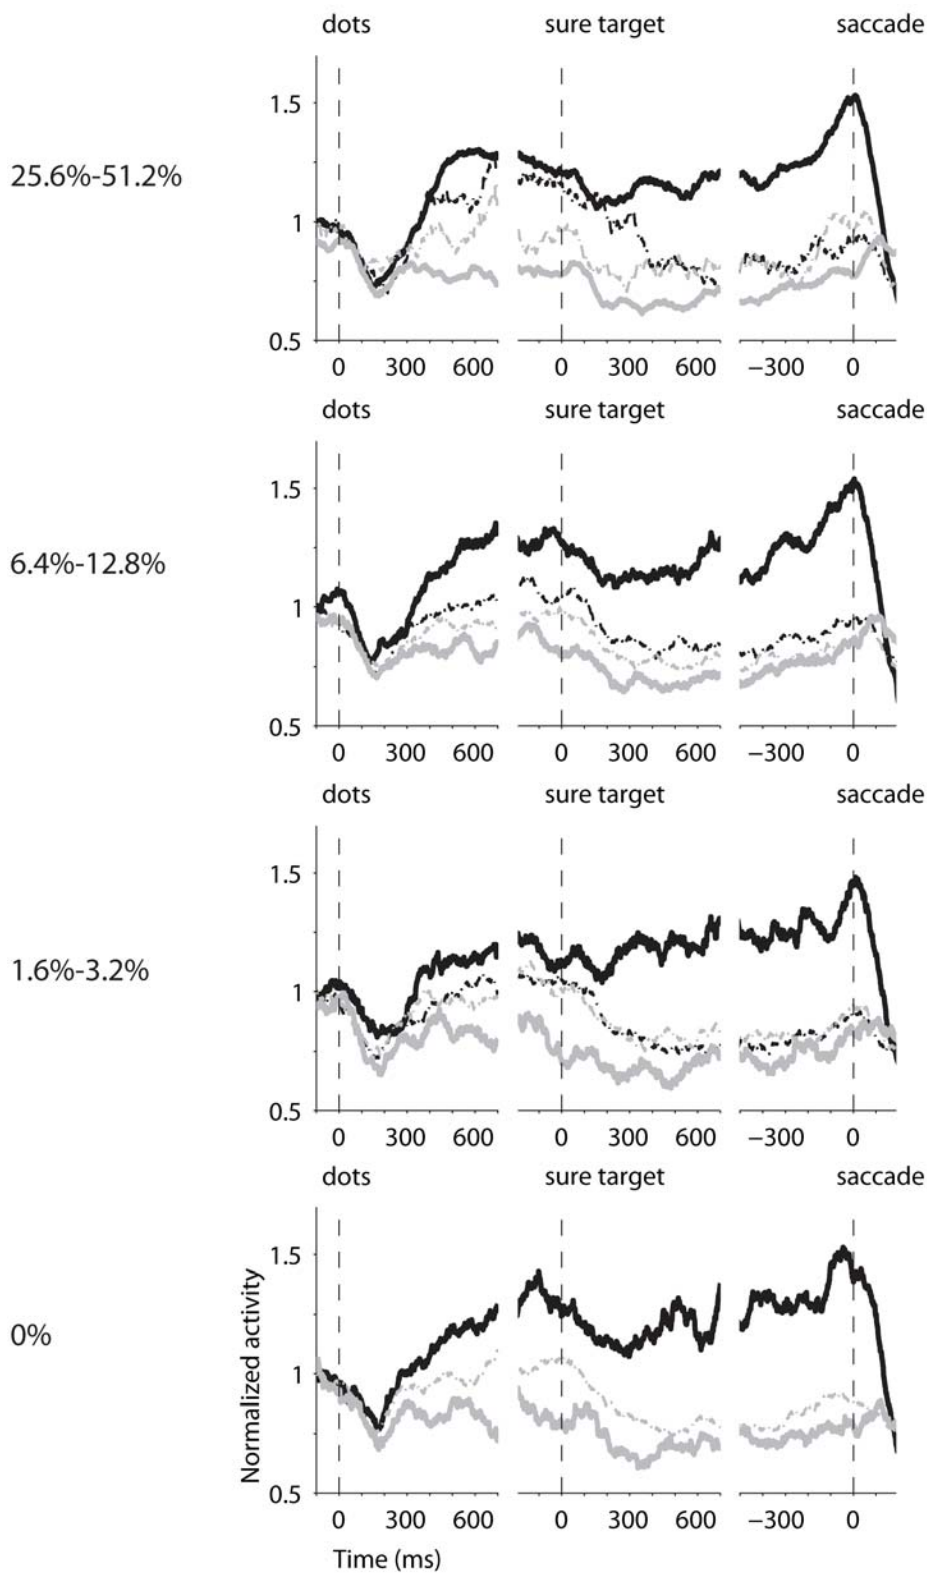


Figure S4.

The level of activity for T_s choices was modulated by motion strength. Population average responses from 70 LIP neurons are grouped by motion strength. One of the high stakes (direction choice) targets was in the RF of the recorded neurons. Plotting conventions are identical to Figure 2D. Coherence levels are paired for illustration purposes. The same trend exists for individual coherence levels.

Table S1. Fit parameters (mean±SE) of the bounded accumulation model.

k	0.255±0.002
B	39.4±10.0
θ	0.591±0.005

References

- S1. W. T. Newsome, E. B. Pare, *J. Neurosci.* **8**, 2201 (1988).
- S2. K. H. Britten, M. N. Shadlen, W. T. Newsome, J. A. Movshon, *J. Neurosci.* **12**, 4745 (1992).
- S3. J. I. Gold, M. N. Shadlen, *J. Neurosci.* **23**, 632 (2003).
- S4. J. W. Lewis, D. C. Van Essen, *J. Comp. Neurol.* **428**, 112 (2000).
- S5. D. C. Van Essen, *Curr. Opin. Neurobiol.* **12**, 574 (2002).
- S6. R. M. Bracewell, P. Mazzoni, S. Barash, R. A. Andersen, *J. Neurophysiol.* **76**, 1457 (1996).
- S7. P. L. Smith, R. Ratcliff, *Trends Neurosci.* **27**, 161 (2004).
- S8. S. Karlin, H. M. Taylor, *A second course in stochastic processes* (Academic Press, 1981).
- S9. K. H. Britten, M. N. Shadlen, W. T. Newsome, J. A. Movshon, *Vis. Neurosci.* **10**, 1157 (1993).
- S10. J. Ditterich, M. E. Mazurek, M. N. Shadlen, *Nat. Neurosci.* **6**, 891 (2003).
- S11. P. Neri, *J. Vis.* **4**, 82 (2004).
- S12. A. J. Ahumada, Jr., *J. Vis.* **2**, 121 (2002).
- S13. E. H. Adelson, J. R. Bergen, *J. Opt. Soc. Am. A.* **2**, 284 (1985).