Neural mechanisms for decisions based on sequential samples of evidence

Shinichiro Kira

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2014

Reading Committee:

Michael N. Shadlen, Chair

Eberhard Fetz

Eric T. Shea-Brown

Program Authorized to Offer Degree:

Neurobiology and Behavior

University of Washington

Abstract

Neural mechanisms for decisions
based on sequential samples of evidence

Shinichiro Kira

Chair of the Supervisory Committee:
Professor Michael N. Shadlen
Graduate Program in Neurobiology and Behavior

Difficult decisions often require evaluation of samples of evidence that are acquired sequentially. If samples of evidence differ in their reliability and an unlimited number of samples can be acquired sequentially, the optimal decision strategy is to accumulate evidence in units of log likelihood ratio (logLR) until the accumulated evidence exceeds a predefined threshold, or bound. For this dissertation, I hypothesized that the brain makes probabilistic inferences and decisions by such bonded accumulation of evidence.

I trained monkeys to make decisions based on a sequential presentation of visual shapes. Shapes conferred probabilistic evidence that differed in predicting which target would furnish a reward. The stream of shape presentation continued until the monkeys indicated their decisions by eye movements. I found that the monkeys terminated their decisions when the cumulative logLR reached a stereotyped bound, on average. Firing rates of neurons in the lateral intraparietal area (LIP) reflected the accumulated evidence (i.e., the cumulative logLR) and also reached a stereotyped level before the monkeys committed to a decision. LIP activity explained the monkeys' choices and reaction times that reasonably approximated those expected by the optimal strategy.

Psychophysical bounds for decision termination often decline in a time-dependent manner, which has been hypothesized to reflect a decision-making strategy that incorporates the cost of elapsed time. In the second part of my dissertation, I tested this hypothesis by manipulating time-cost in a perceptual choice reaction-time task. Specifically, I introduced provisional deadlines while human subjects made decisions about the direction of a random-dot motion stimulus. The pattern of choices and reaction times were explained by bounded accumulation of evidence. Termination bounds were found to decline as a function of time, which was exaggerated in response to the time-cost manipulation.

In most of the previous studies of decision making in LIP, the monkeys were aware of the specific eye movement that would result from their decisions while they deliberated on the decision. Thus, it remained unknown whether evidence accumulation in LIP neurons was an abstract representation of the decision itself or a reflection of the impending eye movement. I modified the sequential sampling task, by preventing the monkey from planning a specific eye movement while it accumulated evidence. LIP neurons did not show evidence-dependent activity during the early phase of evidence accumulation period although the evidence provided in this phase affected the monkey's ultimate choice. The result suggests that LIP activity does not represent accumulated evidence per se, but only in preparation for an impending saccade.

These experiments demonstrate that a bounded accumulation of evidence explains various types of decision-making such as probabilistic inference and perceptual discrimination. For monkeys, the bounded accumulation of evidence is represented in LIP activity when decisions are formed in parallel with a preparation for a specific eye movement.

Table of Contents

List of Figures

List of Tables

Acknowledgements

First and foremost, I would like to thank my graduate training advisor, Mike Shadlen. I joined the graduate program at University of Washington because I wanted to work with him. As I studied neuroscience as a medical student in Japan, I was especially interested in Mike's quantitative approach to study neurobiology of decision-making with physiological recording and computational modeling. When I was awarded a scholarship from my medical school to do research outside of Japan for a month, I immediately contacted Mike and asked if I could work in his lab for that period. He kindly accepted my request and offered me a chance to analyze data even before I joined the graduate program at University of Washington. This was a very fruitful experience and eventually became one of my thesis projects.

Mike has been an excellent mentor and has always taught me tips for how to become a better scientist. For example, Mike would not immediately correct me during our discussions, but gave me time to think what might be wrong. Since English is not my first language, Mike has given me a lot of advice on how to clearly convey ideas in scientific writing. Perhaps, the most important lesson from Mike is that a good scientist has to be always honest and self-critical. I think Mike himself is the most critical reviewer of his own works, and I would also like to maintain such an attitude throughout the rest of my career. Other than science, he has been really thoughtful for my life events and always supportive during my mishaps and setbacks. Owing to his support, I was able to keep focus on my projects.

I also thank other members of my Supervisory Committee: Eb Fetz, Eric Shea-Brown, and Anitha Pasupathy, and my Graduate School Representative: John Palmer. Their advice at committee meetings has helped me refine analyses and revise manuscripts. I enjoyed my rotation in Eric's lab and collaboration with his student, Nick Cain. This rotation gave me confidence to tackle mathematical problems, which seemed intimidating during the early stage of my graduate training. I thank John for always reading my drafts carefully and critically, and helping me set achievable goals so that I can advance towards my goals step by step.

Interactions with the lab members, office mates, and technicians have been a great source of education, motivation, and entertainment. They taught me new techniques and analyses, and pointed me to helpful resources. They read my manuscripts carefully and gave me critical comments (special thanks to Charlie Hass, Luke Woloszyn, and Shushruth). They listened to my speculative ideas with patience and encouraged me to support those ideas with additional analyses rather than rejecting them. They gave me guidance on how to plan a career path (special thanks to Yasmine El-Shamayleh). They shared joy and grief in my work and life. Without their tremendous support, I could not accomplish my goals during the graduate training.

Outside the lab, I thank administrators of Neurobiology & Behavior Program and Department of Physiology & Biophysics—especially Lucia Wisdom, Ann Wilkinson, and Tina Schulstad. Because I was an international student and also transferred to Columbia University, their support was critical to maintain my student status at University of Washington. I also greatly appreciate scholarship and support from Nakajima Foundation throughout my graduate training.

Finally I thank my family members. I am grateful to my wife, Yumi, for her tremendous support. We had a lot of fun together, which I could not experience by myself. We also teamed up to overcome many problems, which I could not accomplish without her help. In addition, she is a great chef and cooks me delicious meals, which always refresh my mind and refuel me to stay focused. I also thank my parents for their understanding. None of my family members have a science or research-related background, but they understand my interest and have always supported me to pursue my dream.

Chapter 1: Introduction

Decision-making allows animals to choose an appropriate option among alternatives, and thus guides them to maximize their reward (Gold and Shadlen, 2002; Kable and Glimcher, 2009; Shadlen and Kiani, 2013). What constitutes the "reward" may vary depending on an individual or situation (Behrens et al., 2009). For example, animals in the natural environment may try to choose an action that maximizes the amount of food or water they obtain during a day whereas some people may attempt to acquire positive appraisals from their peers. In addition, it is also rewarding to minimize the cost or risk associated with the chosen option. For animals, choosing an appropriate path to the nest may save their energy and also reduce the predation risk, thereby increasing their chance of survival. Thus, decision-making is crucial for animals to survive in the natural environment and for humans to live in society.

As an animal's repertoire of possible behaviors expands, decision-making becomes more complex because stereotyped or random behaviors are unlikely to lead to a favorable reward. This can be realized by comparing animals in different levels of evolutionary hierarchy. In lower animals, most of their behavior can be explained by a reflex—that is, responding immediately to the external stimulus in a stereotyped manner. In contrast, higher animals such as humans have acquired a rich repertoire of behaviors, giving rise to their more complex stimulus-response patterns (Miller and Cohen, 2001). For example, we, as humans, respond to the same stimulus with a different behavior depending on our internal state, such as attention to the stimulus. We also change our behavior by learning from our past experience. Moreover, we do not have to

respond to the stimulus immediately as in a reflex. Rather, we can store information in memory and decide later, or store a decision in memory and execute it as an action later. This enables us to accumulate additional pieces of information before reaching a decision, or use the outcome of one decision as a guide for another decision (Shadlen and Kiani, 2013).

My thesis builds upon neural principles of decision-making formulated through past studies in physiology and psychology. The brain can accumulate evidence to improve the decision accuracy; when unlimited amount of evidence is available, the brain accumulates evidence until an amount of accumulated evidence reaches a decision criterion. In this chapter, I will first introduce some of the experiments in perceptual decision-making that revealed these principles. Then I will provide rationale for these principles from a statistical point of view. These discussions will lead to three questions in my thesis: (1) do these principles apply beyond perceptual decisions to account for decisions based on arbitrary cues? (2) do subjects adjust these principles in a normative way when an additional time cost is introduced by a deadline? (3) do neural activities measured during the deliberation period really reflect accumulation of evidence rather than a motor plan for an impending action?

Many important principles of decision-making have emerged from studying perceptual decisions. Since mid-19th century, psychophysics has been a discipline that investigates the quantitative relationship between perceptual decisions and physical stimuli (Fechner, 1860; Green and Swets, 1966). In mid-20th century, owing to advances in neural recording, physiologists started to explain perceptual decision by neural activities (Parker and Newsome,

1998). Their progress largely hinged on the development of the signal detection theory (Green and Swets, 1966). In this theory, subjects discriminate a signal based on whether their internal representation of the signal exceeds their internal threshold. The theory assumes that the internal representation is noisy. When there is a large difference in the two signals to be discriminated, the effect of noise is small and subjects correctly discriminate the signals (i.e. hits and correct rejections). As the difference in two signals becomes smaller, the effect of noise increases and subjects sometimes mislabel the signal (i.e. miss and false alarms). Physiologists reasoned that the firing rate of sensory neurons is the internal representation of stimuli because their firing rate reflects the signal property with noise (i.e. variability in the firing rate) (Barlow et al., 1971). This simple extension of signal detection theory allowed them to explain perceptual decisions by neural activity recorded from sensory neurons.

However, their conclusions were often derived from a comparison of behavioral and neural data collected under different conditions because it was technically difficulty to record neural activity from awake behaving animals or human subjects engaging in a task. Instead, in many studies, neural activity was recorded from different animals or even from different species (LaMotte and Mountcastle, 1975; Tolhurst et al., 1983; Parker and Hawken, 1985). Consequently, their conclusions were inevitably undermined by inter-subject or inter-species differences—neural activity that explained decisions may not have been present when the subjects actually made those decisions.

To eliminate this concern, Newsome and colleagues compared the decision-making behavior and neural activity that are simultaneously recorded from the same animals (Newsome et al., 1989; Britten et al., 1992). They trained monkeys to make a decision on the net motion direction (e.g. rightward vs leftward) of a random dot kinematogram, and to report the decision by a saccadic eye movement to one of two choice targets. When the stimulus had a strong rightward or leftward motion, the monkey accurately reported the motion direction. When the motion strength was weak, however, the monkey made error choices by reporting the opposite direction from the stimulus direction. Thus, decision accuracy increases as a function of the motion strength, which is summarized as a psychometric function. While the monkey performed the experiment, they recorded from motion-sensitive areas in the brain, namely the middle temporal area (MT). During the motion presentation, the firing rate of an MT neuron fluctuates from moment to moment around an underlying rate, which is higher for a particular motion direction (preferred direction) than for the opposite direction (null direction). As the motion strength is varied from the null direction (negative motion strength) to the preferred direction (positive motion strength), the firing rate of MT neurons increases roughly proportionally to the motion strength (Britten et al., 1993). To test whether the motion signal in MT activity is sufficient to explain the monkey's performance, Newsome and colleagues predicted the decision accuracy by applying signal detection theory—that is, by comparing signals from a pair of MT neurons: a recorded neuron that prefers one direction and a hypothetical 'anti-neuron' that prefers the opposite direction. Assuming that the decision is made based on which neuron in the pair generated a greater number of spikes, they predicted the decision accuracy for each motion strength, which is summarized as a neurometric function. As the motion strength increases, a pair

of neurons exhibits more distinct patterns of activity, which would allow more accurate decisions to be made. As a result, the predicted decision accuracy monotonically increases for larger motion strengths in a neurometric function. The psychometric and neurometric functions allowed them to compare the behavior and neural activity on the same probabilistic measure.

Surprisingly, the monkeys made less accurate decisions than those based on the most sensitive MT neurons that exhibited the most distinctive firing rates to the preferred vs null motion direction. This means that if the monkey had made decisions based only on these cells, the monkey could have achieved a higher accuracy. To account for this discrepancy in the decision accuracy, they reasoned that the brain makes decisions by pooling signals not only from these sensitive neurons, but also from less sensitive ones (Britten et al., 1992). However, this idea still did not explain the gap easily. As the brain pools signals from N neurons, the signal-to-noise ratio (SNR) of their summed signal improves by a factor of $\sqrt{N}$ if the noise in each neuron is independent from the rest of the neurons. Therefore, a small pool of neurons (10-100) can outperform the monkey, creating a discrepancy again. To reconcile this gap, one needs to realize that the noise correlation between neurons limits the improvement in SNR. In case of MT neurons, their noise is weakly correlated (r=~0.12) (Zohary et al., 1994), so that the improvement of the SNR plateaus once the group size reaches ~100 neurons. By taking this noise correlation into account, the pooled signals from MT neurons successfully explained the monkey's decision accuracy (Shadlen et al., 1996).

Signals can also be pooled across time, in the form of temporal integration, to improve decision accuracy especially when noisy or ambiguous evidence arrives sequentially in time. Suppose, as described above, that the brain decides the direction of random dot motion based on which of two competing pools of MT neurons generated a greater number of spikes during some time window. Because the MT firing rate fluctuates with noise from moment to moment, the decision becomes more accurate as the time window becomes longer. Consistent with this idea, when humans and monkeys performed the random-dot task with variable length of motion stimulus presentation, the decision accuracy improved as the stimulus length became longer up to a few hundreds ms (Watamaniuk and Sekuler, 1992; Burr and Santoro, 2001; Kiani et al., 2008), which is longer than the time scale of temporal motion integration observed in MT neurons (<150 ms; Bair and Movshon, 2004). This suggests that the brain has a mechanism to accumulate evidence over time.

Because the monkey's choice is indicated by eye movement in the random dot task, one candidate area to study the neural mechanism of temporal integration was the lateral intraparietal area (LIP) in the posterior parietal cortex. Anatomically, LIP receives direct inputs from visual sensory areas, including MT, and sends output to oculomotor areas such as the frontal eye field (FEF) and the superior colliculus (SC) (Lynch et al., 1985; Felleman and Van Essen, 1991; Schall et al., 1995). Physiologically, it was known that LIP neurons show a sustained activity when the monkey is planning to make a saccadic eye movement towards a particular part of the visual field, referred to as a response field (RF) (Gnadt and Andersen, 1988; Barash et al., 1991). Both

anatomical and physiological evidence suggested that LIP neurons might integrate a sensory motion signal into an oculomotor decision signal during the random dot task.

Physiological recording from LIP revealed that the evolution of LIP activity reflects temporal integration during decision-making. Shadlen and Newsome (2001) recorded from LIP neurons while monkeys performed the random dot task. The monkey observed the random dot stimulus for 2 sec and decided the net motion direction. The stimulus offset was followed by a delay period, and then the monkey was required to choose one of two choice targets that correspond to the net direction. One target was placed in the RF of the recorded LIP neuron whereas the other target was placed outside of the RF in the opposite hemifield. After the stimulus onset, LIP neurons increased or decreased their firing rate depending on whether the monkey's impending saccade direction was to or away from the RF target, respectively. Moreover, the rate of increase or decrease in the firing rate varied with the motion strength in the stimulus. This is consistent with an idea that LIP integrates the difference of the firing rates in two opposing MT neurons. Therefore, these results suggested that LIP firing rate reflects accumulation of sensory evidence in favor of a choice target in the RF.

This experiment, however, could not distinguish neural activity associated with decision formation from the rest of activities (e.g. commitment to a choice) because they could not measure how long the monkey spent to make a decision on each trial. The monkey could have reached a decision at any time during the 2 sec of stimulus presentation, or might have not reached a decision by the end of presentation but had to make a choice. By contrast, the time

spent on the decision formation can be unambiguously measured by a reaction time task, in which the monkey is allowed to make a choice whenever it reaches a decision. In fact, many studies reported that the mean reaction time for the random dot task was shorter than 2 sec (Roitman and Shadlen, 2002; Palmer et al., 2005; Cohen and Newsome, 2009). Interestingly, firing rates of MT neurons better explained the monkeys' psychometric performance when their firing rates were computed within the reaction time rather than the entire 2 sec of the stimulus duration (Cohen and Newsome, 2009).

The reaction time task allows experimenters not only to isolate the neural activity associated with decision formation, but also to study neural mechanisms that terminate the decision process. Roitman and Shadlen (2002) trained monkeys on the reaction time version of the random dot task, in which monkeys were allowed to observe the motion stimulus until they reached a decision and chose a target by a saccadic eye movement. While the monkeys engaged in the task, their LIP response was measured. After the stimulus onset, LIP neurons started to change their activity with a build-up rate proportional to the motion strength, consistent with the aforementioned experiment. Importantly, LIP firing rate ramped up to a stereotyped level ~50 ms before saccade to the RF target, which signaled the termination of a decision-making process. Recently, very similar results were observed in physiological recording from FEF when the monkey reported their decision by saccades in the random dot task (Ding and Gold, 2012). Altogether, these physiological results suggest that the neural activity associated with evidence accumulation terminates the accumulation process when the activity reaches a threshold (or a bound) — a mechanism often referred to as the "bounded accumulation".

Furthermore, Shadlen and colleagues quantitatively showed that the bounded accumulation model can explain the speed and accuracy of decisions in the random dot task (Ditterich et al., 2003; Huk and Shadlen, 2005; Palmer et al., 2005; Hanks et al., 2006; Churchland et al., 2008; Kiani and Shadlen, 2009). They attempted to reproduce the performance of the subjects with a computational model, which assumed that the LIP activity reflects the accumulation of noisy sensory signal from MT, namely a difference in the firing rate of two opposing MT neurons. In their model, the 'decision variable' accumulates noisy momentary evidence, which is a normally distributed random variable with the mean proportional to the motion strength. For example, in a right-left discrimination task, the mean may be positive for rightward motion, and negative for leftward motion. As these momentary samples of evidence are summed into the decision variable, the decision variable evolves as if a particle diffuses with a constant drift, giving the model its name: "diffusion model". The decision variable continues to accumulate momentary evidence until it reaches either an upper or a lower bound, corresponding to each choice (see Palmer et al., 2005 and Shadlen et al., 2007 for details of the model). This model successfully explained the distribution of both choice and reaction time, providing additional support for the idea that the bounded accumulation mechanism governs decision-making in the random dot task.

The appeal of the bounded accumulation lies not only in its explanatory power but also in its optimal performance in statistical decision-making when evidence arrives piece by piece (Wald, 1947). Therefore, it has been widely applied in industry and well studied in statistics as

the sequential probability ratio test (SPRT). SPRT achieves a desired accuracy with the fewest pieces of evidence on average (Wald and Wolfowitz, 1948). Because some of my hypotheses in this thesis stem from properties of SPRT, below I will briefly explain how SPRT operates.

In SPRT, the decision variable represents the belief in two hypotheses, and keeps evolving as pieces of evidence arrive until the decision variable reaches one of two decision bounds. Suppose a decision maker has to decide between two hypotheses (H_0 and H_1) based on multiple pieces of evidence. Given a set of evidence, $\mathbf{e} = \{e_1,\ldots,e_n\}$, one can compare the belief in each hypothesis, for instance, by taking their probability odds: $P(H_1 \mid e) / P(H_0 \mid e)$. If the odds are greater than 1 (or smaller than 1), H_1 (or H_0) is the more likely hypothesis. By applying Bayes' rule, the probability odds could be written as:

$$\frac{P(H_1 \mid \mathbf{e})}{P(H_0 \mid \mathbf{e})} = \frac{P(\mathbf{e} \mid H_1)}{P(\mathbf{e} \mid H_0)} \cdot \frac{P(H_1)}{P(H_0)} \qquad (1)$$

If the two hypotheses are equally likely before the arrival of evidence (i.e. flat priors), the second term on the right-hand side vanishes from Eq (1). Furthermore, if each piece of evidence is conditionally independent under each hypothesis—that is if each piece of evidence is drawn independently from a designated sampling distribution under each hypothesis—the first term on the right-hand side of Eq (1) can be broken down into products of likelihood ratios as follows:

$$\frac{P(H_1 \mid \mathbf{e})}{P(H_0 \mid \mathbf{e})} = \frac{P(e_1 \mid H_1)}{P(e_1 \mid H_0)} \cdots \frac{P(e_n \mid H_1)}{P(e_n \mid H_0)} \qquad (2)$$

By taking the log of Eq (2), the log posterior odds, L_n, can be computed by a cumulative sum of log likelihood ratio (logLR) for each piece of evidence:

$$\begin{aligned} L_n &\equiv \log \frac{P(H_1 \mid \mathbf{e})}{P(H_0 \mid \mathbf{e})} = \log \frac{P(e_1 \mid H_1)}{P(e_1 \mid H_0)} + \cdots + \log \frac{P(e_n \mid H_1)}{P(e_n \mid H_0)} \\ &= \sum_{i=1}^{n} \log \frac{P(e_i \mid H_1)}{P(e_i \mid H_0)} \end{aligned} \tag{3}$$

This operation is employed in SPRT—as a new piece of evidence (e_n) arrives, SPRT updates the log posterior odds from L_{n-1} to L_n by summing logLR of e_n:

$$\begin{aligned} L_n &= \sum_{i=1}^{n} \log \frac{P(e_i \mid H_1)}{P(e_i \mid H_0)} \\ &= \sum_{i=1}^{n-1} \log \frac{P(e_i \mid H_1)}{P(e_i \mid H_0)} + \log \frac{p(e_n \mid H_1)}{p(e_n \mid H_0)} \\ &= L_{n-1} + \log \frac{p(e_n \mid H_1)}{p(e_n \mid H_0)} \end{aligned} \tag{4}$$

SPRT makes a decision when L_i exceeds one of decision bounds (+B or -B), and otherwise continues to acquire a new piece of evidence:

$$\begin{aligned} &L_i \le -B : \text{choose } H_0 \\ &L_i \ge B \ : \text{choose } H_1 \end{aligned} \tag{5}$$

$-B \le L_i \le B$: acquire a new sample of evidence

If the bounds are set at the same level ($B>0$) for both H_A and H_B, a desired level of accuracy (a proportion of correct trials α) can be obtained by setting bounds at

$$B = \log \frac{\alpha}{1-\alpha} \qquad (6)$$

Here, the log posterior odds, L_i, is the decision variable, to which a decision termination criterion (i.e. bound) is applied. As a new piece of evidence arrives, the decision variable evolves through a random walk as in Eq (4) until it reaches either the upper bound (+B) or lower bound (-B).

The key idea is that, in the random dot task, the bounded accumulation reflected in LIP activity conforms to SPRT because the momentary evidence from MT represents logLR (Gold and Shadlen, 2001). Suppose that the motion stimulus induces normally distributed firing rates with different means in two opposing MT neurons: one preferring rightward and the other preferring leftward. If the firing rate of the leftward-preferring neuron is subtracted from that of the rightward-preferring neuron, the difference in the rates is proportional to logLR in favor of right choice over the left choice. Thus, by integrating this difference in MT firing rate, LIP firing rate can represent the log odds of two alternative motion directions. Because SPRT achieves a desired accuracy with the least number of samples of evidence on average, monkeys can maximize the number of correct trials in a given time and thus achieve the maximum "reward rate" if they set decision bounds at appropriate level (Gold and Shadlen, 2002).

The firing rate of a sensory neuron is suitable for accumulation in many other perceptual decision-making tasks. Perceptual decisions often require subjects to make decisions between two hypotheses (e.g. signal vs no signal). Suppose the stimulus induces a sensory neuron to exhibit a normally distributed firing rate with a different mean under each hypothesis but with the same variance. For these conditions, the firing rate from even a single sensory neuron represents logLR:

$$\log \frac{P(r \mid H_1)}{P(r \mid H_0)} \propto r - c \qquad (7)$$

where r is the firing rate of a neuron, c is a constant value, and H_0 and H_1 are the two hypotheses being tested. This proportionality also holds when firing rates are distributed according to a Poisson or exponential distribution (Gold and Shadlen, 2001). For such cases, the firing rate can be used as appropriate evidence for accumulation regardless of sensory modality. Therefore, even when the subject is required to make decisions based on neural activities represented in several sensory modalities, the brain can compute the log posterior odds by a weighted sum of the firing rates of sensory neurons in these modalities (Ma et al., 2006; Fetsch et al., 2013).

In contrast, the firing rate of a sensory neuron is not suitable for accumulation in more complex decision-making tasks, where the decision-maker has to learn probabilistic association between multiple arbitrary cues and hypotheses. Consider a case where a doctor makes a diagnosis whether a patient has a specific disease or not based on the patient's symptoms. Here,

it is unlikely that the firing rate of a single sensory neuron is systematically distributed (i.e. according to a normal, exponential, or Poisson distribution) under each hypothesis, which implies that its firing rate cannot represent logLR. Theoretically, its downstream neurons could represent logLR if they receive inputs from sensory neurons with appropriate weights, but it was unknown whether the brain can represent and integrate probabilistic information.

Recently, physiological recording from LIP suggested that its neural activity reflects probability integration even in such a complex task. Yang and Shadlen (2007) trained monkeys to make binary decisions based on the sequential presentation of four shapes, each of which was randomly sampled from a set of eight unique shapes. Each shape altered the probability of obtaining a reward from one choice target or the other. Behavioral analyses showed that monkeys learned the probabilistic association between the shapes and choice target, and made decisions based on accumulated evidence. In physiology, LIP firing rate reflected the cumulative logLR of the presented shapes, indicating that LIP activity reflected the probability integration in units of logLR. These findings suggest that the brain can combine probabilistic information in a form that can be used for SPRT.

Chapter 2 of my thesis assesses whether the evolution of LIP activity during probabilistic inference conforms to SPRT. Methods used in Yang and Shadlen (2007) had two limitations that precluded this assessment. First, they could not measure when the monkeys reached a decision because the experimenters, rather than the monkeys, controlled the amount of evidence available to the monkeys in each trial. This concern arises for the same reason that Shadlen and Newsome

(2001) could not address termination mechanism when they presented the moving dots for a fixed duration in every trial. Second, this study could not accurately quantify how much evidence was accumulated by the monkey at each moment. Each shape remained on the screen until all four shapes were presented. Thus, instead of accumulating evidence from a new shape as it appears on the screen, the monkey could have waited until all shapes appear on the screen and accumulated evidence all at once as a "snap shot". By ameliorating these limitations, I show that LIP activity represents accumulation of noisy evidence, and signals the end of decision process when the activity reaches a threshold level. The results demonstrate similarities between the evolution of LIP activity and SPRT.

Chapter 3 shows that increased time cost (i.e. time pressure) dynamically lowers the subject's decision criterion as time elapses during a trial. Accumulation of evidence to a fixed bound, as in SPRT, achieves the maximum reward rate when the accumulation of evidence does not incur any cost. However, time-dependent declining bounds achieve a better reward rate than fixed bounds when the accumulation of a new piece of evidence incurs a cost such as mental fatigue (Drugowitsch et al., 2012) or increased chance of unfavorable outcome. To study how subjects incorporate time pressure in their decision-making process, I compared the performance of subjects with and without deadlines. As the deadlines increased time pressure, subjects made decisions faster, and their decision accuracy was lower for a given range of reaction time. Subjects' decision speed and accuracy were well explained by a diffusion model with decision bounds that become lower as time elapses during a trial, referred to as the "collapsing bounds".

The model indicated that the time pressure induced subjects to lower their decision bounds more sharply, which helped subjects improve their performance.

Chapter 4 addresses whether the "decision-related" activity in LIP represents the decision formation itself or an evolving motor plan for an impending saccade. In many decision-making tasks, choice targets are displayed from the beginning of each trial. As a result, the decision formation is coupled with the development of an oculomotor plan. If the "decision-related" activities in LIP reflect an oculomotor plan, this activity may vanish when the monkeys cannot plan where to make an eye movement during evidence accumulation. To test this possibility, I recorded from LIP neurons during a sequential inference task similar to the ones used in Chapter 2 and Yang and Shadlen (2007). I prevented the monkey from planning a specific eye movement during the evidence accumulation period by hiding choice targets until the end of the evidence-accumulation period and by randomizing the target locations from trial to trial. Whereas the evidence provided in the early phase of the evidence-accumulation period affected the monkey's choice, the LIP response did not reflect this evidence until the late phase that preceded the time of saccade. The dissociation in the leverage of evidence on choice and on LIP response indicates that LIP response does not represent the evidence accumulation itself. Rather it represents an oculomotor plan that is modulated by the accumulated evidence.

Chapter 2: Neural mechanism for decisions based on sequential samples of evidence that differ in reliability

Abstract

I recorded from neurons in the intraparietal area (LIP) while monkeys made binary decisions that required probabilistic inference based on a sequential presentation of shapes. Shapes conferred probabilistic evidence that differed in reliability in predicting which target would furnish a reward. The monkeys were allowed to observe additional samples of shapes until they indicate their choice by a saccade. The firing rate of LIP neurons reflected the accumulation of evidence in probability units, and signaled the end of decision when the firing rate reached a threshold level. A physiologically-constrained model demonstrated that the monkey's decision speed and accuracy could be explained by a race between two neural populations, each of which accumulates evidence up to a threshold for the two alternative choices. This neural mechanism conforms to a normative decision-making strategy employed in sequential hypothesis testing.

Introduction

Complex decision-making often requires the collection of multiple pieces of evidence before committing to a decision. Along the way, the brain must evaluate each piece of evidence, combine them together, and determine whether more evidence is needed or the current evidence is sufficient for a final decision. Recent advances in neuroscience have begun to reveal the neural correlates of this process (Gold and Shadlen, 2007). When an animal is trained to indicate a decision with an eye movement, neurons in the parietal and prefrontal cortex, which are associated with response selection, represent the accumulating evidence during deliberation. The same neurons achieve a stereotyped level of firing rate upon completion of the decision (Roitman and Shadlen, 2002; Churchland et al., 2008). Thus these neurons are thought to participate in the conversion of evidence to a decision variable (DV) suitable for accumulation and comparison to a threshold (or bound) for terminating the decision process with a choice. Although the neural mechanisms are less well understood, a similar "bounded evidence accumulation" framework explains a variety of perceptual and mnemonic decisions in animals and humans (Ratcliff and McKoon, 2008; O'Connell et al., 2012; Shadlen and Kiani, 2013).

The idea is appealing because the accumulation of evidence might be likened to the evolution of belief in a proposition. However, it presupposes that the brain possesses a mechanism to convert sensory evidence into probabilistic value associated with degree of belief (Gold and Shadlen, 2001; Pouget et al., 2013). It has been shown that humans and nonhuman primates rationally combine simultaneous cues in accordance with their reliability (Fetsch et al.,

2013), such as judging slant form texture and disparity (Knill, 2007); or heading from motion and vestibular sensation (Fetsch et al., 2012). However such rational combination of cues has not been studied extensively in the setting of decision making from a sequence of cues which are separated in time. This is because most studies of decision making employ a single stimulus whose reliability is fixed (i.e., statistically stationary) over the course of a decision.

To overcome this limitation, Yang & Shadlen (2007) trained monkeys to observe a sequence of shape-cues that furnished probabilistic evidence bearing on a binary decision. They showed that the monkeys based their decisions on the combined evidence from four cues, giving more weight to the more reliable cues. Moreover, as the shapes appeared sequentially during a trial, the firing rates of neurons in area LIP tracked the running sum of the evidence, in units proportional to logLR, for and against the choice alternatives. This suggests that the brain can optimally combine cues from sequential samples. However, two aspects of this study preclude a direct connection to the "bounded evidence accumulation" mechanism mentioned above. First, there was no measure of decision termination (e.g., reaction time), because four shape cues were shown on each trial. Second, although the animal based its choices on the total evidence from the four shapes, there was no actual requirement to integrate evidence in time. This is because each shape remained visible from the time it was presented until the monkey made a decision. Thus, it was possible that the monkeys based their decision on the combination of the four cues present at the end of each trial.

Here I employed a modified version of this probabilistic classification task in which a sequence of shape-cues are presented transiently, until the monkey terminates the sequence with a decision. The task makes explicit demands on working memory and evidence accumulation. Moreover, to perform this task optimally, the monkey should terminate decisions when the accumulated logLR reaches a threshold level, or bound (Wald and Wolfowitz, 1948). This process, termed the Sequential Probability Ratio Test (SPRT; (Wald, 1947)), is optimal in the sense that it requires the least number of samples, on average, to achieve any given level of accuracy. I hypothesized that the primate brain possesses the capacity to support just such a stratagem.

Methods

Two male rhesus monkeys (*Macaca mulatta*) were implanted with a head post and a recording chamber above the intraparietal sulcus. Throughout training and recording sessions, horizontal and vertical eye positions were recorded by a EyeLink infrared video-tracking system (sampling rate, 1000 Hz). Timing of task-related events were monitored and controlled by NIH Rex system. Visual stimuli were displayed using Psychtoolbox on a CRT monitor with a refresh rate of 75Hz. All animal procedures complied with 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.

Task

Monkeys were trained to perform a choice-reaction time task in which they evaluated a sequence of shapes to choose the better of two saccade targets. The monkeys were trained extensively to interpret eight shapes as evidence bearing the likelihood that one or the other choice target would furnish a reward. Monkey J learned that half of the shapes favored the red or green target, whereas monkey E learned that half of the shapes favored the left or right target. For ease of presentation I refer to these targets as Ta and Tb, and in both situations one or the other target was placed in the RF of the LIP neuron under study.

After acquiring fixation on a central fixation point (FP), two choice targets appeared in the periphery equidistant from the FP. After a random delay, a sequence of highly visible shapes appeared every 250ms centered on a vertex of an invisible 3ºx3º grid centered on the FP. The shapes were approximately 1.5 by 1.5 deg, high contrast line art with equal perimeter (Fig. 2.1b). The shapes were displayed sequentially, until the monkey initiated a saccade to one of the choice targets. Successive shapes were not shown in the same grid location. The computer randomized which target would be rewarded on each trial, and sampled the eight shapes accordingly (Fig. 2.1b). Thus on any one trial, any of the 8 shapes could be displayed, but four were more likely if Ta or Tb was assigned the reward.

After the first shape's onset, the monkey was allowed to make a saccadic eye movement to one of the targets to indicate its choice. The time from the first shape's onset to the time of

saccade was registered as the reaction time (RT) in each trial. If the monkey chose a reward-assigned target, it received a drop of water or juice immediately (monkey J). For the second monkey (E), I imposed a minimum time of 1.8 s from the onset of the first of shape to reward in order to encourage this monkey to view more shapes. After the monkey indicated its decision with an eye movement, the video screen was set to black through an inter-trial interval (2.2 s) with an additional 3 s to penalize errors. Note that unlike Yang and Shadlen (2007), the more shapes the monkey observes, the more likely the monkey can choose a correct target. For example, an ideal observer given 10 shapes would choose correctly on 99% of trials.

For monkey J, I included 2 extra "trump" shapes that the animal had been trained on previously. These shapes, though rarely shown (~2%), would guarantee reward at one or the other choice target. However, the monkey did not always terminate decisions upon seeing these shapes, consistent with my previous inference that the monkey did not learn the significance of these shapes. Data using these "trump" shapes are not included in this report. I mention it here because the differences in training and task details bear potentially on the interpretation of the differences in results from the two monkeys. More details on logLR and related Baysian inference procedures can be found in Yang and Shadlen, 2007.

Data acquisition and neuron selection

Quartz-platinum/tungsten electrodes (Thomas Rec. / Alpha-Omega Co. with 1-3 MΩ impedance) were used for recording. Recorded signals from the electrode were amplified and

bandpass filtered (150Hz–8kHz for monkey E; 0.1Hz–10kHz for monkey J) before action potentials (spikes) were detected by a dual voltage-time window discriminator (Plexon and Bak Electronics).

I recorded from neurons in the posterior one third of LIP in the lateral bank of the intraparietal sulcus and in its ventral part (LIPv) (Lewis and Van Essen, 2000) near the fundus of the sulcus. LIPv was identified by both its anatomical and physiological properties. Recording sites were planned using a postoperative MRI that shows my recording chamber and grid. I registered these images with the standard MRI supplied with the CARET software (Van Essen, 2002) and targeted the posterior one third of the flat-map representation of LIPv. I then identified LIPv using physiological criteria and selected isolated single units that showed a robust and spatially selective persistent activity while the monkeys waited to execute a saccadic eye movement towards a remembered target location during the memory period of the memory-guided saccade task (Hikosaka and Wurtz, 1983). The property is common in the region of the IPS ~4-8 mm below the cortical surface. The criterion is qualitative and applied on line, but it is typically confirmed in delayed saccades performed during the main experiment.

Analysis of behavioral data

I used the logistic regression to estimate the leverage of the shapes on the monkey's choice. The probability of choosing the "A" target [*P(A)*] was modeled as a function of a "sum of leverages", Q, that should affect P(A) monotonically:

$$P(A) = \frac{1}{1+10^{-Q}} \quad (1)$$

which is fit as a general linear model (GLM) with binomial error. To determine whether shapes that appeared in a specific time window affected the monkey's choices (Fig 2.2b), I considered all the other shapes shown in the trial as confounders:

$$Q = \beta_1 w_t + \beta_2 W \quad (2)$$

where w_t is the assigned logLR of shapes that appeared during the time window of interest, W is the total logLR of the rest of the shapes that appeared, and β_i are the fitted coefficients. β_1 quantifies the leverage of shapes in the specific time window after accounting for the influence of other shapes. To estimate the leverage of shapes presented at the beginning of the trial, w_t was obtained from logLR of the k^{th} shape (k = 1, 2, ...) (Fig. 2.2b left). To estimate the leverage of shapes presented at a time t relative to the saccade, w_t was obtained only from trials in which a shape appeared withing ±20 ms of t. This analysis was repeated by sliding the time window in 10ms steps (Fig. 2.2b right). I also performed a variant of this analysis that incorporated a constant (bias term) at each time point. It is difficult to justify an independent bias at each time point, but it is reassuring that the results are consistent with the one I report based on Eq. 2.

To estimate the subjective weight that the monkey assigned to each of the eight shapes (Fig 2.2f), I fit 8 coefficients:

$$Q = \sum_{i=1}^{8} \beta_i n_i \,. \qquad (3)$$

where n_i is the count of each unique shape during a trial (omitting the final shape as explained elsewhere). The fitted logistic coefficients (β_i) quantify the leverage of each unique shape in units of logLR.

Analysis of physiological data

I calculated the firing rate of LIP neurons by convolving spike trains with a causal kernel: $\alpha(t) = t / \tau_K^2 \cdot \exp(-t / \tau_K), \quad \forall t \geq 0$ (τ_K = 0.01s). I formed population averages from these traces. To establish the epoch in which a shape influences the LIP response, I examined the time course of the response to the first shape to identify the point, τ_s, when the responses begin to diverge as a function of evidence.

Since shapes were presented every 0.25 s, the epoch associated with the k^{th} shape in a sequence extends from $\tau_s + 0.25(k-1)$ to $\tau_s + 0.25k$ seconds. The change in the firing rate induced by the k^{th} shape was estimated from the mean firing rates relative to the one preceding it:

$$\Delta r_k = r_k - r_{k-1} \qquad (4)$$

where r_k is the mean firing rate at the end of the k^{th} shape epoch (at $\tau_s + 0.25k$ s). For the first shape, the baseline firing rate (r_0) was estimated from the mean firing rate during τ_s s after the first shape onset. I computed $\Delta r_{i,k}$ for each unique shape i ($i \in \{1,\cdots,8\}$), and pooled them across shape epochs to obtain a set of Δr_i. The mean values of Δr_i are shown in Figure 2.3c and 2.3d.

Bounded Accumulation Model

The model presented in Figures 2.5 adopts the bounded evidence accumulation framework to explain choice and reaction time. My strategy was to constrain the model as best I could, using neural data, subject to a variety of simplifications in the interest of parsimony. The models were evaluated separately for the two monkeys. The basic structure is a race between two accumulation processes. One of these uses the ΔFR, average changes in the firing rate in my data set 'detrended' by their grand average:

$$\Delta FR_i = \langle \Delta r_i \rangle - \langle \Delta r \rangle \qquad (5)$$

where Δr represents a pool of Δr_i across the eight unique shapes. This detrending was necessary to remove the effect of evidence-independent signal, u(t) from Δr (see below). The largest ΔFR is associated with the triangle, the shape that furnishes the strongest evidence for the choice target in the neuron's RF (i.e., Tin), whereas the smallest ΔFR accompanies the T-shape, the shape that provides the strongest evidence for Tout. For the competing process, I reversed the association

between shape and ΔFR. For this process, the largest ΔFR is induced by the T-shape and the smallest by the triangle. I use the same 8 values of ΔFR. I recognize that this is an oversimplification (I have recorded from only a few neurons with ipsilateral RFs), but it allows me to constrain the model with neural responses instead of fitted parameters for the influence of the shapes on an abstract decision variable.

These ΔFR_i values can be thought of as the expected momentary evidence supplied by the shape. I assume that the actual change induced by each shape on each trial is both noisy and dynamic. To achieve this, I add an unbiased sample of a Normal distribution $\varepsilon \in N(0,\sigma_\varepsilon)$ to the expectation and allow the sum to scale an impulse of evidence, x(t), illustrated in Figure 2.5c. I approximate the template as the integral of an exponential decay

$$x(t) = \begin{matrix} 1-\exp\left(-[t-\tau_s]/\tau_\Delta\right) & for\ t \geq (t_k+\tau_s) \\ 0 & otherwise \end{matrix} \qquad (6)$$

where t_k is the time of onset of the k^{th} shape, τ_Δ=50 ms and τ_s is the delay described above. Note that both tau terms are approximated to the neural data. I used the same τ_Δ for the two monkeys.

The firiing rate on a trial is described by an accumulation of noisily scaled impulses added to a time-dependent, evidence-independent baseline, termed the urgency signal, u(t). This is evident in the rising firing rate for the neutral quantile in Figure 2.3a for monkey E whereas it

was negligible for monkey J. To model the accumulation process, it is convenient to represent x(t) and u(t) by their derivatives. This yields the stochastic differential equation:

$$\begin{aligned} \dot{r} &= \dot{u}_k + (\Delta FR_k + \varepsilon_k)\dot{x} \\ &= \dot{u}_k + \dot{x}\Delta FR_k + \dot{x}\varepsilon_k \qquad for\ t \in [t_k, t_{k-1}] \end{aligned} \tag{7}$$

In this equation, the value of $\dot{u}_k$, ΔFR_k, and ε_k are updated every 0.25 s, and only the last term contains the stochastic elements, which are independent and identically distributed, described by a single free parameter σ_ε. I reasoned that the amount of urgency signal differs from epoch to epoch, and estimated the amount of urgency signal added in the k th epoch with the grand average of changes in the firing rate during that epoch:

$$u_k = \langle \Delta r_k \rangle \tag{8}$$

I implemented a non-absorbing (i.e., nonterminating, also known as reflecting) lower bound at zero for both stochastic processes so that r(t) does not take a negative value. The boundary conditions are derived from the physiology. The firing rate before the first shape-induced change (i.e., $t < \tau_\Delta$) was approximated by the mean firing rate in these intervals (20 and 30 Hz for monkey E and J, respectively). The process stops when one of the rates exceeds a threshold (θ). The value of θ was based on an estimate of the firing rate at a point of convergence preceding Tin choices (Fig. 2.4). I obtained this estimate by first grouping the responses into 10 quantiles based on the cumulative logLR at τ_s + 100ms prior to that time point. I smoothed each

quantile response curve by running mean (50ms window), and computed the standard deviation of the 10 smoothed response curves (Fig. 2.4, inset). During the pre-saccadic period, this standard deviation reached its minimum shortly before the time of saccades (Fig 2.4, arrows). For both monkeys, the mean firing rate establishes the threshold for terminating decisions (θ = 50 Hz for both monkeys) and an approximate latency from decision termination to saccade in initiation, which I interpret as a motor latency ($\overline{\tau}_m$, ~200 ms and ~50ms for monkey E and J, respectively). I model τ_m as gamma distributed with mean $\overline{\tau}_m$ and standard deviation σ_m, which is a free parameter in the model. The predicted RT distribution is the convolution of this gamma distribution with the predicted distribution of decision time times. I am least secure about the estimates of $\overline{\tau}_m$, in part because of the discrepancy between monkeys and in part because the analysis depicted in Figure 2.4 does not identify a unique termination point. However, this concern is mitigated somewhat because I assume that the motor latencies are themselves variable. Thus, I model τ_m as gamma distributed with mean $\overline{\tau}_m$ and standard deviation σ_m, which is a free parameter in the model.

To fit the free parameters of the model, σ_ε and σ_m, I used Monte-Carlo method to simulate trials and used the maximum log likelihood method to fit the RT distributions for correct and error trials (Fig. 2.6a, black histograms). I simulated 10,000 repetitions per iteration and used the same random number seeds for each choice of parameters σ_ε and σ_m. Fitted parameters are shown in Table 2.1.

Results

Behavior

Two monkeys made decisions between two peripheral targets (Target A and B) based on a sequential presentation of shapes (Fig. 2.1a, see Methods). Each presented shape conferred probabilistic evidence that differed in reliability in predicting which target would furnish a reward. Although each of the eight shapes could be shown on any trial, they were sampled differently depending on the reward-containing target (Fig. 2.1b). Four shapes favored one of the targets ("Target A") because they were more likely to appear when Target A was assigned the reward as opposed to Target B (Fig. 2.1b). The other four shapes favored Target B. Each shape's evidence can be quantified by logLR whose sign indicates a favored target and whose magnitude represents the reliability of evidence. While the monkey maintained fixation, shapes appeared parafoveally every 250 ms, one at a time, until the monkey indicated its decision with a saccadic eye movement to one of the two targets (Fig. 2.1a).

As shown in Figure 2.2a, both monkeys displayed a wide range of reaction times (RT) on this task. The serrated quality of the histograms is consistent with the rate at which the shapes were shown (4 Hz), and there is considerable variability that blurs this discrete schedule. Evidently, monkey J made some decisions based on only a few shapes, but both monkeys often delayed their responses until after viewing several shapes. I therefore wished to ascertain

whether they used information from all the shapes or waited for a particular shape, or perhaps used information from a favored epoch.

To address this empirically, I performed an analysis to ascertain whether the evidence (i.e. shapes) presented at the beginning through the end of the trial had leverage on the choice. The leverage is summarized by a single weighting coefficient (β_1, Eq. 2 in Methods; logistic regression) displayed in Figure 2.2b. The left half of these graphs displays the effect of evidence presented at the beginning of the trial on the monkey's ultimate choice. Because the shapes appear (and disappear) every 250 ms, this portion of the graph contains only three discrete values, corresponding to the onset times of the first three shapes. Shapes presented in each of these epochs had coefficient significantly above zero, indicating that they affected the monkeys' choices ($p<0.01$). The right half of the graph plots the leverage of the shapes at the time they appeared relative to the saccadic choice. Since saccade initiation is not fixed to the time of appearance of a shape, I applied logistic regression at finer time steps. At each time point, I extracted the leverage of shapes that were presented within ±10 ms of the time point displayed on the graph. Importantly, shapes that were presented long before the saccade had leverage on the choice (up to ~1.5s and ~0.75s for monkey E and J, respectively; $p<0.05$), indicating that early shapes affected the choice even when they were followed by multiple shapes. The analysis also demonstrates that shapes that are shown within the last 200-300 ms had negligible effects on the choice (last 270 ms and 180 ms for monkey E and J, respectively; $p>0.05$). This observation implies that, in most of trials, the monkeys had committed to a decision before the evidence from the last shape was incorporated into the deliberation process. For monkey E, this non-decision

interval was long enough to allow for the possibility that on some trials even the penultimate shape arrived after commitment (i.e. if a saccade occurred <20 ms after the final shape).

I applied these adjustments (see Methods) to generate the "number of samples" histograms in Figure 2.2c. These discrete RT distributions depict more clearly the number of shapes the monkeys used to reach their decision, which I term N*. Although monkey E used more samples than monkey J (mean 5.3±1.4 vs. 2.2±1.5 shapes, respectively), both monkeys exhibited considerable variance in the number of shapes they relied upon. A possible reason for this will be clear in a moment. Importantly, when either monkey's decision was based on more than one shape, the entire succession of shapes, from the first to the N^{*th}, was influential. From here on, I perform all analyses on these shapes.

By basing their decisions on more than one shape, the monkeys improved their performance. Had they based their decisions on just one shape, even perfect knowledge of the likelihoods would achieve only 63% correct (i.e., rewarded) choices, owing to the overlap of the sampling distributions (Fig. 2.1b). However, monkeys E and J achieved 85% and 80% correct choices, respectively. Perhaps a more telling observation is that the fraction of trials in which the monkey chose the option supported by the total logLR from the shapes —regardless of what was actually rewarded —were 90% and 97% of trials (E and J, respectively; Fig. 2.2d). More striking yet, the accumulated logLR was remarkably similar, on average, whether the monkeys used a few or many shapes to make a decision (Fig. 2.2e). The consistency of this mean as a function of the number of shapes suggests that that the decisions might terminate when the representation of

accumulated logLR reaches a threshold level. I infer that the representation is noisy, however, because the standard deviations (error bars in Fig. 2.2e) are substantial. Although this threshold does not appear to be perfectly flat ($p<10^{-5}$; weighted regression; dashed lines, Fig 2.2e) the deviation is subtle and markedly inconsistent with what would be expected if the reaction time were independent of the accumulated evidence (Fig. 2.S1).

Finally, I assessed how much weight the monkeys assigned to each of the eight shapes (Fig. 2.2f). For this analysis, instead of extracting a weighting of the evidence at each time, I used all the shapes presented in the trial (up to N^{*th}) to extract the relative weight given to each of the individual shapes ($\beta_{1\text{-}8}$, Eq. 3). Although the estimated weights differ from the assigned weights, they have the correct sign, and they are ordered appropriately with just one exception for each monkey ($p<10^{-3}$ for both monkeys, Spearman's rank correlation test). From this, I conclude that both monkeys successfully learned to assign an appropriate weight to each shape.

Physiology

I recorded from single neurons in the ventral part of the lateral intraparietal area (LIPv) while the monkeys performed the RT-classification task in Figure 2.1. I selected all well-isolated neurons that exhibited spatially-selective persistent activity during memory guided saccades (n=62; 38 in monkey E; 24 in monkey J). During recording, I mapped the response field (RF) of a recorded neuron, and placed one choice target inside the RF (T_{in}) and the other choice target outside the RF (T_{out}). The two targets were in symmetric positions around the fixation point. The

targets differed in color for monkey J, whereas both were red for monkey E, reflecting the two training strategies (see Methods).

Shortly after the presentation of the first shape, LIP activity began to reflect the accumulated evidence in favor of the choice associated with the neuron's RF. Figure 2.3a shows the evolution of the firing rate as a sequence of up to 5 shapes were presented. To make this graph, I grouped trials into five quantiles based on the cumulative assigned logLR in each epoch. To avoid contamination of these averages with activity near saccade preparation, the averages exclude the firing rate induced by the N^{*th} shape. Note that any one trial will tend to contribute to different quantiles in each epoch, and while almost all trials contribute to the responses to the 1st shape epoch, the number declines in the later epochs. There are three salient observations. First, the LIP firing rate began to diverge as a function of logLR after a short delay from the first shape onset (200ms and 130ms for monkey E and J, respectively; Fig. 2.3a 1st epoch). I refer this delay as a "sensory" delay, τ_s—in the sense that it reflects the time for the presented evidence to reach LIP through a sensory pathway—and use this delay in subsequent analyses. Second, with each successive shape, the range of the cumulative logLR expands, and this is reflected in the separation of the firing rates, grouped by quintile of logLR. Third, for monkey E, even when the evidence is neutral (middle green quantile), there was a gradual increase in the firing rate as a function of time. Such evidence-independent rise is thought to incorporate temporal cost into the DV (Churchland et al., 2008; Drugowitsch et al., 2012), as explained further below.

The quintile grouping helps to simplify the graphs, but it masks a more refined relationship between the gradation of evidence and LIP firing rate. The insets in Figure 2.3a show that the firing rate varies approximately linearly with the cumulative logLR throughout its range. Each subplot shows the mean (and running mean) of the firing rate as a function of cumulative logLR in each epoch. Although the shapes appear for only 250 ms and then disappear, the firing rate reflects the accumulated evidence through the penultimate shape used by the monkey. Notice that the correspondence between firing rate and logLR is stable for the first 5-6 shape epochs, which accounts for ~80% of the trials. For later epochs, the slope attenuates. As will be clear later, this is readily explained by conditionalization — that is, selection of the trials that required more evidence to terminate.

During decision formation, each new shape adds a quantity of evidence and should therefore affect the LIP activity accordingly. I calculated the change in the LIP population response (Δr_i, see Methods) induced by each shape and compared these increments and decrements to the logLR values assigned to the eight shapes (Fig. 2.3c,d). Whereas shapes with larger logLRs were associated with larger increases in the firing rates, shapes with negative logLR values were associated with decreases in firing rate for Monkey J. In Monkey E, however, this decrease was countered by the evidence-independent trend of rising activity noted above. Although the correspondence between ΔFR and assigned logLR is only qualitative, there is a striking similarity between ΔFR and subjective weights deduced from behavior (Fig. 2.2f, 2.3c,d, $\rho = 0.99$ and 0.98 for monkey E and J, respectively; $p<10^{-5}$ Pearson's correlation test). The imperfect learning of the shapes' logLR is faithfully reflected by the LIP neurons' activity.

The reaction-time task allows me to examine the neural events associated with termination of a decision. In contrast to the divergence of responses associated with evidence accumulation, the activity near the end of the decision process appears stereotyped. Figure 2.4 shows the average firing rates sorted by the choice that the monkey made (T_{in} or T_{out}) and by the relative support for that choice (grouped into sets of quintiles). By aligning the responses to saccade initiation, the appearances of the shapes are not aligned as they were in Figure 2.3a. The averaging is therefore somewhat less straightforward. Using the sensory delay (τ_s) assessed in Figure 2.3a, I first established the cumulative logLR contributing to each time point in each trial, and then formed the quantile groups from trials in which the onset of a shape would be expected to affect that time point (see Methods for details). This grouping allows me to see whether the neural responses at each moment reflected the cumulative evidence.

For T_{in} choices, the diversity of neural activity associated with strength of evidence gradually dissipated near the time of the saccade. We interpret this as a sign of a threshold or bound that would terminate decisions in favor of T_{in}. Notably, responses converged to the same level of threshold for both correct and error Tin choices ($p > 0.2$), suggesting that error choices occurred because the LIP firing rate misrepresented the cumulative logLR. For T_{out} choices, the neural activity did not converge, suggesting that activity of other neurons, presumably those with the RFs aligned to T_{out}, precipitated termination of these decisions. The effect of cumulative evidence on firing rate can be captured at each time point by regressing the firing rate (associated with one choice) against the cumulative logLR up to that time point minus τ_s. For trials ending in

T_{in} choices, this modulation is no longer detectable before saccade initiation (red arrows, Fig. 2.4a insets; ~210ms and ~50ms before saccades for monkey E and J, respectively, $p > 0.05$), consistent with a bound or threshold. We interpret the time from the red arrows to saccade initiation as an average motor latency from decision termination to saccade initiation. While these latencies are supported by other analyses (e.g., decline in firing rate variance), we feel they are at best rough approximations (e.g., for monkey J, there is a competing nadir in the regression slope to the left of the red arrow). Nonetheless, the loss of evidence dependence is consistent with the idea that the decision process ends when firing rate reaches a threshold level.

A possible concern is that the convergence of LIP firing rate before T_{in} saccades is simply a saccadic burst rather than a sign of decision termination based on evidence accumulation. To test this possibility, I sorted the peri-saccadic responses based on the evidence from the last effective shape (N*). For T_{in} trials, the firing rate ramped up to a stereotyped value after the onset of the last shapes (Fig. 2.S2a,b). Importantly, the excursion of these ramp-ups were larger when the last shapes conferred stronger evidence for T_{in}, indicating that the last shape induced a logLR-dependent change in the firing rate to the threshold ($p<0.001$; Fig. 2.S2c, d). This implies that the evidence induced the change in firing rate that I measure as stereotyped because it signals — to other neurons — the termination of the decision.

Finally, it is worth noting that the graded representation of evidence in Figure 2.4a, even when trials are grouped by choice, implies that the graded responses evident in Figure 2.3a do not arise as a consequence of averaging responses comprising different mixtures of T_{in} and T_{out}

responses. Thus, taken together, Figures 2.3a and 2.4a imply that LIP neurons represent an evolving decision variable that will ultimately lead to a T_{in} or T_{out} choice.

Model

The observations from both behavioral and neural recordings suggest that decisions are made when a sufficient amount of evidence has been accumulated in favor of one of the alternatives. For example, Figure 2.2e indicates that on average, deliberation ends when the absolute value of the cumulative log likelihood reaches a common level, regardless of the number of shapes it takes to achieve this level. Furthermore, the stereotyped firing rate before T_{in} choices also supports a stopping rule based on the neural representation of this cumulative evidence. However, there are also anomalous observations that cannot be overlooked. Across trials, there is considerable variation in the state of the accumulated evidence at termination. This is evident in the standard deviations depicted in the same Figure 2.2e. Indeed the choice function displayed in Figure 2.2d spans a range of cumulative evidence levels, which at face value contradicts the premise that all decision stop at or very near a positive or negative threshold. I hypothesized that these contradictory observations can be explained by the noisy representation of logLR by LIP neurons and a termination mechanism that operates upon the LIP firing rate and not on logLR directly.

I modeled the decision process as a race between two accumulators: one accumulating evidence in favor of Target A and the other accumulating evidence in favor of Target B. The

model structure, illustrated in Figure 2.5, exploits values obtained from the neural recordings (Fig. 2.3c,d), as follows. Each shape produces an expected increase or decrease in the firing rate adduced from the average firing rates in the data (detrended; see Methods). These expectations, $\langle \Delta FR_i \rangle$, (Fig. 2.5b) are assumed to be of opposite sign for the two accumulators. I assume that each shape induces an actual change in firing rate, ΔFR_i that is determined by this expectation plus unbiased noise (standard deviation σ_ε; Fig. 2.5b, orange distribution), which is assumed to be independent in the competing processes. Sampled values of ΔFR_i scale a temporal impulse shape, thereby capturing the delay and dynamics of the response to single shapes (e.g., Fig. 2.3a,b). The noisy impulses are integrated and added to a time-dependent urgency signal, u(t), which is also derived from the physiology (Eq. 8, Fig. 2.S3). The process terminates when one of the accumulators reaches a threshold firing rate level, θ, thereby determining the choice and decision time on a trial. The measured reaction time is the decision time plus a gamma distributed delay, τ_m, representing the component of "motor delay" that is not incorporated in the impulse shape. I fit this model to the choice and reaction time data for each monkey using just two degrees of freedom (σ_ε and σ_m; Table 2.1). See Methods for additional details.

I fit this model to the choice and reaction time data for each monkey. In each case, the model had just 2 degrees of freedom (Table 2.1): σ_ε, the standard deviation of the noise added to the $\langle \Delta FR_i \rangle$ terms, and the standard deviation of τ_m (σ_m). In principle, one might expect that σ_ε could be measured from the firing rate variability of the recorded neurons. However, single

neurons are more variable than a population of neurons, and unfortunately the degree of reliability in the population depends on factors I did not measure, such as the degree of correlated variability among neurons (Zohary et al., 1994; Shadlen et al., 1996).

The model achieves a reasonable fit to the choice and RT distributions for both monkeys (Fig. 2.6a). Furthermore, it explains the anomaly mentioned above, namely the variation in the state of the accumulated evidence at decision termination (Fig. 2.6b,c). It does so by postulating that decisions terminate when a population firing rate reaches a critical level, and these firing rates are noisy reflections of the underlying evidence.

The model also furnishes insight into the differences between the behavioral patterns exhibited by the two monkeys. Monkey E assigns smaller changes in LIP firing rate to the shapes than monkey J, and this is compounded by the larger noise term (Table 2.1). However, monkey E compensates by applying a more conservative bound. That is, the difference in firing rate from beginning of evidence accumulation to the terminating bound is larger in monkey E, consistent with the longer RT, or equivalently, the larger number of sampled shapes, on average. This excursion also incorporates an evidence-independent, time-dependent "urgency" signal, u(t) (Fig 2.5d), which is thought to instantiate a "cost of elapsed time" (Drugowitsch et al., 2012). It implies that as time elapses, the amount of evidence required to terminate the decision is reduced and thus explains the decline in accuracy as a function of RT (Fig. 2.6d). Interestingly, these differences led monkey E to a higher overall accuracy rate (85% vs 80% correct) but to similar reward rates (0.179 vs. 0.171 s^{-1}).

Finally, the model predicts the linear relationship between the firing rate and the cumulative logLR (Fig. 2.S4). As has been observed in the physiology (Fig. 2.3a), the slope of the linear regression lines remains stationary for earlier epochs but gradually flattens out for later epochs. Based on my model, this is because the monkeys tend to use more shapes in trials in which the evidence changes the firing rate by a smaller magnitude than average. The agreement between the physiology and the model confirms that LIP response represents the accumulation of discrete changes, where each change reflects logLR of a shape with noise.

Discussion

When a decision is based on samples of evidence of different reliability, then it is rational to give more weight to the more reliable cues (Ernst and Banks, 2002; Fetsch et al., 2013). If the evidence arrives sequentially and more is expected, then the decision maker might also decide when to terminate the process and commit to a choice. Here, the notion of rationality is less well prescribed, as it depends on the desired level of accuracy and the cost of time. In many instances the speed-accuracy tradeoff is thought to be explained by a common mechanism in which evidence is accumulated to a threshold level or bound. Previous studies of perceptual decision making have demonstrated a neural correlate of both evidence accumulation and termination in area LIP. In these studies the reliability of evidence was statistically stationary during a trial, and as best I can tell, all samples of evidence, during the course of a single decision are weighted identically. On the other hand, while many behavioral studies have demonstrated a rational

combination of cues of different degrees of reliability, nearly all such studies present the cues simultaneously. The one exception, to my knowledge, is the study upon which the present experiment is based.

Yang and Shadlen (2007) demonstrated that monkeys were capable of combining evidence from a sequence of four shapes, like those used in the present experiment, by assigning greater weight to the more reliable cues. Although the LIP responses in that study appeared to encode the cumulative evidence from the sequence of shapes, this was the only support for the conjecture that the monkeys actually reasoned sequentially from the four samples. The monkey was required to make his only choice after all four shapes were visible on the display.

The present study enforced a sequential strategy in two ways. First, the task requires integration in time because only one shape is visible at a time. Second, the monkey was free to terminate each trial when it had seen enough shapes. The behavioral data indicate that shapes from early epochs influenced the monkeys' choices, even on trials in which they are followed by many more shapes. Moreover, the speed and accuracy of the monkeys' decisions are well described by bounded evidence accumulation. Thus monkeys demonstrated a capacity to form decisions based on sequential analysis of evanescent cues of varying reliability.

A well known strategy to solve this class of decisions, known as the sequential probability ratio test (SPRT), would place termination criteria on the cumulative log likelihood supplied by the shapes. The level of these criteria determines the error rates and the distribution

of the number of shapes supporting decisions, across trials. In fact SPRT is optimal in the following sense: given a desired error rate, it ensures that the number of samples is minimized on average. A natural candidate for "desired error rate" is one that would maximize the rate of rewarded decisions over many trials (Gold and Shadlen, 2002; Simen et al., 2009). Neither of the monkeys achieves this optimal level, but they came close — Monkey E and J achieved a reward rate that is 87% and 81% of the optimal reward rate.

However, it is probably incorrect to assume that monkeys even care about maximizing the reward rate. As has been previously shown, decision time itself appears to be costly to monkeys. For example, when random dot motion stimuli are presented for durations exceeding 0.5 s, monkeys (and humans) tend to ignore late information, as if performing a RT task (Kiani et al., 2008; Tsetsos et al., 2012). This time cost can be realized by collapsing the termination bounds as a function of time, or equivalently, by adding a monotonically rising firing rate as a function of time, termed urgency, to the accumulated evidence represented by competing accumulators (Ditterich, 2006; Churchland et al., 2008; Cisek et al., 2009; Drugowitsch et al., 2012). The signature of such a process is the declining accuracy as a function of RT for monkey E (Fig. 2.6c). The neural correlate of this is the gradual rise in firing rate that is apparent even when the accumulated evidence favored neither choice (Fig. 2.3a, green trace).

Other explanations for the less than optimal performance are more informative about the underlying process. The most obvious is that neither monkey assigned the 8 shapes the correct weights. This leads to systematic error in the choices, although the effect is modest. For example,

if I assume weighting like that shown in Figure 2.2f, and apply the optimal bound, the overall rate of reward would be reduced by only ~3%. This reduction is due to imperfect assignment of relative weights (Fig. 2.2f) rather than the absolute values of weights. Another discrepancy highlighted by the model is that the firing rate reflects the accumulation of evidence, but it is corrupted by noise. This introduces a diminution in SNR that can only be overcome by increasing integration times (e.g., a higher bound setting) at the cost of more time per trial.

The same idea resolves the paradox inherent in the behavioral observations. On the one hand, the monkeys' choice and RT functions suggest an accumulation of evidence in units of logLR to a terminating threshold level. On the other hand, the level of evidence at termination is highly variable, which seems entirely inconsistent with such a threshold. The resolution is simply that the brain does not accumulate in units of logLR but in units of firing rate that supply a noisy and imperfect rendering of the cumulative evidence. The modeling exercise (Fig. 2.6d) thus reconciles quantitatively the neural and behavioral measurements. One might argue that the brain is incapable of achieving the optimal solution to this problem, but given limitations in training and proximity to optimal performance, I tend to take the rosier view that despite limitations of biology (e.g., noisy neurons) and learning, the brain can achieve a reasonable approximation to optimal cue combination.

I have mainly focused on aspects of the decision processes that were shared by the two monkeys, but the model also sheds light on differences between the monkeys, which may have been induced by their training histories. I introduced a delay to reward during training in order to

discourage very fast RT, which is the natural tendency of many monkeys on RT tasks (e.g., see Roitman and Shadlen, 2002). I do not know if this training manipulation also led monkey E to give less weight to the shapes or if instead the lower signal to noise led to compensatory change in termination criterion. What I can say, however, is that the combination of bound and weight conferred monkey E with the better accuracy than monkey J.

Another difference between the monkeys can be attributed more clearly to a difference in their training regimen. Monkey J had been trained extensively on a version of the task in which the shapes were associated with colored targets. In an earlier report, there was a hint that the neural responses were more reliable when the red choice target was displayed in the RF. This trend evolved during training on the RT task such that the monkey now appears to use only one population of LIP neurons on this task. That is, when the green target was in the RF, though the monkey performed the task equally well, the neurons did not modulate their activity until 200ms before the saccade, thereby rendering about half of the trials uninformative for physiology (Fig. 2.S5). To avoid repeating this dilemma, I trained monkey E to associate the shapes with the location — left or right — of identical red targets. The symmetrical design ensured that all trials were informative, and this was supported by the occasional neuron with an ipsilateral RF.

I considered and rejected a variety of alternative models to explain the monkeys' behavior. First, the monkeys did not base their choices or decision times on the occurrence of particular shapes. With this strategy, the subjective weight should be zero for all shapes except for those particular ones, whereas each unique shape was assigned a non-zero weight with an

appropriate sign and rank order (Figure 2.2f). Second, the monkeys did not rely on evidence conferred in a particular epoch during a trial because shapes impacted the choice whether presented early or late (Fig. 2.2b). Third, decision times were not dictated by some type of deadline that is independent of the state of the cumulative evidence. This idea would not explain the consistent level of evidence at termination (Fig. 2.2e; see Fig. 2.S1 for further explanation).

That said, the model I proposed is undoubtedly inaccurate in detail. I do not know how the monkeys learn to associate differential weights to the 8 shapes, and it is not clear how the appearance of a shape leads to an increment or decrement in the LIP firing rate that is of the appropriate magnitude for that shape. Comparison to the well studied random dot motion task is instructive. In that task, neurons in extrastriate area MT (V5) represent the momentary sensory evidence, and the accumulation is thought to be of the difference in firing rates of such neurons tuned to opposite directions. The noisy difference signal is itself proportional to units of logLR (Gold and Shadlen, 2001). I assume that the shapes are differentiated in extrastriate cortical areas in the ventral stream (Tanaka, 1996; DiCarlo et al., 2012), but there is no natural subtraction operation that would pit the star against the pentagon shape, as assumed in my model (but see Soltani & Wang, 2009).

What seems to be required is an operation resembling memory retrieval in which the shape is associated with a quantity that bears on another process. In memory retrieval the quantity bears on decision about satisfaction of a match (e.g., similitude) (Ratcliff, 1978a; Ratcliff and McKoon, 2008; Wimmer and Shohamy, 2012; Shadlen and Kiani, 2013), whereas in

my experiment it bears on the relative merit of a left/right or red/green choice. Perhaps this insight bears on the success of sequential sampling models to a variety of cognitive tasks, including memory retrieval, that do not appear to involve any obvious need to integrate independent samples of evidence as a function of time.

Figure 2.1
Sequential inference task
a, Trial flow. The monkey stares at a central fixation point at the beginning of a trials, and two choice targets appear in symmetric peripheral locations around the fixation spot. After a short latency, a sequence of shapes is shown in succession every 250ms. Each shape supplies evidence bearing on whether a reward is associated with one or the other choice target. The sequence continues until the monkey initiates an eye movement to a choice target.
b, The sequence of shapes shown on a trial are independent, random draws from either of the two discrete sampling distributions (indicated by color). On each trial, the designation of rewarded choices is randomized to the A or B target, and the shapes are sampled from the designated distribution. The ratio of sampling probabilities associated with each shape implies that half of the shapes support one or the other choice (red and blue rectangles). Target A refers to the left choice target for monkey E and the red choice target for monkey J.

Figure 2.1

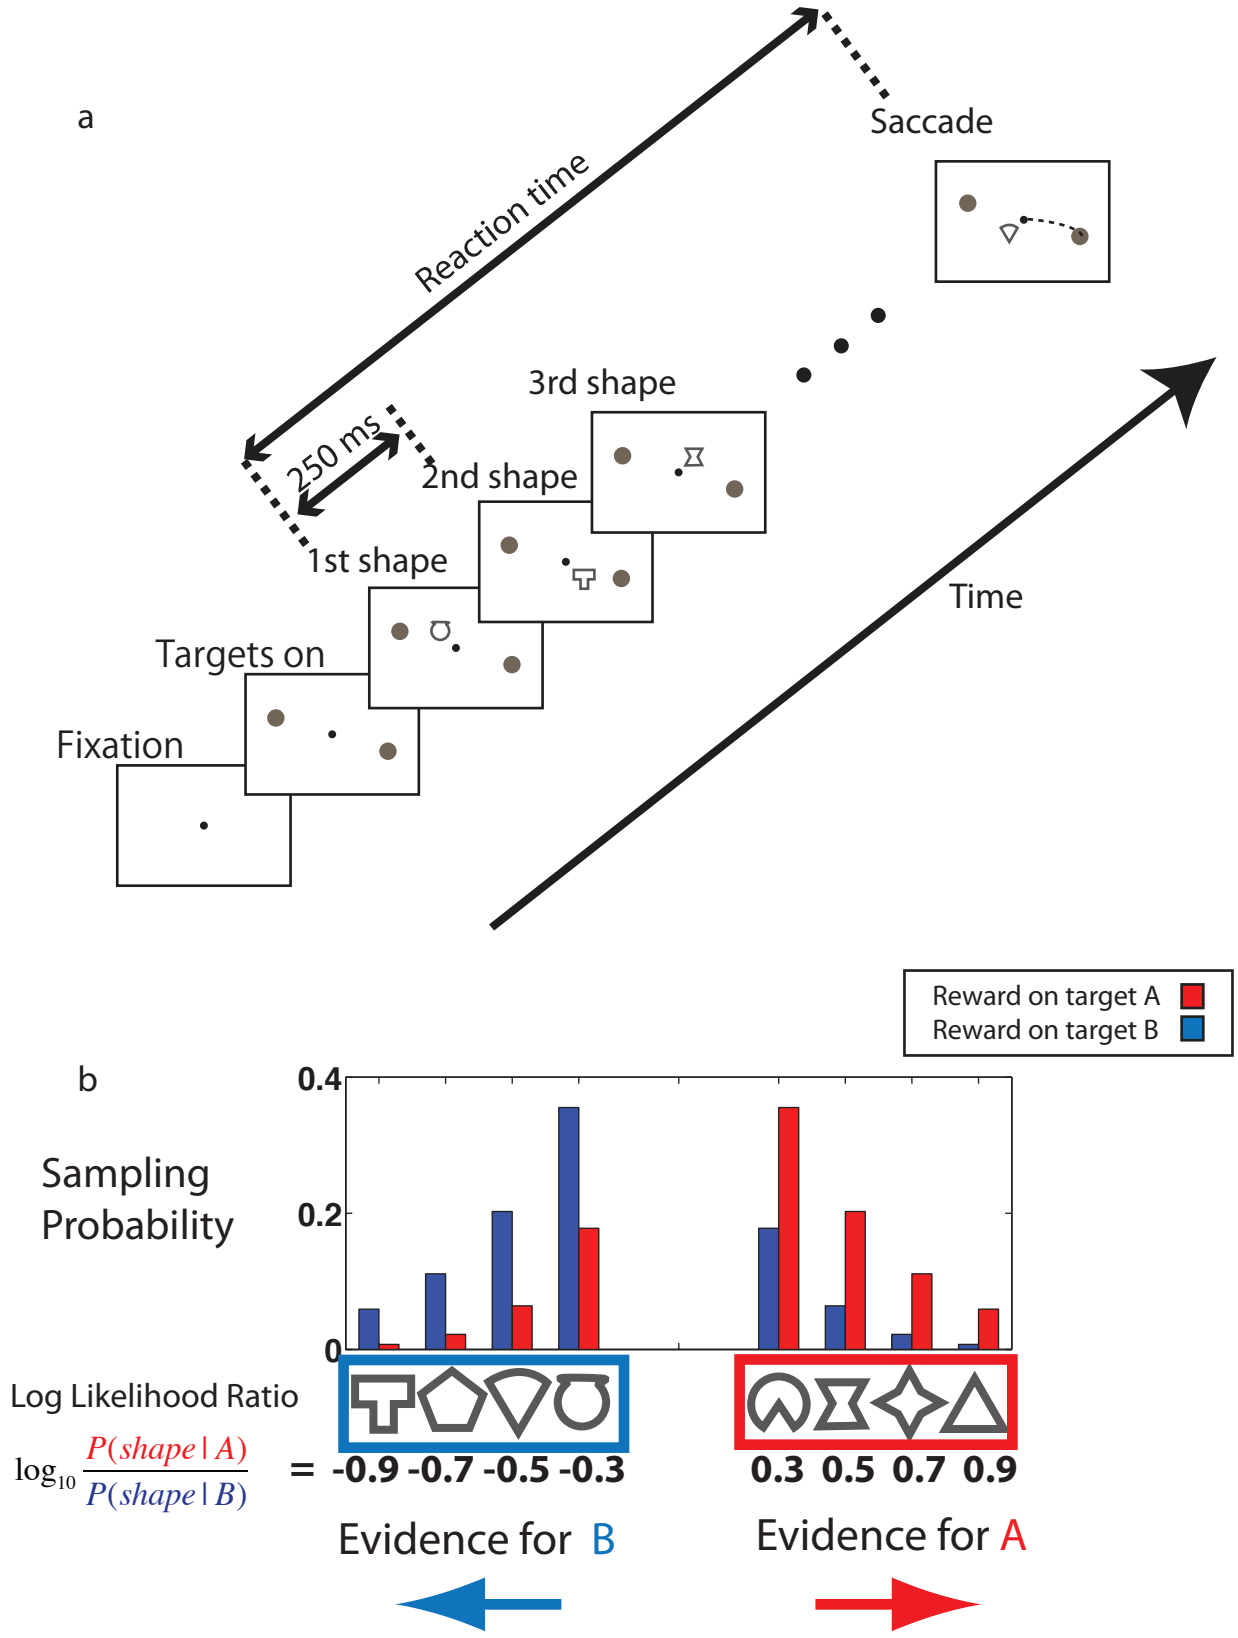
a
Saccade
Reaction time
3rd shape
250 ms
2nd shape
1st shape
Time
Targets on
Fixation
Reward on target A
Reward on target B
b
0.4
Sampling
Probability
0.2
0
Log Likelihood Ratio
$\log_{10}\frac{P(shape\,|\,A)}{P(shape\,|\,B)}$
= -0.9 -0.7 -0.5 -0.3
0.3 0.5 0.7 0.9
Evidence for B
Evidence for A

Figure 2.2
Behavior
The two monkeys are depicted in the left and right columns.
a, Reaction time distributions for the two monkeys. The RT is the time of the saccade relative to the onset of first shape. Shapes appeared at a rate of 4 Hz, until saccade initiation (histogram binwidth = 10ms).
b, All shapes, except the last, influence the monkeys' decisions. Graphs show the relative influence of shapes apearing within ±10 ms of the time indicated on the abscissa. Relative leverage is derived from a logistic regression (Eq. 1,2), which incorporates the other shapes shown on a trial as potential confounders. On the left portion of the graph, time is aligned to the onset of the first shape; shapes appear every 250 ms and onsets are shown by the thick line (shading shows s.e.). On the right portion of the graph, time is aligned to saccade onset. The leverage is calculated every 10 ms, using all trials in which a shape appeared within ±10 ms of the time indicated. Shading indicates s.e. of the leverage (β_1, Eq. 2).
c, Distribution of the number of shapes that affected the decision in each trial. This is a discrete version of the RT distribution adjusting for late shapes that did not influence the decision. I refer to this number of shapes as N* (see Methods).
d, Choice probabilities were governed by the total evidence supplied by the shapes. The total evidence is the sum of the logLR assigned to the N* shapes shown on each trial. Choice A refers to the left choice target for monkey E and the red choice target for monkey J.
e, Consistency of total evidence at the end of decision. Points are the mean of the total logLR from the N* shapes leading to choice A (circles) or B (squares). Error bars show the standard deviation (s.d.). Although the means exhibit consistency, the s.d. expose considerable variation.
f, Subjective weight the monkeys assigned to each unique shape is plotted against the true weight in units of logLR. The weights are the coefficients from logistic regression ($\beta_{1\text{-}8}$, Eq. 1,3). Error bars are s.e. Some error bars are smaller than the data symbols.

Figure 2.2

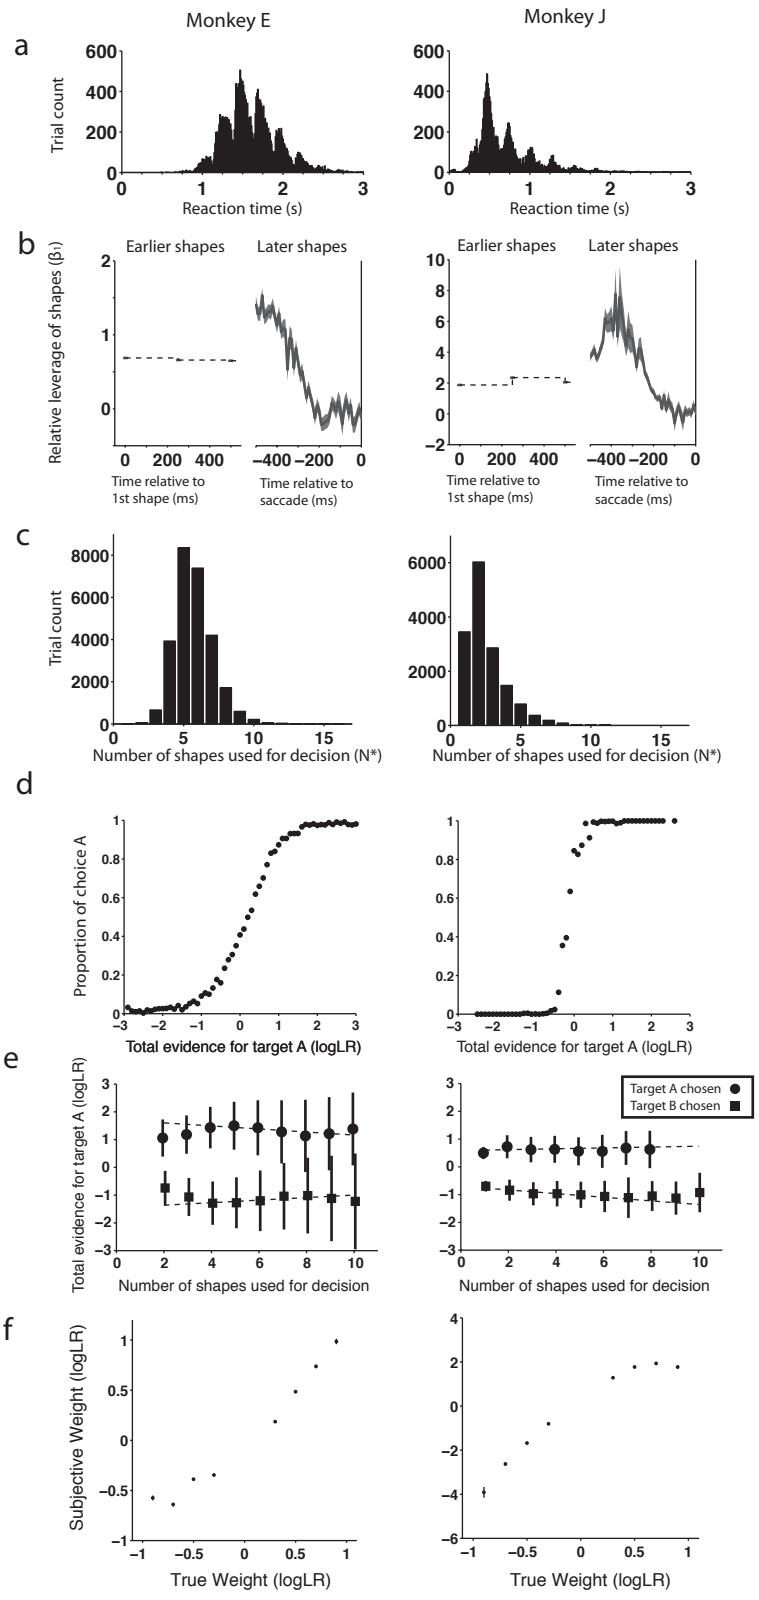
Monkey E
Monkey J
a
Trial count
Reaction time (s)
b
Relative leverage of shapes (β1)
Earlier shapes
Later shapes
Time relative to 1st shape (ms)
Time relative to saccade (ms)
c
Trial count
Number of shapes used for decision (N*)
d
Proportion of choice A
Total evidence for target A (logLR)
e
Total evidence for target A (logLR)
Number of shapes used for decision
Target A chosen
Target B chosen
f
Subjective Weight (logLR)
True Weight (logLR)

Figure 2.3
LIP neural activity accompanying decision formation.
a, **b**, Firing rate averages for reflect the cumulative evidence. Lower row shows average firing rates plotted as a function of time from the onset of the first shape. The five curves in each panel correspond to quantiles grouped by the cumulative logLR in favor of Tin. Quantiles are redefined in each epoch, denoted by the break between panels, which respect the sensory delay evident in the leftmost panel. Insets (top row) plot the mean firing rate at the end of the corresponding epoch as a function of total logLR. Color-filled circles represent the corresponding quantiles, and the gray trace shows the running mean ± s.e.
c, **d**, The average change in firing rate induced by each of the 8 shapes. Eight data points correspond to logLR assigned to each of eight unique shapes. The last effective shape in each trial was excluded from the analysis.

Figure 2.3

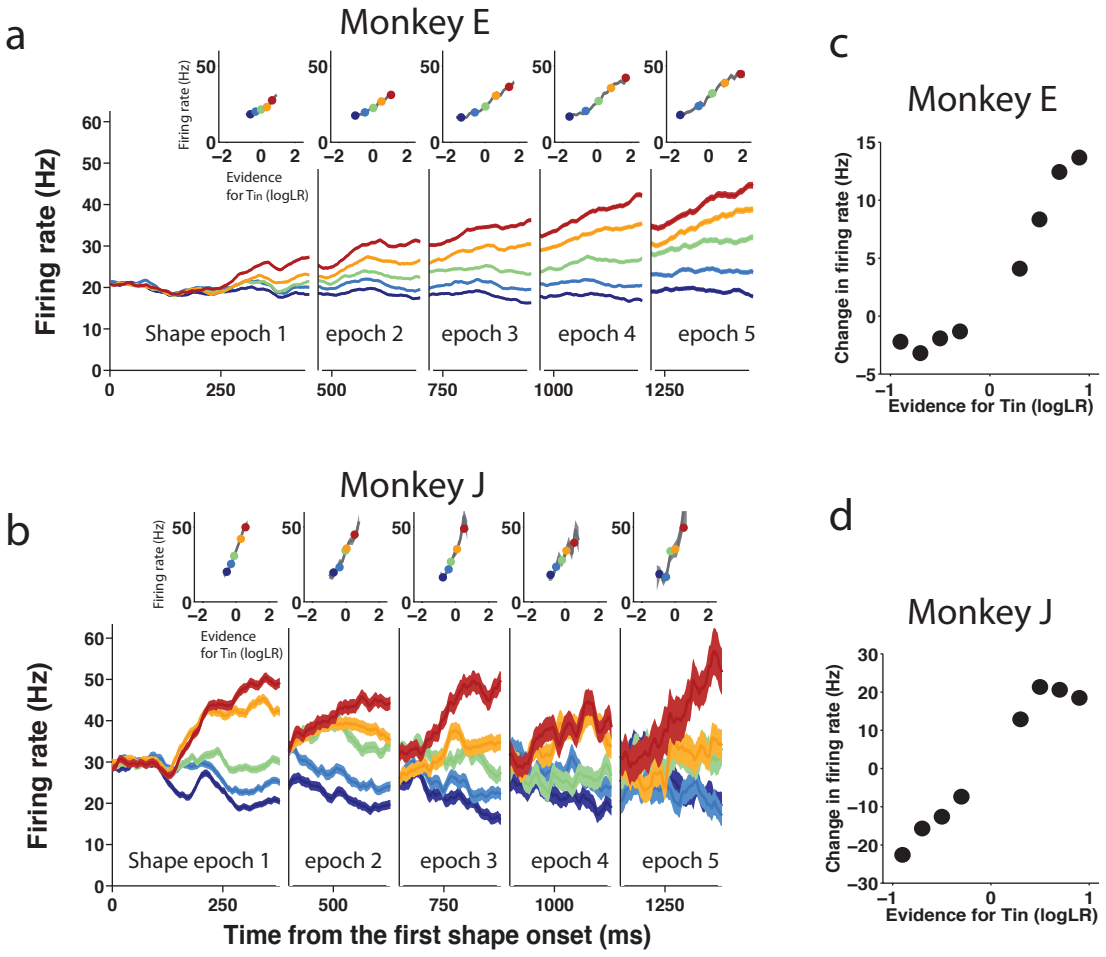
a
Monkey E
Firing rate (Hz)
Evidence for Tin (logLR)
Shape epoch 1
epoch 2
epoch 3
epoch 4
epoch 5
c
Monkey E
Change in firing rate (Hz)
Evidence for Tin (logLR)
b
Monkey J
Time from the first shape onset (ms)
d
Monkey J

Figure 2.4
LIP neural activity accompanying decision termination.
Average firing rates were obtained from data aligned to saccade initiation, grouped by choice (Tin or Tout). Within this grouping, the trials were sorted by quintile of the assigned cumulative logLR at the point in time τ_s+100ms prior to each plotted time point (i.e., 100 ms after the responses begin to diverge in the leftmost panel of Fig. 2.3a,b). Insets: Standard deviation of T_{in} responses. Responses are grouped into 10 quantiles and smoothed by running mean with 50ms bin. The standard deviation of quantile response curves is plotted at the earliest (left) edge of the bin. The arrow indicates the time when the standard deviation reaches its minimum.

Figure 2.4

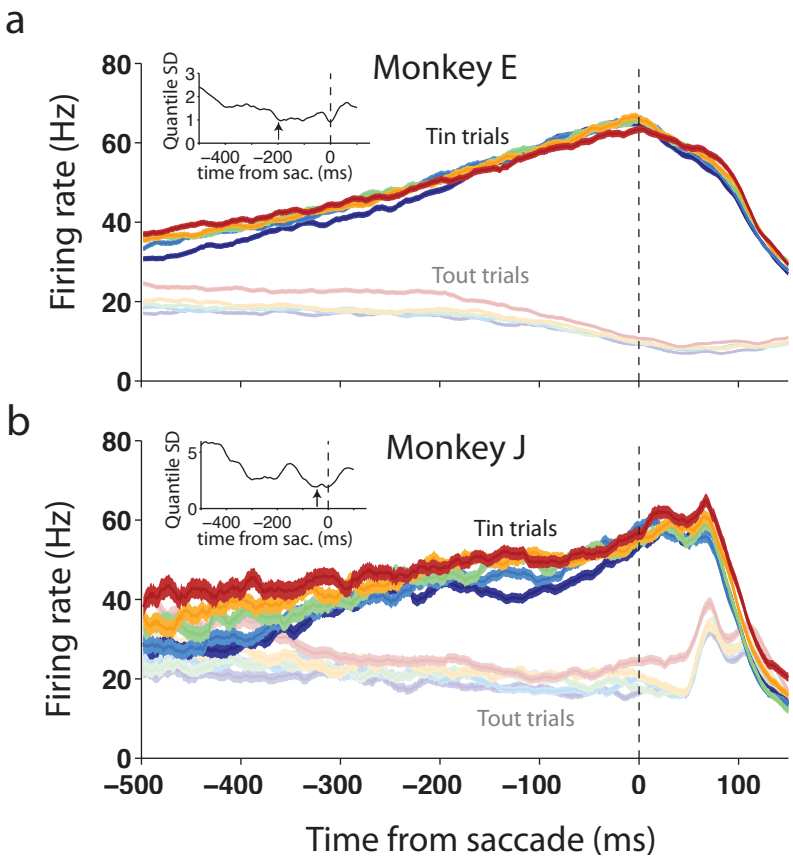
a
Monkey E
Quantile SD
3
2
1
0
−400 −200 0
time from sac. (ms)
Tin trials
Tout trials
Firing rate (Hz)
80
60
40
20
0
b
Monkey J
Quantile SD
5
0
−400 −200 0
time from sac. (ms)
Tin trials
Tout trials
Firing rate (Hz)
80
60
40
20
0
−500 −400 −300 −200 −100 0 100
Time from saccade (ms)

Figure 2.5

A bounded accumulation model informed by the neural measurements

Two competing neural populations accumulate a noisy representation of evidence bearing on the likelihood that their favored choice target, A or B, will be rewarded. The panels establish the logic of the model simulation in stages via an example trial in which a correct decision for A is made with RT=1.3 s, based on the evidence from N*=4 shapes. A 5th shape does not affect the decision.

a, At the beginning of each trial, a reward is assigned randomly to T_A or to T_B. The assignment determines the sampling distribution for the shapes in the trial. For example, if the reward were assigned to T_A the shapes would be sampled from the distribution in red, perhaps giving rise to the sequence of shapes: star, pacman, wedge, hourglass, pacman.

b, The presentation of each shape gives rise to a change in firing rate in each of the racing accumulators, termed population A and B. The expected ΔFR for population A is taken from the data (detrended; see Methods). The expectations are the opposite sign for the B population. These expectations are corrupted by unbiased noise, represented by the Guassian distributions. For example, the pacman presented as the 5th shape (orange shading) gives rise to a change change in firing rated in the A population modeled as a random draw from the orange Gaussian distribution on the left. The B population undergoes a change modeled as an independent draw from the orange Gaussian on the right. This values scale the dynamic templates (lower insets), which captures the sensory delay and rise time of the response.

c, The accumulated evidence is the sum of these scaled templates. The arrows point to the steady state response 100 ms after the sensory delay. The portion of the response that would be contributed by the 5th shape (orange) .

d, The time-dependent urgency signal is added to x in each neural population. This undoes the detrending step in (b).

e, The urgency signal pushes both accumulations upward from the value in panel c (reproduced here by the thin lines). When x in either population reaches the threshold (θ) (marked by a red star), the model chooses the corresponding target (T_A in this example). Notice that it is based only on the first 4 shapes. The RT occurs after a motor delay (τ_m), a random variable drawn from a Gamma distribution (blue).

Figure 2.5

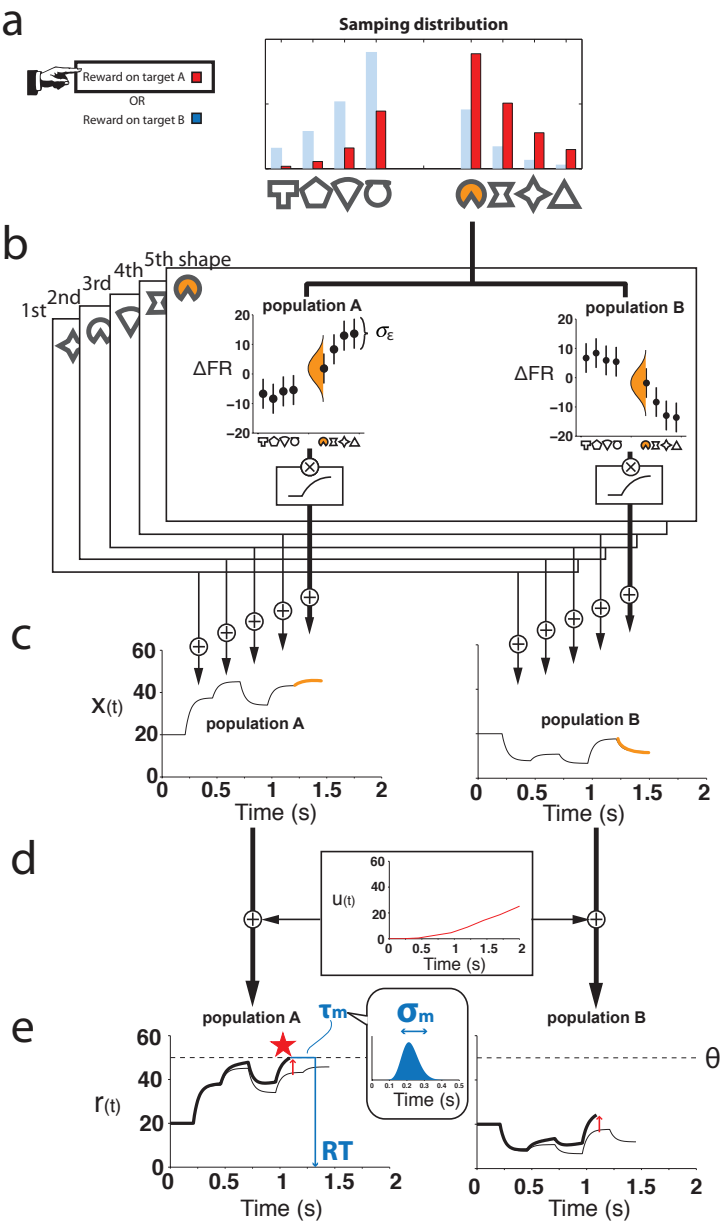
a
Samping distribution
Reward on target A
OR
Reward on target B
b
1st
2nd
3rd
4th
5th shape
population A
population B
ΔFR
σε
c
X(t)
population A
population B
Time (s)
d
U(t)
Time (s)
e
population A
Tm
σm
population B
r(t)
RT
θ
Time (s)

Figure 2.6
Evaluation of the physiology-inspired model (red) to data (black).
Monkeys E and J are shown in the left and right columns of graphs, respectively.
a, The reaction time distributions. Data (black) are the same as in Fig. 2.2a, but split by correct and incorrect trials.
b, Dependence of choice on the cumulative evidence supplied by the sequence of shapes. Data (black) are the same as in Fig. 2.2d.
c, Mean and standard deviation of the total logLR at the end of decision. Data (black) are the same as in Fig. 2.2e.
d, Time-dependent accuracy (TDA). Data (black) are the running proportions of correct choices for 100 trials sorted by RT. Values are plotted at the median RT. Red curves are model predictions. The decreasing TDA for monkey E is explained by the evidence-independent "urgency" signal that is added to the representation of accumulated evidence. The red TDA function is a prediction (not a fit) of the model framework.

Figure 2.6

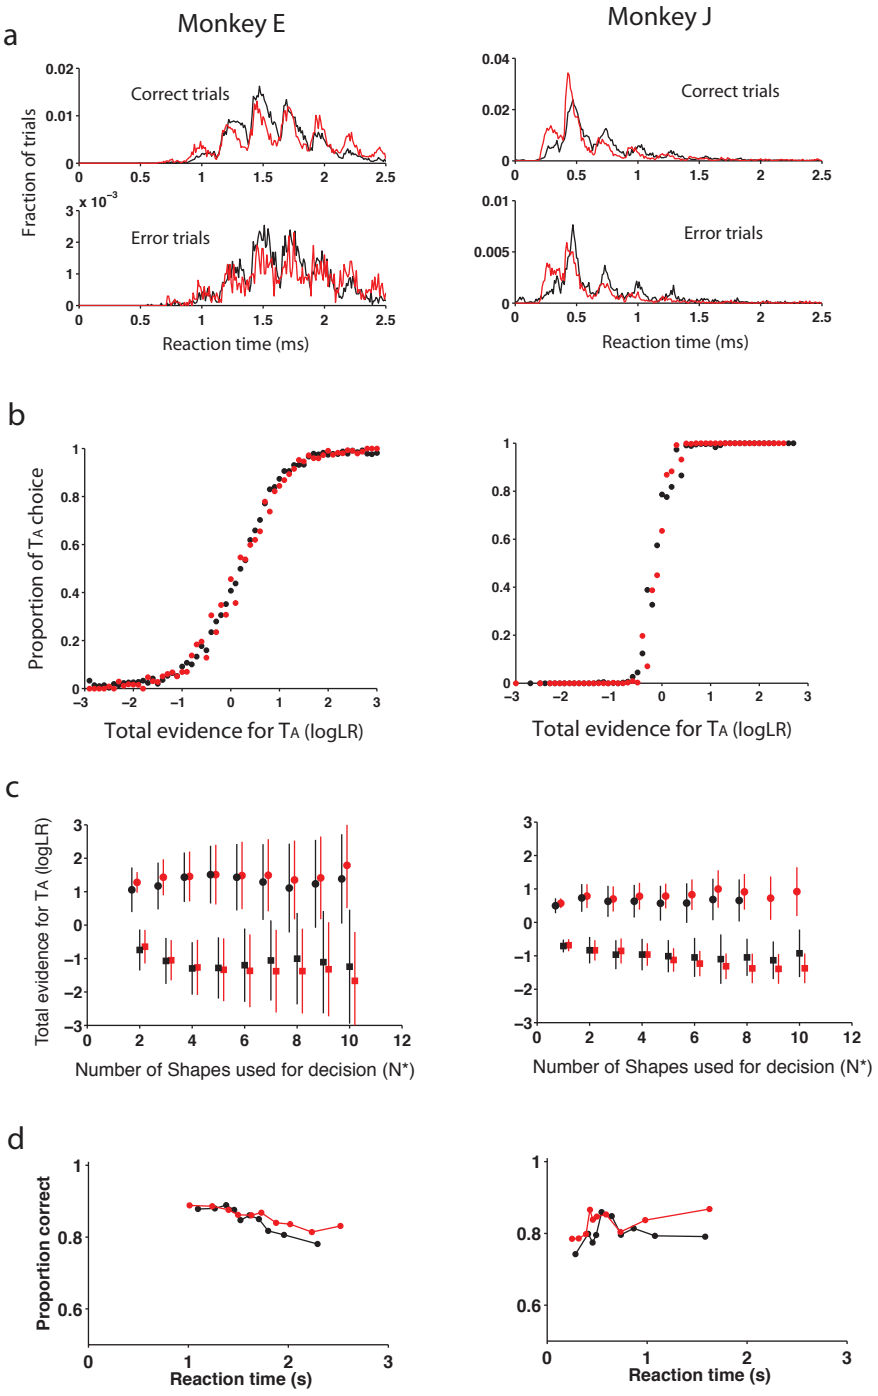
a
Monkey E
Monkey J
Correct trials
Error trials
Fraction of trials
Reaction time (ms)
b
Proportion of TA choice
Total evidence for TA (logLR)
c
Total evidence for TA (logLR)
Number of Shapes used for decision (N*)
d
Proportion correct
Reaction time (s)

Figure 2.S1.
Evaluation of an alternative model
Plots show observed mean total logLR at the end of decision conditioned on the reward assignment (Target A (circles) or Target B (squares)) and the number of shapes used for decision (N*). Dashed lines show the expected mean total logLR when RT is independent of evidence. Solid lines show constant levels of total logLR that best fit the observed data. Both Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) support that the total evidence remained at a constant level regardless of the number of shapes (N*) used for decision (solid line; H_0) rather than increased at a rate expected from the sampling distribution (dashed line; H_1) (ΔAIC= -1990 and -1348; ΔBIC=-1988 and -1345 for monkey E and J, respectively, where the negative value favors H_0). Error bars show s.e.m., but most of them are smaller than the data points.

Figure 2.S1

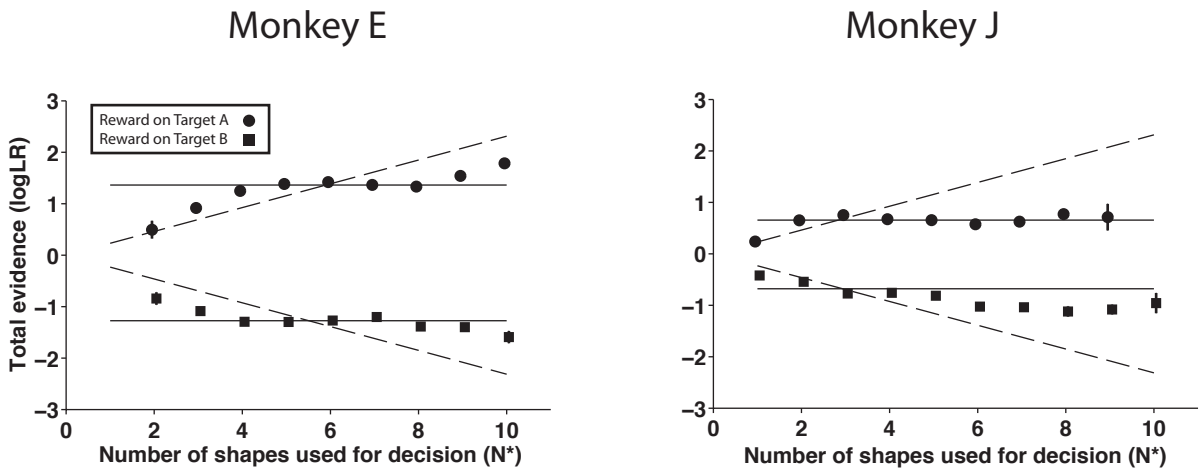
Monkey E
Monkey J
Reward on Target A
Reward on Target B
Total evidence (logLR)
Number of shapes used for decision (N*)
Number of shapes used for decision (N*)

Figure 2.S2

LIP neural activity at the end of decision aligned to the onset of last shapes

a, **b**, The response curves are sorted into quintiles as in Fig. 2.3a,b, but the response curves are aligned to the onset of the last effective shapes. The shaded blue rectangle indicates the timing of saccades (mean±S.D.). The Inset above each column shows the mean firing rate during a period indicated by gray bars as a function of total logLR. Color-filled circles represent the quintiles in each epoch, and the gray trace shows the running mean ± s.e. The regression slope value (± 95% confidence interval) is also shown.

c, **d**, The change of the firing rate induced by the last effective shape (the N* th shape). The regression slope value (± 95% confidence interval) is shown for changes induced by shapes that support T_{in}.

Figure 2.S2

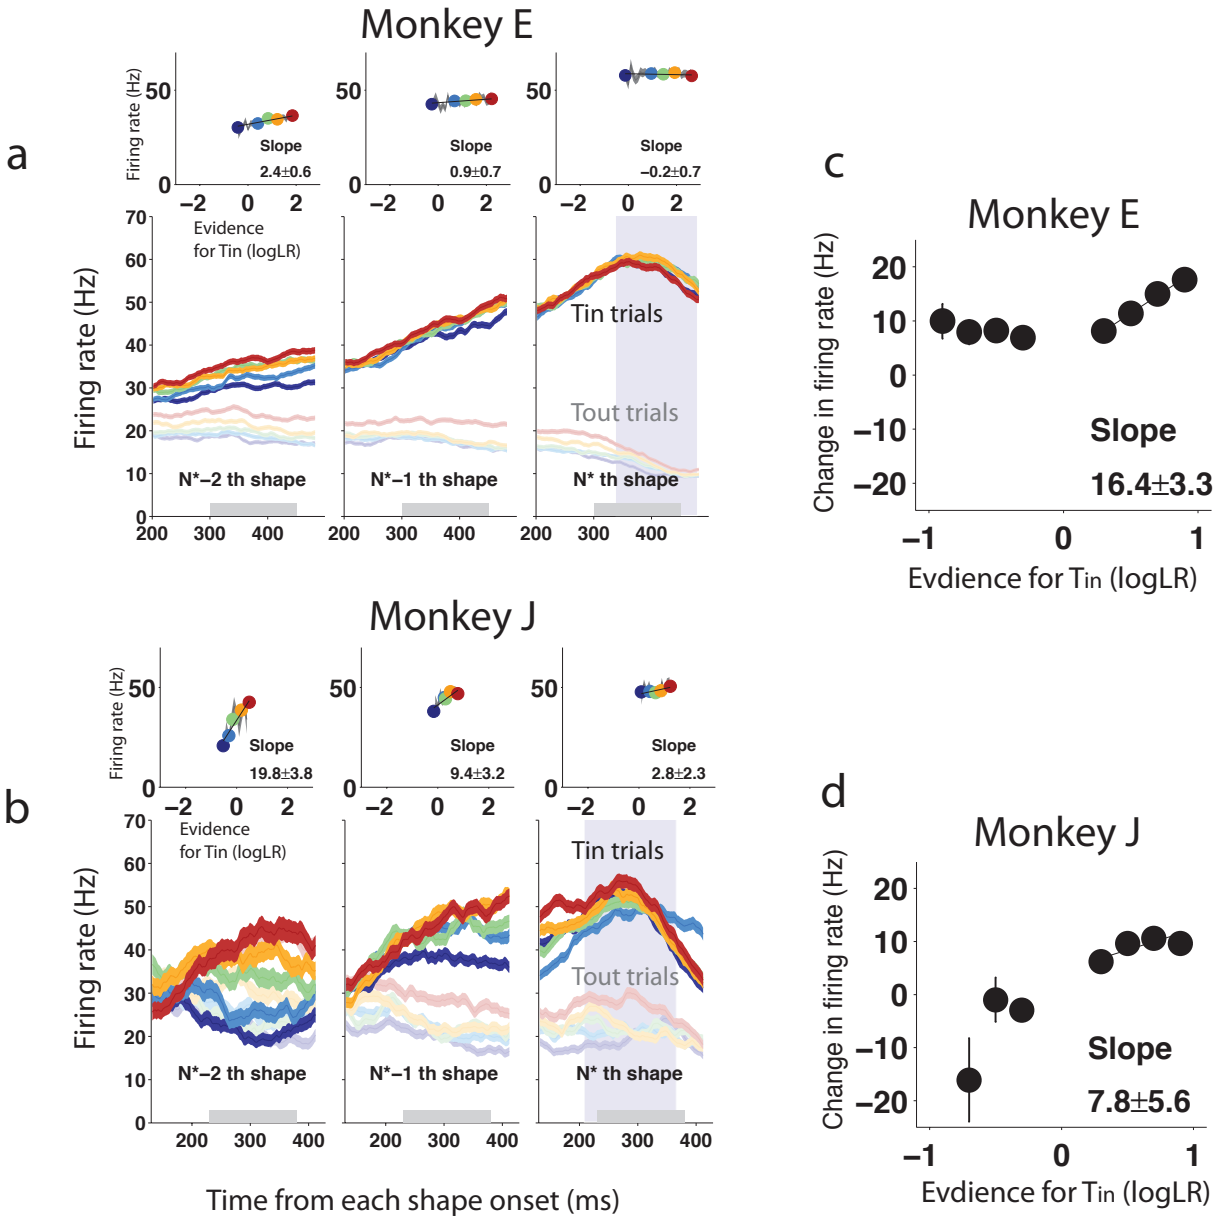
a
Monkey E
Firing rate (Hz)
Slope
2.4±0.6
Slope
0.9±0.7
Slope
−0.2±0.7
Evidence
for Tin (logLR)
Tin trials
Tout trials
N*−2 th shape
N*−1 th shape
N* th shape
b
Monkey J
Slope
19.8±3.8
Slope
9.4±3.2
Slope
2.8±2.3
Time from each shape onset (ms)
c
Monkey E
Change in firing rate (Hz)
Slope
16.4±3.3
Evdience for Tin (logLR)
d
Monkey J
Slope
7.8±5.6
Evdience for Tin (logLR)

Figure 2.S3
Shape-induced change in the LIP firing rate (ΔFR) in each epoch
The change induced by the k th shape is derived from the difference of the firing rate in the k-1 th and k th shape epoch. Each data point shows mean ± s.e. of the change in a similar fashion to Fig. 2.3c,d.

Figure 2.S3

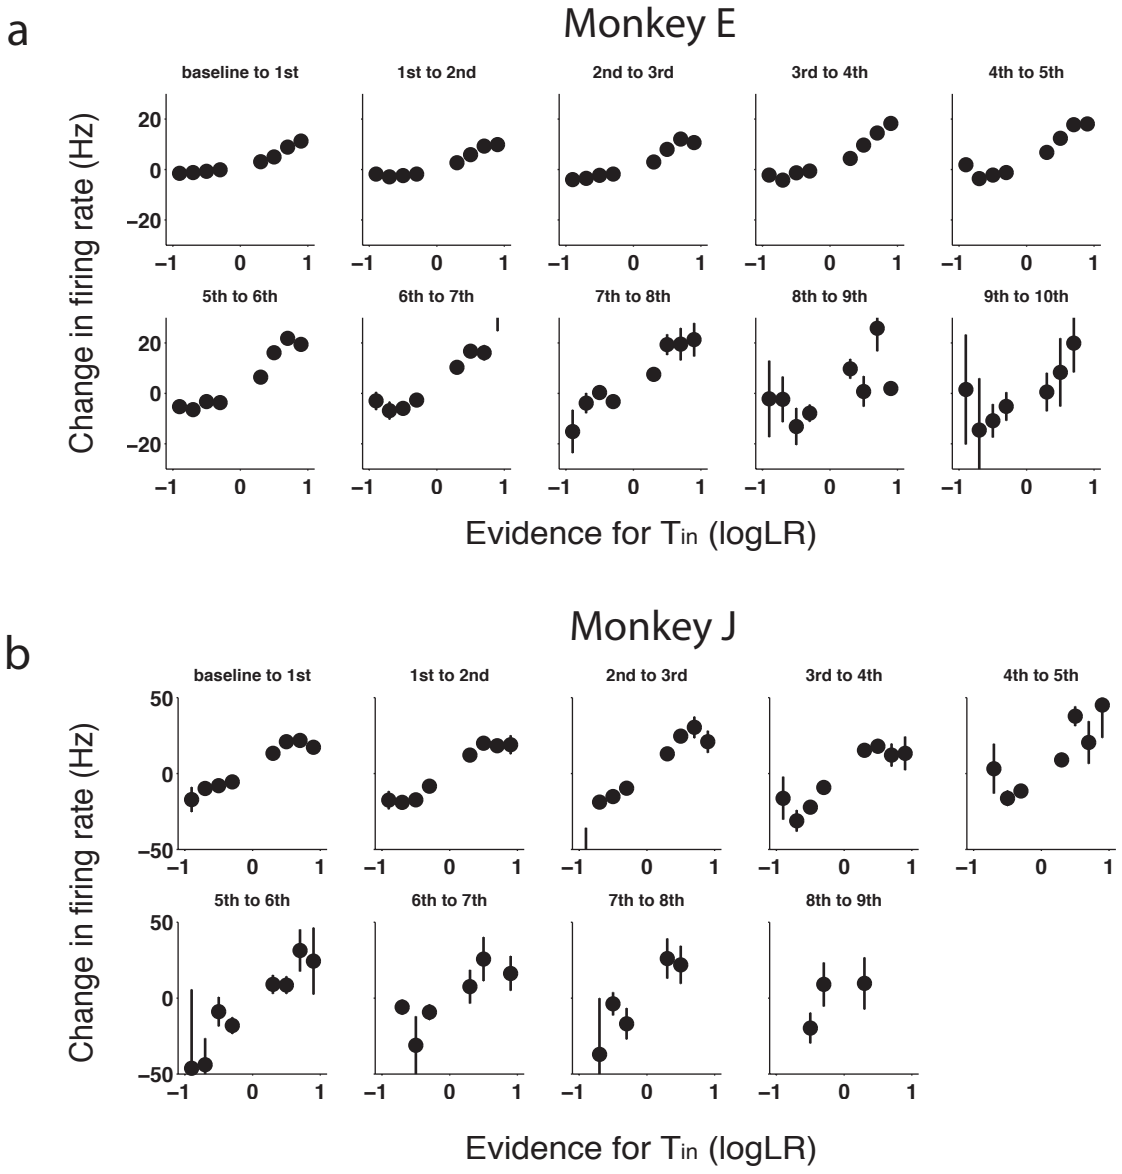
a
Monkey E
baseline to 1st
1st to 2nd
2nd to 3rd
3rd to 4th
4th to 5th
5th to 6th
6th to 7th
7th to 8th
8th to 9th
9th to 10th
Change in firing rate (Hz)
Evidence for Tin (logLR)
b
Monkey J
baseline to 1st
1st to 2nd
2nd to 3rd
3rd to 4th
4th to 5th
5th to 6th
6th to 7th
7th to 8th
8th to 9th
Change in firing rate (Hz)
Evidence for Tin (logLR)

Figure 2.S4.
Epoch dependence of LIP firing rate on cumulative evidence
Bounded evidence accumulation explains an apparent time/epoch dependence of firing rate on cumulative evidence. Each panel shows the mean firing rate (±s.e.m.) as a function of mean cumulative assigned logLR (black). For the first epoch, the 8 points correspond to the 8 shapes. For all subsequent panels, the points represent quantiles of the cumulative evidence (displayed at each quantile median along the abscissa). The model (red) uses the same 8 values of ΔFR for all epochs. yet it shows the same apparent decrease in slope as the data. This is because the monkeys tend to use more shapes in trials where the evidence changes the firing rate by a smaller magnitude than average. The plots exclude the firing rates induced after the last effective shapes (the N*th shapes).

Figure 2.S4

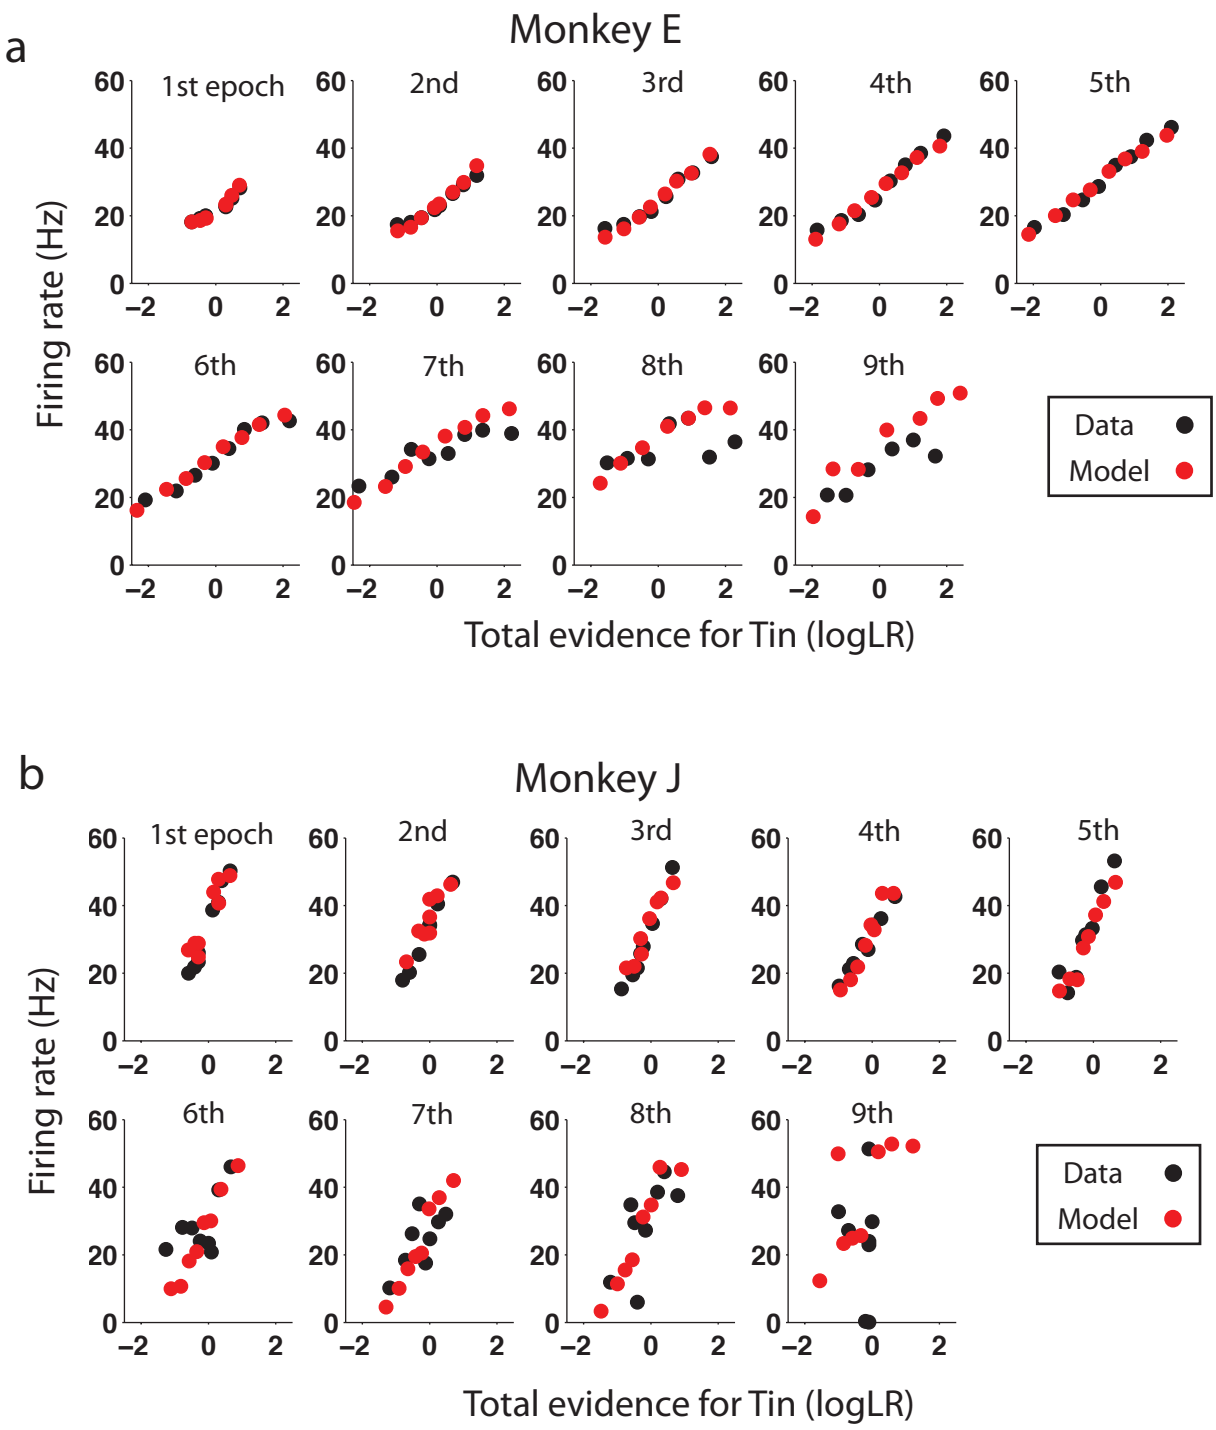
a
Monkey E
1st epoch
2nd
3rd
4th
5th
6th
7th
8th
9th
Firing rate (Hz)
Total evidence for Tin (logLR)
Data
Model
b
Monkey J
1st epoch
2nd
3rd
4th
5th
6th
7th
8th
9th
Firing rate (Hz)
Total evidence for Tin (logLR)
Data
Model

Figure 2.S5.
LIP firing rate for monkey J when the green target was placed in the RF
Responses are grouped into quintiles as in Fig. 2.3.
a, The response curves aligned to the onset of the first shape.
b, The change in the firing rate induced by each unique shape.
c, The response curves are divided into T_{in} and T_{out} trials and aligned to the onset of the last three effective shapes.

Figure 2.S5

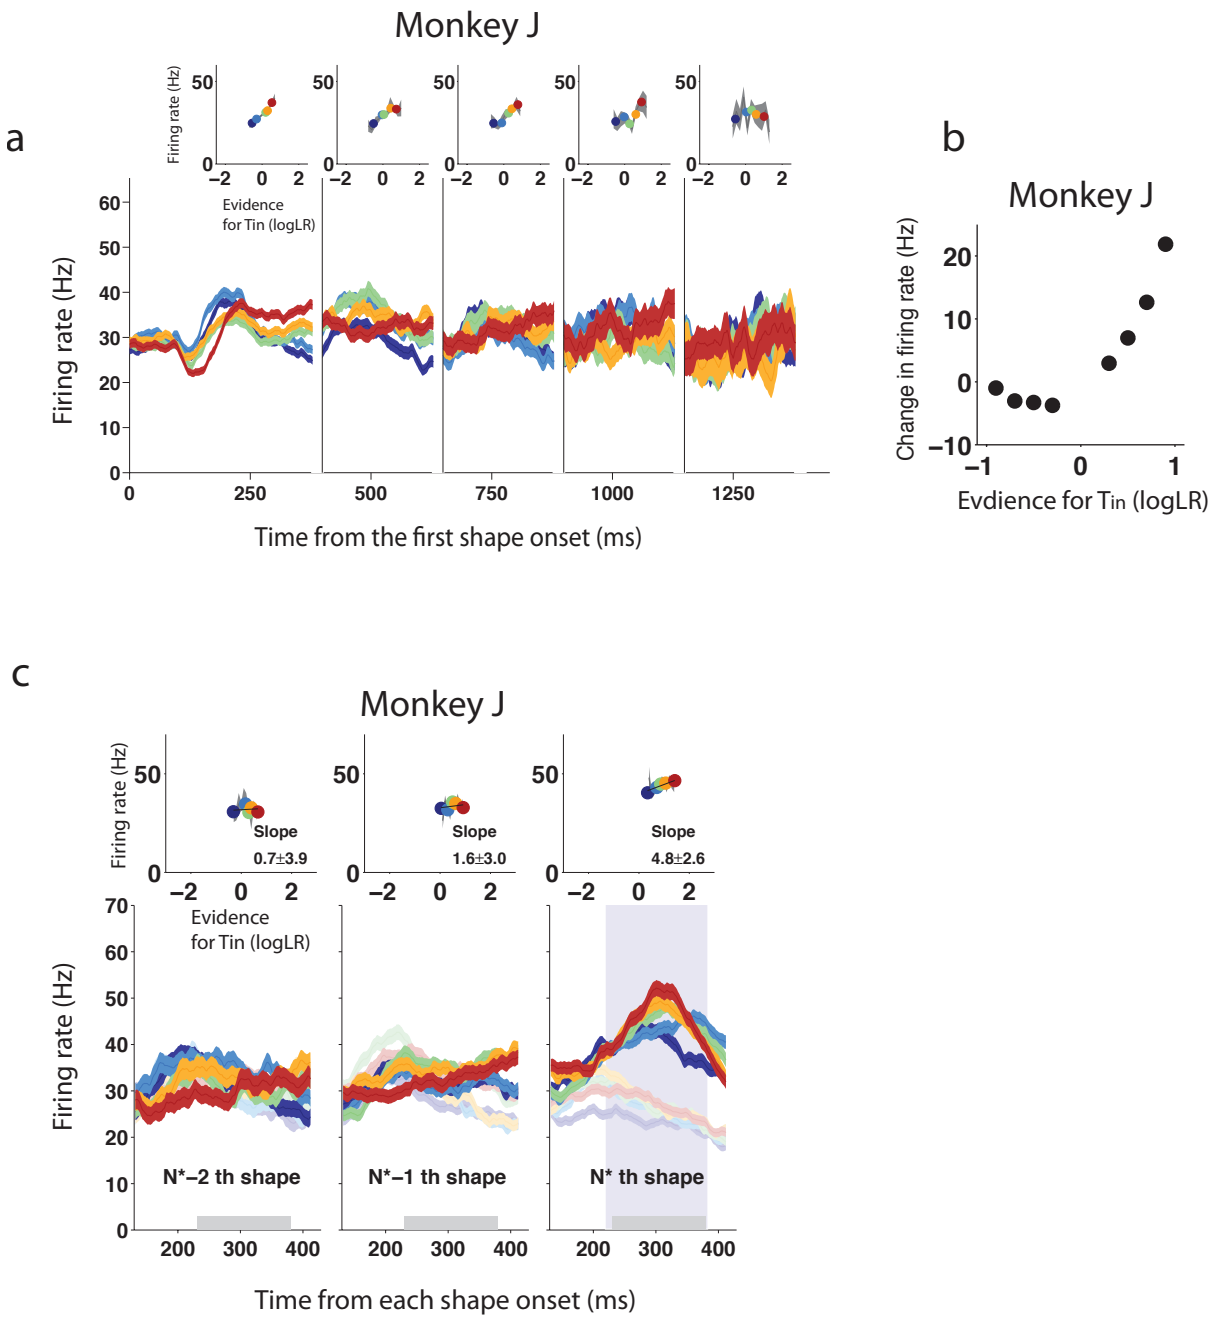
a
Monkey J
Firing rate (Hz)
Evidence for Tin (logLR)
Time from the first shape onset (ms)
b
Monkey J
Change in firing rate (Hz)
Evdience for Tin (logLR)
c
Monkey J
Slope 0.7±3.9
Slope 1.6±3.0
Slope 4.8±2.6
N*–2 th shape
N*–1 th shape
N* th shape
Time from each shape onset (ms)

Table 2.1

	σ_ε (Hz)	σ_m (ms)
monkey E	6.0	34
monkey J	3.7	69

Table 2.1. Fitted parameters for each monkey

Chapter 3: The effect of time pressure can be explained by a dynamic change in the evidence required for a decision

Abstract

In many decision models, evidence is accumulated until it reaches a bound. Changing the bound height results in a trade-off between decision time and accuracy. Subjects can improve their accuracy by raising their bound to accumulate more evidence before committing to a choice, but this takes more time. I investigated how effects of time pressure can be accounted for by such models of evidence accumulation. To vary time pressure, I manipulated the cost of time during a simple perceptual task: the discrimination of the direction of motion. On each trial, I provided feedback of performance in terms of points and instructed the subjects to maximize their points per unit of time. Unbeknownst to the subjects, halfway through the experiment I increased the time pressure by canceling a small fraction of trials if they had not made a decision by a randomly-chosen deadline. This manipulation induced the subjects to make faster decisions while lowering their decision accuracy. To fit their behavior with the evidence accumulation model, I had to allow the decision bound to decline as time elapses during a trial. Increased time pressure was accounted for by a more rapid decline in the decision bound. In summary, the effects of time pressure can be accounted for by changes in the decision bound over time rather than other parameters of the evidence accumulation model.

Introduction

Acquiring an additional sample of evidence can improve the accuracy of many decisions, albeit at the cost of time. One idea is that decisions end when the accumulated evidence reaches a threshold level, or bound (Ratcliff, 1978b; Gold and Shadlen, 2002; Shadlen and Kiani, 2013). Humans and monkeys can adjust the level of such a bound and trade off speed and accuracy of decisions (Palmer et al., 2005).

When time is limited by a deadline, as is often the case in real life, time pressure urges a decision-maker to finish making a decision before the deadline. However, it is unknown how a decision maker incorporates time pressure that dynamically changes over time. Theoretical studies suggested that if the accumulation of evidence over time incurs a time cost, the time-dependent decline of decision bounds achieves an optimal performance in the sense that it maximizes the net reward rate (i.e. the sum of expected rewards and costs per unit of time) (Franzier and Yu, 2008; Drugowitsch et al., 2012). Yet, experimental evidence is lacking as to whether time pressure can induce subjects to adjust their bounds in a time-dependent manner.

To test this, I manipulated the cost of elapsed time during a simple perceptual decision-making task. Human subjects discriminated the direction of random dot motion in a choice-reaction time paradigm. After each trial, I provided subjects with feedback that showed a score based on their decision results (i.e. correct, wrong, or aborted) and instructed them to maximize their score per unit of time. When the subjects reached a stable performance, I increased their

time pressure by introducing deadlines (in half of trials) that cancelled a small fraction of trials. Compared to their baseline performance, the increased time cost induced the subjects to make faster decisions, and their decision accuracy decreased more quickly as time elapsed during a trial. These behavioral changes were well explained by a time-dependent decline of the bounds in a bounded accumulation model.

Methods

Subjects

Subjects were four young adults (3 males and 1 female) with normal or corrected-to-normal visual acuity. They either volunteered to participate in the experiment or were paid $15 per hour. They provided informed written consent before participation in the study. All procedures were approved by the institutional review board at the University of Washington.

Apparatus

Subjects were seated in front of a video display. A chin rest and forehead bar were used to stabilize the head of the subject for the duration of each experimental session. The stimuli were displayed on a flat-screen CRT video monitor (19-in View Sonic PF790, 800 by 600 pixels, viewing distance = 57 cm, subtending 34° by 25.5° with 22.8 pixel/deg at screen center; refresh rate = 75 Hz) controlled by a Macintosh G5 (2x2.8GHz, Mac OS 10.5 with an ATI Radeon HD

2600 graphics card). Stimuli were generated using the Psychophysics Toolbox Version 2 (Brainard, 1997; Pelli, 1997) for MATLAB (Version 7.5, Mathworks). Eye movements in one eye were recorded with 1000 Hz sampling rate by a noninvasive infrared video-tracking system (EyeLink 1000, SR Research).

Motion direction discrimination task

Subjects discriminated the direction of dynamic random dot motion in a choice-reaction time paradigm (fig. 1a). Trials began with the appearance of a fixation spot (0.3° diameter) in the center of the screen. Subjects were required to look at the fixation spot and maintain their fixation within 1°x1° of their initial fixation. Once fixation was attained, two red targets (0.5° diameter) were presented 6° to the left and right of the fixation spot. After a variable delay (200–500ms, exponentially distributed; mean 300ms), the random dot motion stimulus appeared within a 5° circular aperture centered on the fixation spot. Dots were 2 by 2 pixels (0.1° square), with a density of 16.7 dots/deg^2/s. On each trial, the direction of motion was randomly assigned as right or left, and the motion strength was randomly selected from a set of six coherence levels (0%, 3.2%, 6.4%, 12.8%, 25.6%, and 51.2%). The coherence specifies the fraction of dots that are displaced in motion (5°/s) as opposed to randomly repositioned. Details on the generation of the random dot motion have been described in previous studies (Roitman and Shadlen, 2002; Palmer et al., 2005).

After the motion stimulus onset, subjects were required to decide the direction of the motion stimulus and make a saccadic eye movement to the corresponding choice targets. When a saccade was detected within 6°x6° around the choice target, the trial was classified as either correct or error based on whether a chosen target direction and an assigned motion direction was the same or opposite, respectively (the designation of correct/error was random for the 0% coherence trials). In these trials, time from the stimulus onset to the initiation of a saccade was registered as a reaction time. When subjects broke a fixation prematurely (e.g. eye blinks) or made a invalid saccade, the trial was classified as "*aborted*" and terminated immediately.

After each trial, 2s of an inter-trial interval followed. Subjects performed trials repeatedly during a 13min session (typically ~250 trials), and completed two sessions a day with a short break (5–10min) between them.

Scoring system

While subjects performed the task, their performance was scored as follows. The subjects earned one point for a correct choice and lost one point for an error choice. Their total point did not change for aborted trials. During the 2s of an inter-trial interval, subjects received visual and auditory feedback about their performance. The visual feedback indicated the total points and the earned points per minute (point rate). The feedback presented a histogram that summarized the current point rate during the latest 2 min, and a history of the point rate for every 2 min from the

beginning of a session (Fig. 3.1b). The subjects were instructed to maximize their point rate during a session.

Experimental schedule

I divided the experiment into two phases (Fig. 3.1b). In the first phase, subjects were allowed to make a decision whenever they were ready within 5s from the motion stimulus onset. These trials are referred to as *standard trials*. Because subjects never viewed the stimulus for the whole duration, subjects performed the task as if there was no time limit. While the subjects performed the task, I monitored their decision speed and accuracy by their average reaction time for 0% coherence trials (for subject A and B only) and average point rate (for all the subjects). The subjects performed the task repeatedly until their speed and accuracy were deemed stable for 10 consecutive sessions. These 10 sessions were categorized as the *normal phase*, and all the trials during this phase were the standard trials.

In the next *time-pressure* phase (20 sessions), deadlines were imposed in a half of trials, which are referred to as *deadline trials*. The deadline trials and standard trials were randomly interleaved during this phase. If subject made a decision before a deadline, the trial ended in the same way as in the standard trial. If subjects had not made a decision by this deadline, however, the trial was classified as "*cancelled*" and terminated immediately as if they had broken fixation. For each deadline trial, a deadline was randomly drawn from a shifted Rayleigh distribution:

$$P_{DL}(t) = \frac{t - t_0}{\sigma_{DL}^2} \exp\left(-\frac{(t - t_0)^2}{2\sigma_{DL}^2}\right)$$

where t_0 is the time shift and sets the minimum time for the deadline, and σ_{DL} determines the variance of the distribution. These parameters were tailored for each subject based on his or her reaction time distributions measured in the normal phase (fig. 1c). Specifically, the parameters were chosen to cancel ~10% of the trials in the time-pressure phase if the subjects had performed with the same speed as in the normal phase. Table 3.1 summarizes t_0 and σ_{DL} used for each subject. Note that the even if the subjects performed the task with the same speed, their mean reaction time for the deadline trials is shorter than that for the standard trials. This is because their reaction time has to be shorter than a deadline in each trial. To rule out a possibility that this attrition accounts for the shorter reaction time in the time-pressure phase than in the normal phase (Fig. 3.2a), only the standard trials were used for the analyses, and deadline trials were excluded.

Subject performed additional 10 sessions without deadlines. I did not include results in these sessions, however, because subjects showed inconsistent behavioral change from the time-pressure phase. I did not explain the experimental schedule to the subjects, so that the prior knowledge of the task would not bias their behavior.

Diffusion model

I analyzed subjects' performance in the framework of a bounded drift-diffusion model (Fig. 3.4a,b). In this model, a positive and negative sign are assigned to a rightward motion coherence (+C) and leftward motion coherence (-C). The decision variable (X) evolves over time

by following a Wiener process with a constant drift:

$$dX = \mu_d\, dt + dW$$

The drift term μ_d is proportional to the motion strength: $\mu_d = kC$, where C is the motion coherence and k is a free parameter that determines a signal-to-noise ratio in my model. The stochastic term W follows the Wiener process: $dW \sim N(0, dt)$. The right-hand side of the equation represents momentary noisy evidence. When the motion coherence is rightward (C>0), it is likely to furnish positive momentary evidence. The model works by accumulating momentary evidence into a decision variable until this variable reaches one of two decision criteria—that is, either to a right-choice bound or a left-choice bound set at B and -B, respectively. The bound reached first by the decision variable determines the choice, and the decision time is determined by how long the decision variable took to reach that bound. The total reaction time is a sum of the decision time and a non-decision time (t_{nd}), which accounts for sensory and motor delays. I assumed that the non-decision time has a normal distribution over time:

$$p_{nd}(t) = \frac{1}{\sqrt{2\pi\sigma_{nd}^2}} \exp\left\{\frac{(t-\mu_{nd})}{2\sigma_{nd}^2}\right\}$$

where its mean (μ_{nd}), and variance (σ^2_{nd}) do not depend on a choice direction or motion coherence.

I compared two types of models: one is a conventional model where the bound height remains stationary during a trial (stationary bound model). The other model dynamically changes

the bound height as a function of elapsed time during a trial (collapsing bound model). I estimated a shape of the collapsing bounds by three types of functions (i.e. linear, quadratic, and exponential functions), among which the exponential function best fitted behavioral data in most of the subjects. Thus, I used the following form of exponential function in my analyses:

$$u(t) = B\exp(-at)$$

where a determines the rate of collapse. Whereas the stationary bound model has four free parameters: $\theta = [k, B, \mu_{nd}, \sigma 2_{nd}]$, the collapsing bound model has one additional parameter (a), for a total of five free parameters: $\theta = [k, B, \mu_{nd}, \sigma^2_{nd}, a]$.

Model fitting

I fitted the two types of diffusion models to the data as follows. First I numerically solved a Fokker-Plank equation to calculate the distributions of decision time (p_d) for right and left choices at each coherence level given a set of parameters θ: $p_d(t, right \mid \theta, coh)$ and $p_d(t, left \mid \theta, coh)$. Because a reaction time is a sum of a decision time and a non-decision time under my assumption, I computed the probability of observing an actual reaction time in each trial by a convolution:

$$p_{RT}(t, choice \mid \theta, coh) = \int_0^{RT} p_d(\tau, choice \mid \theta, coh)\, p_{nd}(RT - \tau \mid \theta) d\tau$$

I computed this probability for all trials, and fitted the free parameters (θ) by the maximum log likelihood method. Standard error of each parameter was obtained by fitting the model to

bootstrapped data 100 times.

Point rate calculation

The expected point rate in each phase can be calculated as follows:

$$\left\langle \text{Point Rate}(\theta) \right\rangle = \frac{\left\langle P(correct \mid \theta) \right\rangle - \left\langle P(error \mid \theta) \right\rangle}{\left\langle RT(\theta) \right\rangle + T_{ITI}}$$

where RT is the reaction time, and T_{ITI} is the inter-trial interval. In the normal phase, where all the trials are the standard trials, the expected reaction time is,

$$\left\langle RT_{ST}(\theta) \right\rangle = \int_0^\infty t \cdot p_{RT}(t \mid \theta) dt$$

The proportion of correct/error for the standard trials can be calculated as follows:

$$\left\langle P_{ST}(correct \mid \theta) \right\rangle = \int_0^\infty p_{RT}(t, correct \mid \theta) dt$$

$$\left\langle P_{ST}(error \mid \theta) \right\rangle = \int_0^\infty p_{RT}(t, error \mid \theta) dt$$

In the time-pressure phase, the expected reaction time and the proportion of correct/ error must be calculated separately for the standard trials and deadline trials. For the standard trials, I already calculated them above. For the deadline trials, I calculate them for trials where the subject made a choice before the deadline (valid trials), and for trials that were cancelled (invalid trials). The time distribution of valid and invalid trials can be calculated as follows:

$$p(t, valid \mid \theta) = p_{RT}(t \mid \theta) \cdot \left(1 - \int_0^t p_{DL}(\tau) d\tau\right)$$

$$p(t, invalid \mid \theta) = \left(1 - \int_0^t p_{RT}(\tau \mid \theta) d\tau\right) \cdot p_{DL}(t)$$

where $p(t, valid \mid \theta) + p(t, invalid \mid \theta) = 1$.

I can compute the expected "reaction time", which includes cancellation times, as follows:

$$\left\langle RT_{DL}(\theta) \right\rangle = \int_0^\infty t \cdot \left(p(t, valid \mid \theta) + p(t, invalid \mid \theta) \right) dt$$

The proportion of correct and error trials for the deadline trials can be calculated as:

$$\left\langle P_{DL}(correct \mid \theta) \right\rangle = \int_0^\infty p_{RT}(t, correct, valid \mid \theta) dt$$

$$\left\langle P_{DL}(error \mid \theta) \right\rangle = \int_0^\infty p_{RT}(t, error, valid \mid \theta) dt$$

Note that $\left\langle P_{DL}(correct \mid \theta) \right\rangle + \left\langle P_{DL}(error \mid \theta) \right\rangle < 1$ because some fraction of trials are cancelled and counted as neither correct nor error.

Finally, by taking the average of the standard trials and deadline trials, I can compute the mean reaction time and the proportion of correct/error trials in the time-pressure phase:

$$\left\langle RT_{TP} \right\rangle = \rho \cdot \left\langle RT_{\mathrm{ST}}(\theta) \right\rangle + (1 - \rho) \cdot \left\langle RT_{DL}(\theta) \right\rangle$$

$$\left\langle P_{TP}(correct \mid \theta) \right\rangle = \rho \cdot \left\langle P_{\mathrm{ST}}(correct \mid \theta) \right\rangle + (1 - \rho) \cdot \left\langle P_{DL}(correct \mid \theta) \right\rangle$$

$$\left\langle P_{TP}(error \mid \theta)\right\rangle = \rho \cdot \left\langle P_{\mathrm{ST}}(error \mid \theta)\right\rangle + (1-\rho) \cdot \left\langle P_{DL}(error \mid \theta)\right\rangle$$

where ρ= 0.5 is a proportion of standard trials. The rest of the calculations for the point rate follow those for the normal phase.

Maximum-rate bound

I estimated the bound height that theoretically maximizes the point rate for this model, which I define as the *optimal bound*. To find the optimal bound, I first fitted the collapsing bound model to each subject in each phase, and then adjusted the bound while keeping other parameters (k, μ_{nd}, and σ^2_{nd}) fixed., To estimate the shape of the optimal bound more accurately, I added one more free parameter (b) which allowed the bounds to collapse after a time offset (b) as follows:

$u(t) = B$ if t < b

$u(t) = B\exp(-a(t-b))$ otherwise

Results

Four human subjects performed a reaction-time version of the random dot motion discrimination task (Fig. 3.1a). Their performance was scored with points (see Methods), and they were instructed to maximize their point earning rate. Figure 3.2 shows how the mean reaction time and decision accuracy varied with the motion strength in both the normal (black) and time-pressure phase (red). The mean reaction time decreased from its highest value at zero motion strength to its lowest value at the highest motion strength (Fig. 3.2a). The mean reaction time was longer for error trials than for correct trials (Fig. 3.2a). The accuracy increased from a chance level at zero motion strength to a nearly perfect level at the highest motion strength (Fig. 3.2b).

Despite these similarities, the introduction of deadlines induced the subjects to make faster decisions in the time-pressure phase than in the normal phase (Fig. 3.2a, $p<0.001$, t-test for average reaction time across all coherence trials). Even though the timing of deadlines was designed to cancel ~10% of trials, subjects quickly increased their decision speed within the first few sessions so that less than 5% of trials were cancelled (Table 3.1). The reduction in the mean reaction time ranged from 33% at zero motion strength (from 1314±17ms to 874±9ms) to 6% at the highest motion strength (from 434±2ms to 406±1ms), indicating that the reduction was more pronounced for lower coherence trials than for higher coherence trials.

The decision accuracy in the psychometric curve did not significantly differ across the two phases (Fig. 3.2b, $p>0.4$, likelihood ratio test for Weibull fit). However, a closer inspection revealed that the time pressure significantly affected the subjects' decision accuracy in a time dependent manner (Fig. 3.3). I show this using time-dependent accuracy functions that plot accuracy conditional on binned reaction time. For each level of weaker motion strengths trials were sorted into trial bins and the decision accuracy in each bin was plotted against its mean reaction time. In both phases, the accuracy declined as a function of reaction time for a given motion strength ($p<10^{-5}$; linear regression). However, the magnitude of the decline was greater in the time-pressure phase (red) than in the normal phase (black). These results indicate that the subjects made less accurate decisions as time elapsed during a trial in both phases, and this trend was more prominent under time pressure.

Bounded Drift-Diffusion Model

To gain insights into the underlying decision-making mechanism, I analyze the behavioral data using the bounded drift-diffusion model (Fig. 3.4, see Methods). This model explains a wide range of psychophysical experiments in humans and monkeys (Smith and Ratcliff, 2004; Palmer et al., 2005; Gold and Shadlen, 2007). In this model, noisy sensory evidence is accumulated over time into a decision variable that undergoes diffusion with an overall drift that is proportional to the motion coherence. When the decision variable reaches one of the two decision bounds,which correspond to the two alternative choices, the model makes a decision.

In a bounded drift-diffusion model, the bound height determines the decision accuracy. Because the conventional model has a constant bound across time during a trial, this model predicts that the decision accuracy is also constant across time for a given motion strength. Contrary to this prediction, the decision accuracy of my subjects declined with time (Fig. 3.3), suggesting that the subjects lowered their decision criteria through the time course of a trial. Thus, I generalized the model such that the two bounds dynamically change their height and collapse together in a time-dependent manner, which I call a collapsing bound model.

Collapsing bounds better explained several aspects of subjects' data than stationary bounds. First, a collapsing bound model reproduced the decline in the time-dependent accuracy (Fig. 3.3 solid curves) whereas the stationary bound model predicted the same level of accuracy regardless of reaction time (data not shown). Second, the collapsing bound model also predicted that the mean error reaction time was longer than mean correct reaction time at each motion strength (Fig. 3.2a solid curves), whereas the stationary bonds predicted the same mean reaction time for both correct and error trials. Third, compared to the stationary bound model, the collapsing bound model better predicted the shape of reaction-time distributions by reducing a fraction of long reaction-time trials at each motion strength ($p<0.001$, likelihood ratio test). Altogether, the better explanatory power of the collapsing bound model affirms the plausibility of the collapsing bound model.

One important property of a bounded drift-diffusion model is that a change in the bound height can produce a tradeoff in decision speed and accuracy. Lowering the bound height increases the decision speed but decrease the accuracy, which are consistent with the subjects' behavioral change from the normal to time-pressure phase (Fig. 3.3). Moreover, the time-dependent accuracy decreased more rapidly in the time-pressure phase than in the normal phase (Fig. 3.3). Therefore, I hypothesized that time pressure induced the subjects to decline their bound height more rapidly in the time-pressure phase.

To test this hypothesis, I examined how bound height changed across the two phases by fitting the collapsing bound model to each subject (Fig. 3.5a, solid lines, see Methods). For each subject, when the deadline was introduced in the time-pressure phase, the initial bound height became lower and the rate of collapse became larger than the normal phase (Fig. 3.5b), indicating that the bound height was lower in the time-pressure phase than in the normal phase throughout the trial. This supports my hypothesis that time pressure induced a larger magnitude of bound collapse, explaining the faster but less accurate decisions in the time-pressure phase.

These behavioral changes may also be explained by a change in other parameters in the model. However, an increase in a scaling factor (k) increases both decision speed and accuracy whereas a decrease in non-decision time (μ_{nd}) increases decision speed but does not change the accuracy. Thus, a change in these parameters cannot explain both the increased speed and decreased accuracy simultaneously.

Optimal strategy

Finally, I asked whether collapsing bounds is a rational strategy in terms of maximizing the point rate. To test this, I evaluated how closely the subjects set their bound height to the one that maximizes the point rate, which I call an "optimal bound". I first fitted a collapsing bound model to each subject in each phase, and then searched for an optimal bound by changing the bound parameters to change freely while keeping the rest of parameters fixed. Figure 3.5a compares the subjects' bounds (solid curves) with the optimal bounds (dashed curves). The optimal bounds showed a time-dependent decline both in the normal and time-pressure phase. Furthermore, the optimal bound was lower in the time-pressure phase than in the normal phase. The subjects' bounds also showed these trends and their expected point rate was >97% of the maximum point rate, suggesting that the subjects changed their bound height in a rational way.

Discussion

In this study, response deadlines increased time pressure and induced the human subjects to make faster but less accurate decisions. To understand how time pressure affected their decision-making process, I analyzed their performance with a bounded drift-diffusion model. The model indicated that the subjects' decision bound declined as time elapsed during a trial (Fig. 3.5a). This change of the bound was more pronounced when they were under the time pressure. In short, time pressure lowers the decision criterion in a time-dependent manner.

I estimated the time-dependent bound by a simple exponential function, which made the model fitting tractable and the fitted model easier to interpret. I also used other families of functions (e.g. linear, quadratic, and logistic functions) to estimate the bound change and found qualitatively similar results, reassuring that my results do not depend on a specific choice of the function for the estimation. Moreover, a time-dependent decline of the decision accuracy provides a model-free support for my conclusion (Fig. 3.3).

The scoring system in my study allowed me to objectively assess advantages of collapsing bounds. Collapsing bounds helped the subjects to achieve a higher point rate in several ways. One obvious effect was to reduce a fraction of trial cancellations, which decreases subjects' point rate. Because the cancellation was more likely to happen as time elapsed during a trial, collapsing bounds efficiently helped the subjects make decisions before deadlines to avoid trial cancellations. Furthermore, even without trial cancellations in the normal phase, collapsing

bounds achieve a higher point rate than a stationary bound (Fig. 3.5a). When the reliability of the evidence is fixed—that is, when only one motion strength is used in the task—a decision variable in a bounded drift-diffusion model represents logLR of either choice being correct (Gold and Shadlen, 2007). In this case, stationary bounds can achieve a desired accuracy with the least number of samples (Wald and Wolfowitz, 1948), and thereby maximizes the point rate with an appropriate bound height (Gold and Shadlen, 2007). In contrast, when more than one motion coherence is interleaved as in my task, collapsing bounds achieves the highest point rate (Drugowitsch et al., 2012; Huang and Rao, 2013). Thus, I conclude that collapsing bound is advantageous for the subjects, and an increased rate of collapse is a rational way to incorporate time pressure.

Given that subjects' decision accuracy was not markedly different between the Normal and Time Pressure phases as shown by the psychometric curves in Figure 3.2b, one might wonder why the subjects did not set their bounds lower and make decisions faster in the Normal phase. For lower coherence levels, however, the difference in decision accuracy was large enough (~3%) for three out of four subjects, so that their bound height in the Normal phase still allowed them to achieve a higher point rate than their lower bound height in the Time Pressure phase.

It is unknown how the brain adjusts the decision bound. However, recent neurophysiological studies suggested the presence of a neural signal that lowers the decision criterion in a time-dependent manner. Neural correlate of evidence accumulation has been

observed in physiological recording from the monkey brain during the random-dot experiment. Neurons in the lateral intraparietal area (LIP) exhibit a firing rate that reflects the accumulation of evidence when one target is placed in their response field (RF) and the other target is placed outside. When the motion direction favors the choice target in the RF of the recorded neuron, the firing rate builds up, on average, proportionally to the motion strength. When the firing rate reaches a threshold value, it signals the end of the decision and the monkey makes a saccadic eye movement to the target in the RF (Roitman and Shadlen, 2002). This time evolution of the firing rate resembles that of the decision variable in the diffusion model.

In addition to this activity pattern, LIP firing rate also encodes an evidence-independent signal. When 0% coherence trials were combined, regardless of monkey's choice, there is no net sensory evidence that favors either direction. Yet, the observed LIP firing rate kept increasing in a time-dependent manner, on average, which cannot be explained by evidence accumulation (Churchland et al., 2008). This time-increasing signal, termed the "urgency" signal, would facilitate monkeys to make decisions as more time elapsed during a trial. In fact, because a physiological threshold for the LIP firing rate is constant across time, adding an evidence-independent signal has the same effect on behavior as lowering bound in the drift-diffusion model. Thus, this evidence-independent signal may be a neural substrate of collapsing bounds in a drift-diffusion model.

Adjustments in the evidence-independent signal have also been observed in recent physiological studies. One study showed that LIP firing rate can reflect the timing of events by

encoding its hazard rate (Janssen and Shadlen, 2005). When the timing of events changes, the dynamics of the firing rate also changed to reflect the new timing. Another study demonstrated that when the monkey was trained to make decisions with shorter reaction time, LIP neurons exhibited an increased rate of build-up in the evidence-independent signal (Hanks). These results suggest that such flexibility in the evidence-independent signal could explain the change in the subjects' bound height across the two phases in my experiment.

Figure 3.1.
Task design.
a, Subjects made decisions about the net motion direction of a dynamic random dot display and made a choice by eye movement to one of the peripheral targets. Subjects received points and feedback based on the result in each trial (see Methods). The feedback indicated the current point rate (green bar), and the history of the past point rates in the session (blue bars).
b, Experimental schedule. The subjects performed practice sessions until their performance became stable. Then they went through two phases of the experiment. The first normal phase had 10 sessions where all of the trials were the standard type. The second time-pressure phase had 20 sessions where a half of trials were the standard type (yellow box) whereas the other half of trials had potential deadlines (white box). **c**, Distribution of reaction time (green) and deadline (red) for subject A. The distribution of deadline is normalized for illustration although deadlines were set only in a half of trials.

Figure 3.1

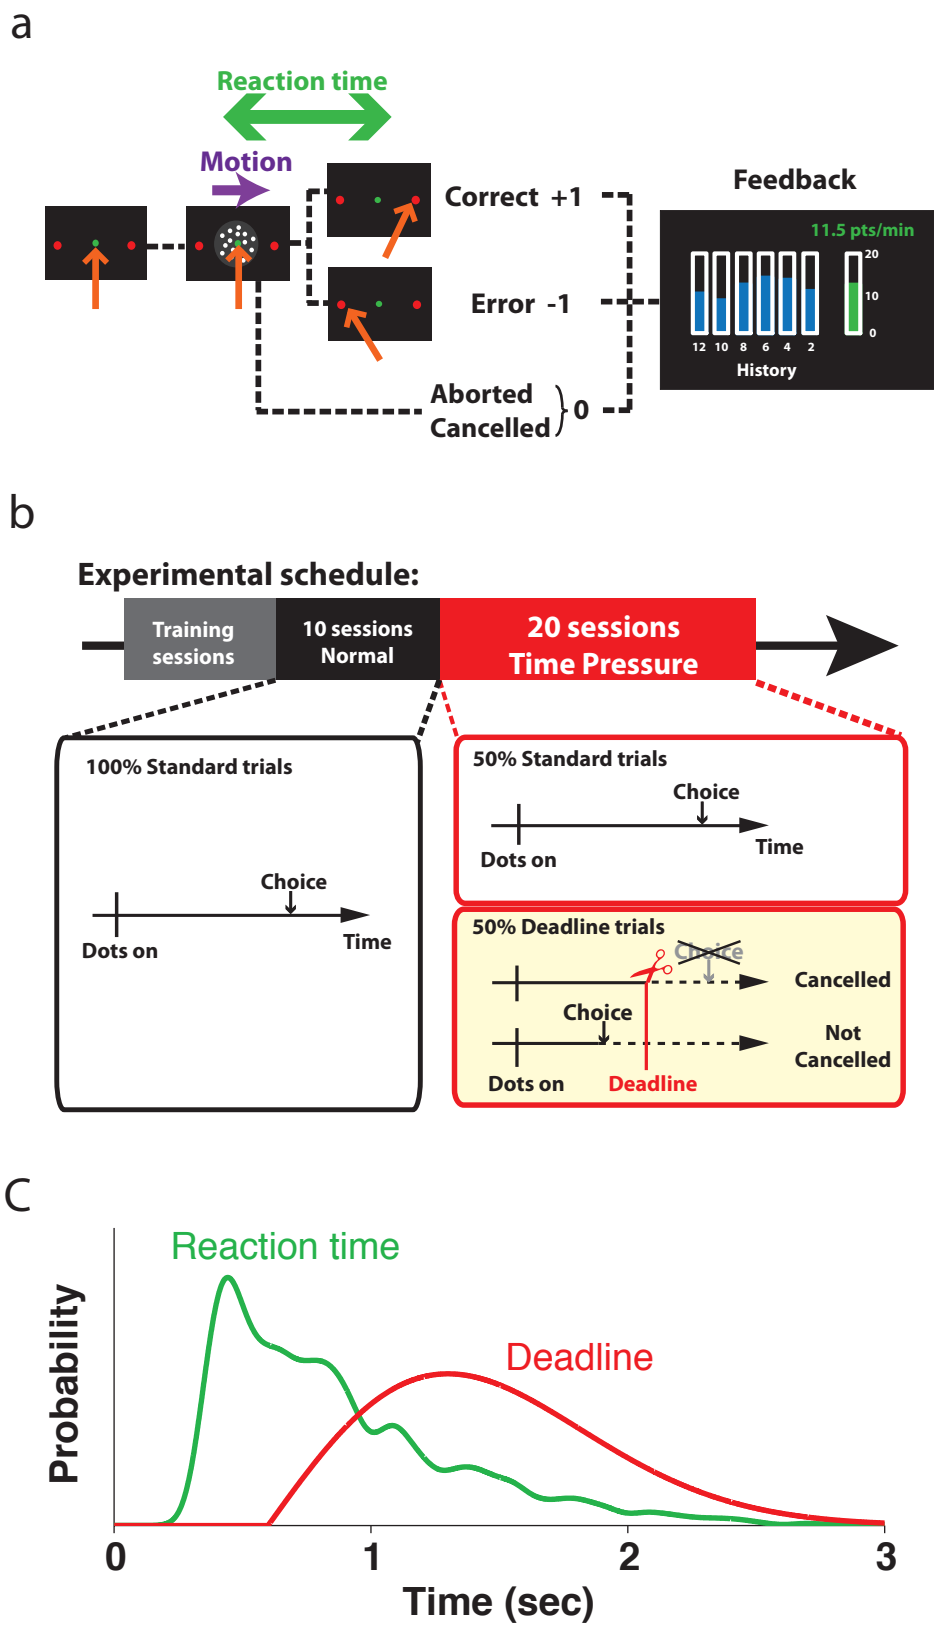
a
Reaction time
Motion
Correct +1
Error -1
Aborted
Cancelled
0
Feedback
11.5 pts/min
20
10
0
12 10 8 6 4 2
History
b
Experimental schedule:
Training sessions
10 sessions Normal
20 sessions Time Pressure
100% Standard trials
Choice
Dots on
Time
50% Standard trials
Choice
Dots on
Time
50% Deadline trials
Choice
Cancelled
Choice
Not Cancelled
Dots on
Deadline
C
Reaction time
Deadline
Probability
0
1
2
3
Time (sec)

Figure 3.2. Effects of time pressure on reaction time and accuracy
in one subject (top) and average across subjects (bottom).
a, Chronometric curve. Mean reaction time for correct (circles) and error choices (squares) are shown for the normal phase (black) and the random cancellation phase (red). Error bars show SEM. The curves indicate the mean reaction time predicted by the collapsing bound model.
b, Psychometric curve. The proportion of correct trials is shown for the normal phase (black) and the random cancellation phase (red). The curves indicate the proportion correct predicted by the collapsing bound model.

Figure 3.2

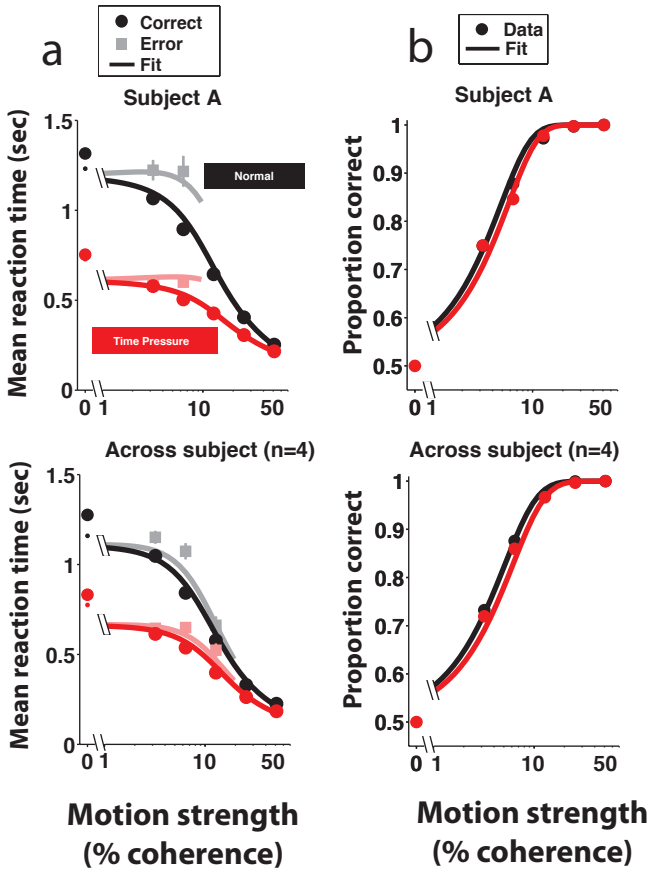
a
Correct
Error
Fit
Subject A
Mean reaction time (sec)
1.5
1
0.5
0
Normal
Time Pressure
0 1
10
50
Across subject (n=4)
Motion strength
(% coherence)
b
Data
Fit
Subject A
Proportion correct
1
0.9
0.8
0.7
0.6
0.5
Across subject (n=4)
Motion strength
(% coherence)

Figure 3.3. Time-dependent accuracy
a, For subject A, trials with each motion strength in each phase are sorted in order of the reaction time and divided into trial bins (100 trials/bin shifted by 10 trials). Each data point shows the mean reaction time and the proportion correct of trials in each bin. Plots for the normal and time-pressure phase are distinguished by black and red, respectively, and the motion strengths are distinguished by the brightness of the color. The curves indicate the time-dependent accuracy predicted by the collapsing bound model.
b, Trials are pooled across subjects, sorted in order of the reaction time, and divided into trial bins (400 trials/bin shifted by 40 trials). The mean reaction time and the proportion correct are plotted in the same way as in A.

Figure 3.3

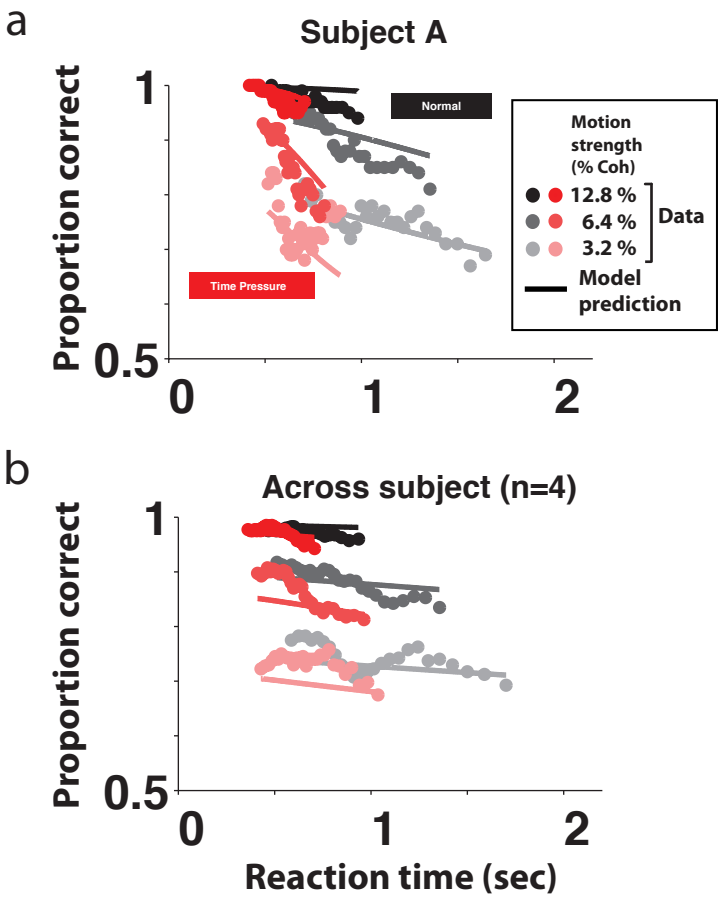
a
Subject A
Proportion correct
1
0.5
0
1
2
Normal
Time Pressure
Motion strength (% Coh)
12.8 %
6.4 %
3.2 %
Data
Model prediction
b
Across subject (n=4)
Proportion correct
1
0.5
0
1
2
Reaction time (sec)

Figure 3.4. Bounded drift-diffusion model
a, Stationary bound model. A stream of evidence is accumulated as a decision variable, which undergoes diffusion with an overall drift until it reaches one of the bounds. The height of the bounds remains stationary during a trial.
b, Collapsing bound model. The dynamics of a decision variable is the same, but the height of the bounds declines as a function of elapsed time (purple curves). The decline was modeled by an exponential decay. The decision time becomes shorter as illustrated by the length of the green arrows.

Figure 3.4

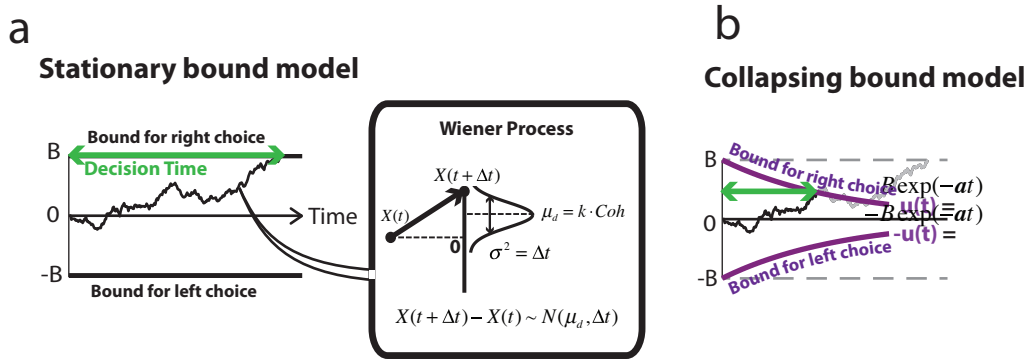
a
Stationary bound model
Bound for right choice
Decision Time
B
0
-B
Time
Bound for left choice
Wiener Process
$X(t+\Delta t)$
$X(t)$
$\mu_d = k \cdot Coh$
$\sigma^2 = \Delta t$
$X(t+\Delta t) - X(t) \sim N(\mu_d, \Delta t)$
b
Collapsing bound model
B
0
-B
Bound for right choice
$B\exp(-at)$
u(t) =
$-B\exp(-at)$
-u(t) =
Bound for left choice

Figure 3.5. Subject's bound height
a, Bounds fitted for each subject (solid curves) and the optimal bound for that subject (dashed curves).
b, Values of best fitted parameters for the normal phase and the random cancellation phase. Error bars show the standard error of the mean across the four subjects.

Figure 3.5

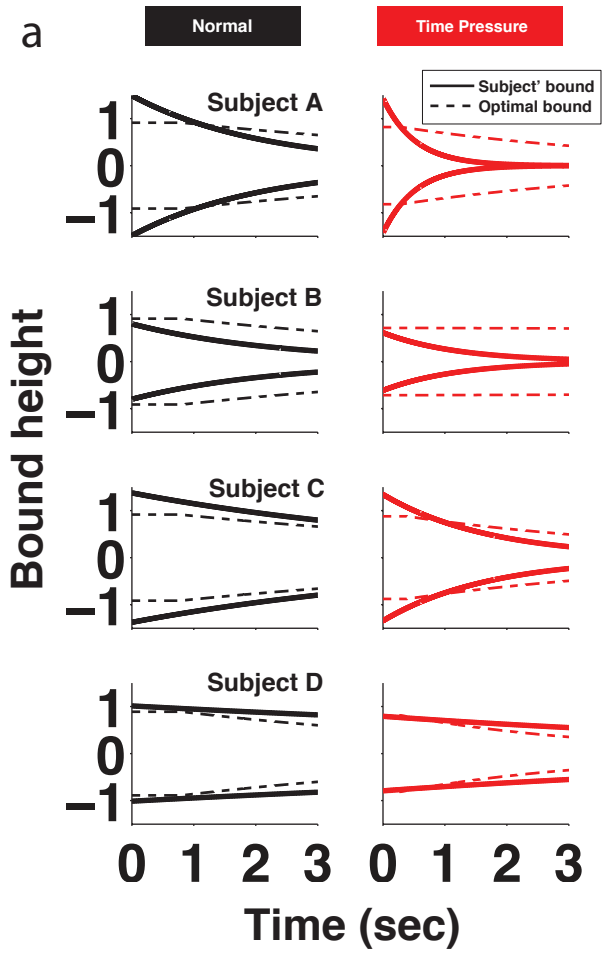
a
Normal
Time Pressure
Subject' bound
Optimal bound
Subject A
Subject B
Subject C
Subject D
Bound height
1
0
−1
0 1 2 3 0 1 2 3
Time (sec)
b
initial bound
rate of collapse
k
mean T_{nd} (sec)
1.5
1
0.5
0
30
20
10
0
0.3
0.2
0.1
0
Normal
Time Pressure

Table 3.1

subject	Sampling distribution for potential deadline				% cancelled	% aborted	
	t_0(s)	σ_{DL}(s)	mean (s)	S.D. (s)		Normal phase	Time-pressure phase
A	0.6	0.7	1.48	0.46	1.4±0.9	1.7±1.3	2.2±0.9
B	0.3	0.6	1.05	0.39	4.3±1.9	3.4±1.1	2.5±0.9
C	0.6	1.0	1.85	0.66	3.6±1.0	2.5±1.3	2.7±2.5
D	0.6	0.6	1.35	0.39	3.7±2.0	0.6±0.5	1.3±1.0

Table 3.1. Potential deadline distribution for each subject

Chapter 4

Seeing is deciding?: comparison of "decision-related" activity in macaque area LIP with and without a preparation for a specific eye movement

Abstract

Putative decision-related activity has been observed in macaque lateral intraparietal area (LIP) in various decision-making tasks that required monkeys to report their decision by a saccadic eye movement. However, it remains unclear whether the LIP activity reflects the decision formation itself or an evolving plan for an impending saccade. To test this, I recorded activities of LIP neurons as I prevented the monkey from planning a specific saccade while it made decisions. We trained a monkey to make binary decisions between right and left saccades based on a sequential presentation of shapes, where shapes provided probabilistic evidence as to which saccade would lead to a reward. Crucially, the exact saccade destination remained unknown to the monkey until the end of shape presentations. Whereas the evidence provided in the early phase of the shape presentation period affected the monkey's choice, LIP activity did not reflect this evidence until the late phase that preceded the time of saccade. The dissociation in the shape's leverage on choice and on LIP activity indicates that LIP activity represents a preparation for an impending eye movement rather than the evidence accumulation itself.

Introduction

Decisions can be flexibly linked to different motor plans or cascaded to make successive decisions. Decisions that are not tied to a 'specific' motor plan are often referred to as "abstract decisions" (Horwitz et al., 2004; Bennur and Gold, 2011; Freedman and Assad, 2011). Understanding the neural mechanisms that transform abstract decisions into actions would reveal neural computations underlying flexible sensorimotor association. To understand where and how such transformation processes occur in the brain, it is essential to find out which areas of the brain are involved in the transformation process and to characterize what type of information is represented in each area.

In most of previous studies on decision-making, the experimenters accessed decisions by requiring subjects to report their decision by actions (e.g. eye movement, lever release). Thus, the formation of abstract decisions was rarely isolated from the development of motor plans (but see Freedman and Assade, 2006 and Bennur and Gold, 2011). As subjects formed a decision, they could also plan a specific action before its initiation, which may explain why neural correlates of decision-making have been often observed in motor-related areas such as the lateral intraparietal area (LIP), frontal eye field (FEF), parietal reach region (PRR), and premotor cortex (Shadlen and Newsome, 2001; Romo et al., 2004; Ding and Gold, 2012) (V. de Lafuente, M. Jazayeri, M. N. Shadlen, unpublished observations).

Among these areas, LIP neurons have been shown to exhibit an elevated persistent activity while monkeys plan a saccadic eye movement to their response field (RF). The same neurons also exhibit decision-related activity while monkeys make decisions and report their decisions by a saccadic eye movement to a choice target (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002). The LIP response reflects the accumulation of evidence in favor of a target in the response field (RF) of a recorded neuron. However, the enforcement of a saccade to report a decision obscured whether the LIP response represents the accumulation of evidence for an abstract decision or a preparation for an impending saccade.

To characterize which type of information is represented in LIP, I recorded LIP responses from a monkey while it made decisions without knowing exactly where to make a saccade to report the decision. While the monkey made decisions based on evidence provided by sequential presentation of shapes, the choice targets were not shown until the monkey finishes acquiring evidence. In addition, I changed the location of the choice targets on a trial-by-trial basis to prevent the monkey from anticipating their exact locations. During this task, LIP response did not reflect evidence in the initial phase of the evidence-accumulation period although this evidence affected the monkey's choice. Towards the end of evidence-accumulation period, LIP response began to reflect evidence robustly as the time of saccade approached. These results suggest that LIP response reflects a preparation for an impending saccade that is modulated by the accumulated evidence, rather than the accumulation of evidence for an abstract decision.

Methods

One male rhesus monkey (*Macaca mulatta*) was implanted with a head post and a recording chamber above the intraparietal sulcus in the right hemisphere. Throughout training and recording sessions, horizontal and vertical eye positions were recorded by an EyeLink infrared video-tracking system (sampling rate, 1000 Hz). Timing of task-related events were monitored and controlled by NIH Rex system. Visual stimuli were displayed using Psychtoolbox (Brainard, 1997; Pelli, 1997) on a CRT monitor with a refresh rate of 75Hz. All animal procedures complied with the National Institutes of Health Guide for the Care and Use of Laboratory Animals, and were approved by the the University ofWashington Animal Care Committee.

Task

The monkey was trained extensively to interpret eight unique shapes as evidence bearing the likelihood that the right or the left choice target would furnish a reward. Among the eight unique shapes, four shapes favored the left target whereas the other four shapes favored the right target (Fig. 4.1, inset).

After acquiring fixation on a central fixation point, putative gray targets (six or twelve targets depending on a session; see Fig. 4.S1) appeared in peripheral symmetric locations around the fixation point. When a total of six targets are presented, three targets appeared on each side of

visual field and they were 45° apart from each other (Fig. 4.S1a,b). When a total of twelve targets are presented, six targets appeared on each side of visual field and they were 30° apart from each other (Fig. 4.S1d). One or two putative targets were placed inside of the RF of a recorded neuron for 6 target sessions and 12 target sessions, respectively.

After a delay of 500ms from the target onset, a sequence of highly visible shapes appeared centered on a vertex of an invisible 3ºx3º grid centered on the fixation point. The shapes were approximately 1.5 by 1.5 deg, high contrast line art with equal perimeter (Fig. 4.1, inset). Four shapes were displayed sequentially one at a time. The length of presentation was 700ms, 500ms, 400ms, 350ms (or 400ms) for the four shapes in a trial (see Fig. 4.S1). Successive shapes were not shown in the same grid location. Each shape was sampled independently and uniformly from the set of eight unique shapes with replacement. In each trial, the computer assigned a reward probabilistically either to the right of left choice target based on the total evidence of the four presented shapes. Specifically, the probability that the left choice target furnishes a reward, $P_r(L)$, is determined as:

$$P_r(L) = \frac{1}{1+10^{-Q}}, (Q = \sum_{i=1}^{4} e_i) \qquad (1)$$

where e_i is the assigned evidence of the *i* th presented shape. The assigned evidence closely approximates logLR of the shape in favor of the left target (see Yang and Shadlen (2007) for more details).

One of gray putative targets was randomly picked on each side of the visual field and turned into red to indicate the location of the right and left choice targets. This disclosure of choice targets occurred either before the first shape onset (early disclosure task) or after the fourth shape offset (delayed disclosure task). The early disclosure task and delayed disclosure task were switched in a block-wise manner in each session. The target disclosure was implemented in slightly different ways in some sessions (Fig. 4.S1b,c). Because the results in each session were qualitatively similar, I included these variants in my analysis to increase statistical power.

After the presentation of four shapes, a disappearance of the fixation point signaled a GO cue, which required the monkey to make a saccadic eye movement to the right or left choice target to indicate the decision. If the monkey chose a reward assigned target, it received a drop of water.

Data acquisition and neuron selection

Quartz-platinum/tungsten electrodes (Thomas Rec. with 1-3 MΩ impedance) were used for recording. Recorded signals from the electrode were amplified and bandpass filtered (150Hz–8kHz) before action potentials (spikes) were detected by a dual voltage-time window discriminator (Plexon).

I recorded from neurons in the posterior one third of LIP in the lateral bank of the intraparietal sulcus and in its ventral part (LIPv) (Lewis and Van Essen, 2000) near the fundus of the sulcus. LIPv was identified by both its anatomical and physiological properties. Recording sites were planned using a postoperative MRI that showed the recording chamber and grid. I

registered these images with the standard MRI supplied with the CARET software (Van Essen, 2002) and targeted the posterior one third of the flat-map representation of LIPv. I then identified LIPv using physiological criteria. That is, I selected isolated single units that showed a robust and spatially selective persistent activity while the monkey waited to execute a saccadic eye movement towards a remembered target location during the memory period of the memory-guided saccade task (Hikosaka and Wurtz, 1983). This property is common in the region of the intraparietal sulcus ~4-8 mm below the cortical surface. A single well-isolated neuron was recorded in each recording session, and a total of 6 neurons were recorded and analyzed.

Analysis of behavioral data

I used the logistic regression to estimate the leverage of the shapes on the choice. A probability that the monkey would choose the left target, $P_c(L)$, was modeled as a logistic function of a "sum of leverages", Q:

$$P_c(L) = \frac{1}{1+10^{-Q}} \qquad (2)$$

which was fit as a generalized linear model (GLM) with binomial error.

To quantify the degree to which the total evidence affected the monkey's choice (Fig. 4.2a,b), I fit two coefficients, β_{0-1} :

$$Q = \beta_0 + \beta_1 \sum_{i=1}^{4} e_i \qquad (3)$$

where e_i is the assigned evidence of the i th shape in each trial. β_0 represents a choice bias and β_1 quantifies the leverage of total evidence on choice.

To estimate how much leverage each of the 1st through the 4th shape had on the choice (Fig 4.2c,d), I fit five coefficients, γ_{0-4} :

$$Q = \gamma_0 + \sum_{i=1}^{4} \gamma_i e_i \qquad (4)$$

γ_0 represents a choice bias and $\gamma_{1\text{-}4}$ quantifies leverage of the first through the fourth shape on choice, respectively.

To estimate the leverage of the eight unique shapes regardless of when they were presented (Fig. 4.2e,f), I fit 8 coefficients, $\hat{w}_{1-8}$:

$$Q = \sum_{j=1}^{8} \hat{w}_j n_j \ . \qquad (5)$$

where n_j is the count of each unique shape during a trial such that $\sum_{j=1}^{8} n_j = 4$. The fitted logistic coefficients ($\hat{w}_{1-8}$) quantify the leverage of each unique shape.

Analysis of physiological data

To plot the peri-stimulus time histograms in Figure 4.3, I aligned a spike train from each trial to the time of five different events (the onset of the first through the fourth shape and the go cue), and divided into five epochs (Fig. 4.3, columns). Each epoch spans the time intervals between the events. In each epoch, the spike trains were sorted into quintiles in order of cumulative evidence in that epoch. For each quintile, the firing rate was computed as a running average of the spike trains (shifting a 50ms average window by 10ms) to plot a response curve.

To assess whether LIP firing rate reflected the evidence from each shape in favor of T_{In}, I used a GLM to quantify the degree to which evidence of each presented shape contributed to the LIP firing rate, FR_k, after the *k* th shape presentation (i.e. during the *k* th shape epoch in Fig. 4.3):

$$FR_k = r_0 + \sum_{i=1}^{k} r_i e_i \qquad (6)$$

where e_i is the assigned evidence of the *i* th shape in each trial. Because all recorded LIP cells had the RF on the left visual hemifield, shapes with positive assigned evidence favored T_{in}, which allowed us to use the same sign convention to model the firing rate. The fitted coefficient r_0 represents the firing rate given neutral evidence (Fig. 4.3, the middle green quintile), and r_i (i = 1,...,k) quantifies a fraction of the firing rate explained by a unit evidence of the i th shape in the k th shape epoch.

Results

Behavior

One monkey was trained to make decisions between a right and left choice based on evidence provided by the sequential presentation of four shapes (Fig. 4.1). Each shape was randomly sampled from a set of eight unique shapes, each of which carried assigned evidence bearing on the likelihood that the left or right choice would furnish a reward. Among the eight unique shapes, four shapes favored the left choice whereas the other four shapes favored the right choice. Their assigned evidence can be quantified in units of log likelihood ratio (logLR) whose sign indicates a favored target and whose magnitude represents the reliability of evidence (Fig. 4.1, inset; see Methods). After the presentation of shapes, a GO cue (disappearance of the fixation point) signaled the monkey to indicate the decision by a saccadic eye movement to the right or left choice target. To prevent the monkey from planning a saccade to a specific target, multiple gray putative targets (6 or 12 targets; see Fig. 4.S1) were presented from the beginning of each trial. One of the putative targets was randomly picked in each visual hemifield and turned into red to disclose the location of the right and left choice targets. The target disclosure occurred either before the onset of the first shape (early disclosure task) or after the offset of the fourth shapes (delayed disclosure task). These two types of task were switched in a block-wise manner in each session. In the early configuration task, the monkey could plan a specific saccade while accumulating evidence supplied by shapes. In the delayed configuration task, however, the monkey cannot plan a specific saccade during the evidence accumulation although the monkey

knows a coarse direction of an impending saccade and could prepare a saccade to several possible locations simultaneously.

The monkey's choice strongly depended on the total evidence of the presented shapes regardless of whether the target locations were revealed before or after the shape presentations (Fig. 4.2a,b). When the total evidence strongly favored the right choice or left choice, the monkey almost always chose the right or left choice target, respectively. In contrast, when the total evidence weakly favored one choice over the other in the middle of the graph, the monkey's choice was distributed more evenly. When each choice function was fitted with a logistic curve (Eq. 2,3), the slope of each curve indicated a significant influence of the total evidence ($p< 10^{-99}$) although the slope was shallower for delayed target trials than early target trials ($p<0.01$). This suggests that, regardless of the uncertainty of the impending saccade destination, the monkey accumulated evidence to make decisions although the accumulation was less efficient when the monkey could not plan a specific eye movement.

To verify that the monkey made decisions in a consistent way for both types of the task, I further assessed whether the monkey accumulated all four shapes presented in a trial and how well the monkey learned the assigned evidence of each unique shape. I first evaluated whether shapes presented in a particular order in the sequence had appropriate leverage on choice by a logistic regression, which models the log odds of a left choice as a weighted sum of the assigned evidence for each of the first through the fourth shape (Eq. 2,4). In this analysis, four fitted logistic coefficients quantified the leverage of the *i* th shape on choice (i = 1,..., 4; Fig. 4.2c,d).

The analysis revealed that shapes presented in each epoch significantly affected the choice in both tasks ($p<10^{-16}$). Shapes presented earlier in a trial had smaller leverage than those presented later. This trend was observed in both tasks although the trend was more pronounced in the delayed disclosure task, explaining the shallower slope in the choice function (Fig. 4.2b). I also assessed how much leverage each of the eight unique shapes had on the monkey's choice regardless of when these shapes were presented. I used a different form of a logistic regression (Eq. 2,5), which models the log odds of a left choice as a sum of the leverage of shapes shown in each trial. The estimated leverage of the eight unique shapes matched their assigned evidence in terms of their sign and rank order, and showed a high correlation with the assigned evidence ($r = 0.94$ and 0.95, $p< 0.001$; Fig. 4.2e,f), suggesting that the monkey learned each shape's weight, though not perfectly. This result also rules out an alternative decision strategy in which the monkey uses only particular shapes for decisions because this strategy would reduce the leverage of other shapes to zero. The leverage of shapes tended to be smaller in the delayed disclosure task than in the early disclosure task, which could be explained by the smaller leverage of earlier shapes in the delayed disclosure task than in the early disclosure task (Fig. 4.2c,d). Overall, despite these subtle differences, the behavioral results suggest that the monkey accumulated evidence of presented shapes during both types of the task in a similar fashion.

Physiology

While the monkey performed the task, I recorded from six well isolated neurons in the ventral part of LIP (LIPv) (Lewis and Van Essen, 2000) that showed spatially selective persistent

activity during the delay period of a memory guided saccade task (see Methods). During the recording, I first mapped the RF of a recorded neuron, and placed one (in 6 target session) or two (in 12 target session) putative targets in the RF (Fig. 4.1, shaded blue area). After the target disclosure, I analyzed trials where a disclosed choice target was presented in the RF, so that I could measure LIP response in favor of the target in the RF (T_{in}). Because all recorded cells had their RF on the left visual hemifield, shapes favored T_{in} if they were assigned positive evidence favored T_{in}.

In the following section, I will first describe LIP response when the monkey was allowed to plan a saccade to a disclosed choice target from the beginning of trials (early disclosure task, Fig. 4.3a). By contrasting with this response, I will then describe LIP response when the monkey was prevented from planning a specific saccade (delayed disclosure task, Fig. 4.3b).

As the shapes were presented during a trial, LIP responses diverged depending on the cumulative evidence. Figure 4.3a shows the evolution of LIP response after each shape presentation and the go cue during a trial. In each epoch, LIP response curves were sorted into quintiles in order of the cumulative evidence in favor of T_{in}. LIP response began to diverge depending on the evidence from the first shape ~250ms after their onset, which agrees well with a response delay reported previously in a similar experiment (Yang and Shadlen, 2007). As each successive shape was presented, the response continued to diverge depending on the cumulative logLR of shapes. In each epoch, LIP response linearly reflected the total evidence in favor of T_{in} ($p<0.001$; Fig. 4.3a, insets).

To quantify how robustly LIP response reflected accumulated evidence, I used a linear regression to model the LIP response in each epoch as a weighted sum of evidence from previous epochs (Fig. 4.4a-e). The fitted coefficients quantify a fraction of LIP response in one epoch explained by a unit of evidence from each previous epoch. LIP response in each epoch reflected evidence from all previous epochs ($p<0.05$), although shapes presented recently were more strongly represented than those presented earlier in a trial. This result indicates that LIP responses reflect the accumulation of evidence throughout a trial.

In contrast, LIP responses were markedly reduced during the shape presentation when the monkey was prevented from planning a specific eye movement in the delayed disclosure task although the monkey could still prepare a saccade to a coarse direction (i.e., to the right or left visual field) (Fig. 4.3b, Fig. 4.4f-j). During the first epoch, LIP response did not reflect the evidence from the first shape ($p=0.48$; Fig. 4.4f) despite its significant leverage on the monkey's choice (Fig. 4.2d). In the second epoch, LIP response reflected the evidence of the second shape, but did not reflect the evidence of the first shape ($p=0.99$; fig. 4.4g). Surprisingly, however, from the third epoch, LIP response reflected the evidence from all shapes that had already been presented by that epoch, including the first shape ($p<0.05$; fig. 4h-j). These results suggest that the evidence from the first shape was retained and accumulated for decision even though it was not reflected in LIP response in the first two epochs. Furthermore, LIP response can reflect the accumulated evidence even when the monkey cannot plan a specific saccade during later epochs.

Our results indicate that LIP response reflects the total evidence when averaged across trials, but this may not happen on a trial-by-trial basis. The evidence-dependent response can also result from a mixture of binary levels of the firing rate, each of which represents a commitment to the right or left choice. Even if the firing rate in each trial took either a high or low binary value, the firing rate would spuriously vary with the cumulative evidence when averaged across trials. Indeed, although the activity seems to depend on the total evidence during the pre-saccadic period (Fig. 4.3, epoch after the go cue), if trials are grouped by the monkey's choice (T_{in} or T_{out}), the presaccadic response (the last 50ms before saccades) in each group loses its evidence-dependence ($p>0.1$). In contrast, while the monkey accumulated evidence in the preceding shape epochs, the LIP response showed graded representation of evidence even when trials are grouped by choice ($p<0.05$ except for the first two epochs in the delayed disclosure task). This observation implies that the graded responses did not arise as a consequence of averaging responses comprising different mixture of T_{in} and T_{out} responses. Rather, the LIP response reflects an evolving decision variable that will ultimately lead to a T_{in} or T_{out} choice.

Discussion

By preventing the monkey from planning a specific saccade, I tested whether firing rates of LIP neurons reflect the accumulated evidence for an abstract decision or a preparation for an impeding saccade. If LIP firing rate reflects an abstract decision, the firing rate should vary with accumulated evidence throughout the four shape epochs even when the monkey could not plan a

specific saccade in the delayed disclosure task. However, I observed that LIP firing rate did not vary with the accumulated evidence during early epochs of the delayed disclosure task (Fig. 4.4f) even though evidence provided in this epoch affected the monkey's choice (Fig. 4.2d). Towards later shape epochs, LIP response began to reflect the accumulated evidence by gradually increasing the response gain, which can be most parsimoniously interpreted as a signal for saccade preparation as will be discussed below. Together, these results indicate that LIP response reflects a preparation for an impending saccade rather than the accumulated evidence for an abstract decision.

The lack of evidence-dependent LIP response in the first shape epoch of the delayed disclosure task (Fig. 4.3b, 4.4f) is a key to my interpretation, but I need to rule out the following case that trivially explains the lack of the LIP response. Suppose that the monkey actually planned a saccade to a non-disclosed target that happened to lie outside of the RF of the recorded neurons. Then, although some group of LIP cells whose RF included that non-disclosed target reflected the accumulated evidence, I recorded from LIP cells in another group that did not show the evidence-dependent activity. As the activity in the former group increased towards the end of accumulation period, its activity partially spread to the recorded group of cells. However, this is unlikely for the following reasons. First, the RF location of a recorded neuron differed from session to session, and the arrangement of possible target locations also varied across sessions. If the monkey was planning a saccade to one of non-disclosed targets during the delayed disclosure task, the monkey would have planned saccades to the target in the RF of a recorded LIP cell in some sessions by chance. Nonetheless, before the target disclosure, none of the recorded cells

showed evidence-dependent responses as strongly as those in the early disclosure task. Second, if the monkey had planned a saccade to a non-disclosed target on the RF side but outside of the RF, LIP cells would have shown a decreased response when evidence favored a choice on the RF side. But, none of my recorded cells showed such a response pattern. Therefore, my results argue against that trivial account for the lack of the evidence-dependent activity. Instead, my interpretation is that the evidence from the earlier shapes was accumulated in the brain but only weakly reflected in LIP response during the earlier shape epochs because of a lack of an oculomotor preparation. If LIP is the locus that accumulates evidence in this task, evidence from the first shape would have been lost because of the low signal to noise ratio in the response. Thus, at least during the earlier shape epochs, I speculate that the accumulation of evidence primarily took place elsewhere in the brain other than LIP. Although it is unknown which part of the brain accumulates evidence during that epoch, one candidate locus is the prefrontal cortex, in which the neural activity is deemed to encode abstract rules (Wallis et al., 2001) and object categories (Freedman et al., 2001) even when monkeys do not have to plan any actions.

After the first shape epoch, the evidence-dependent response increased toward the time of a saccade initiation even before the target disclosure (Fig. 4.3b, 4.4h-i). I argue that this response reflects preparation for an impending saccade rather than the accumulated evidence for an abstract decision because each LIP cell showed a higher activity when the evidence favored a choice on the RF side. Similar preparatory activities were observed in the superior colliculus (SC) (Basso and Wurtz, 1998). A different number of putative targets were presented in each trial to change the probability that a visual stimulus would be selected as the choice target. When

monkeys prepared a saccade for multiple putative targets before the target disclosure, the preparatory activity of 'buildup' SC neurons negatively correlated with the number of putative targets. Recently, similar preparatory activities were also observed during a reach task in the parietal reach region (PRR) (Klaes et al., 2011). In my study, the monkey was in a similar situation. Because the monkey could coarsely estimate the direction of an impending saccade, the monkey could prepare for a saccade to that direction although the exact direction remained unknown until the target disclosure. Preparation for multiple non-disclosed targets might have induced the evidence-dependent LIP response even before the target disclosure.

Importantly, the effect of cumulative evidence on the LIP response depended on the baseline activity induced by neutral evidence (Fig. 4.3 green traces), which differed between the shape epochs and the two tasks. The effect of cumulative evidence on firing rates can be captured in each epoch by regressing the firing rate against the cumulative evidence (Fig. 4.3 insets, dashed lines). As the baseline activity increased, the same level of evidence induced a larger increment or decrement in the LIP firing rate from the baseline (shown as a steeper slope in Fig. 4.3 insets). This observation implies that it was just the gain of LIP activity that differed between the tasks, but the effect of evidence on choice was comparable between the tasks (Fig. 4.2). Figure 4.3 suggests that the gain is higher for the early disclosure task, in which the monkey could prepare an impending saccade as it accumulated evidence. Furthermore, the gain increased for later epochs during a trial as the time of saccade initiation approached.

In contrast with my conclusion, Bennur and Gold (2011) argued that LIP response represents the accumulated evidence for an abstract decision. In their experiment, monkeys were required to discriminate the direction of noisy dot motion. Two categorical decisions about the motion direction (rightward vs leftward) were associated with colors of two targets (red vs green, respectively) and monkeys reported their decision by a saccade to a target in the associated color. Two targets appear in fixed symmetric locations above and below a fixation point, but their color identity was masked in gray until a variable timing of color disclosure—that is, one target turned to red and the other one turned to green. As one of the targets was presented in the RF of a recorded LIP cell, its response reflected the accumulation of motion evidence even before the color disclosure.

Their study reached a different conclusion from mine because of a difference in predictability of an impending saccade direction. On the one hand, I interpreted the emergence of evidence-dependent LIP response before the target disclosure as a preparatory activity for an impending saccade because the monkey could anticipate a coarse direction of saccade in my experiment. Moreover, the evidence-dependent activity was absent in early shape epochs in the delayed disclosure task, suggesting that the LIP response does not represent evidence for an abstract decision during these epochs. On the other hand, Bennur and Gold argued that the evidence-dependent LIP response before the color disclosure represents evidence accumulated for abstract decisions because an impending saccade direction—either to the RF or away from the RF—is statistically unpredictable in their experiment.,

However, their argument is still inconclusive. Because there are only two possible target configurations ("red top target vs green bottom target" or "green top target vs red bottom target"), the monkey could prepare a saccade by assuming one configuration even before the color disclosure. If the assumed configuration turned out to be wrong upon the color disclosure, the monkey could flip a saccade direction from the planned direction. When viewed this way, the observed evidence-dependent LIP response can be interpreted to represent a saccadic preparation rather than evidence for abstract decisions. One way to rule out this interpretation is to simultaneously record from multiple cells that have an overlapping RF on one target. If some cells show an increased response while other cells show a decreased response to the motion stimulus before the color disclosure, such conflicting responses would indicate that the LIP response represents evidence for abstract decisions rather than a saccadic preparation because the monkey cannot simultaneously plan a saccade 'to' and 'away from' the same target.

On the contrary, in order to conclusively show that LIP response represents planning of an impending saccade, one also needs to show that the absence of a saccadic plan abolishes the evidence-dependent LIP response. However, this would be extremely difficult to demonstrate because one cannot prevent monkeys from planning a saccade to an arbitrary location in the visual field, and thus some LIP neurons might exhibit evidence-dependent modulation if their RF overlaps with the location. To show the absence of evidence-dependent LIP response, one needs to record from multiple LIP neurons simultaneously such that their RFs tile the entire visual field, and demonstrate that none of them exhibit evidence-dependent response.

Although such exhaustive experiments are currently difficult to conduct in monkeys, recent advances in *in vivo* imaging in awake behaving rodents might provide a complementary approach. By monitoring response from a large number of cells simultaneously, one can more conclusively address the presence of decision-related activities in a particular brain area (Komiyama et al., 2010; Harvey et al., 2012). In addition to characterizing neural activities in each brain area, it is also important to study communications between two or more brain areas to understand how neural circuits configure specific motor plans based on abstract decisions. Although conventional approaches such as orthodromic and antidromic stimulation are available to find projection cells from one area to another (Berman and Wurtz, 2010, 2011), recent development in optogenetics will allow us to identify them with a higher efficiency (Znamenskiy and Zador, 2013). Combination of these techniques holds promise for unraveling how the brain transforms abstract decisions into appropriate motor plans and actions—a key neural mechanism underlying flexible sensorimotor transformations.

Figure 4.1
Design of the early vs delayed disclosure task. After the monkey stares at a central fixation point, 6 or 12 gray putative targets appear in symmetric peripheral locations around the fixation point (only the 6 target version is shown; see Fig. 4.S1). A shaded blue area (invisible to the monkey) indicates the response field (RF) of the recorded neuron, which includes one (for 6 target sessions) or two (for 12 target session) putative targets. After a short latency, a sequence of four shapes appears near the fixation point one at a time. Each shape gives probabilistic information as to whether the reward is more likely on the right or left target. Among the multiple putative targets, one of the targets is randomly picked on each side of the visual field and turned into red to disclose the location of the right and left choice target. The location of the choice targets are disclosed either before the onset of the 1st shape (early disclosure task) or after the offset of the fourth shape (delayed disclosure task). After the presentation of four shapes, the offset of fixation point signals the go cue for the monkey to make a saccade to the right or left disclosed choice target.

Figure 4.1

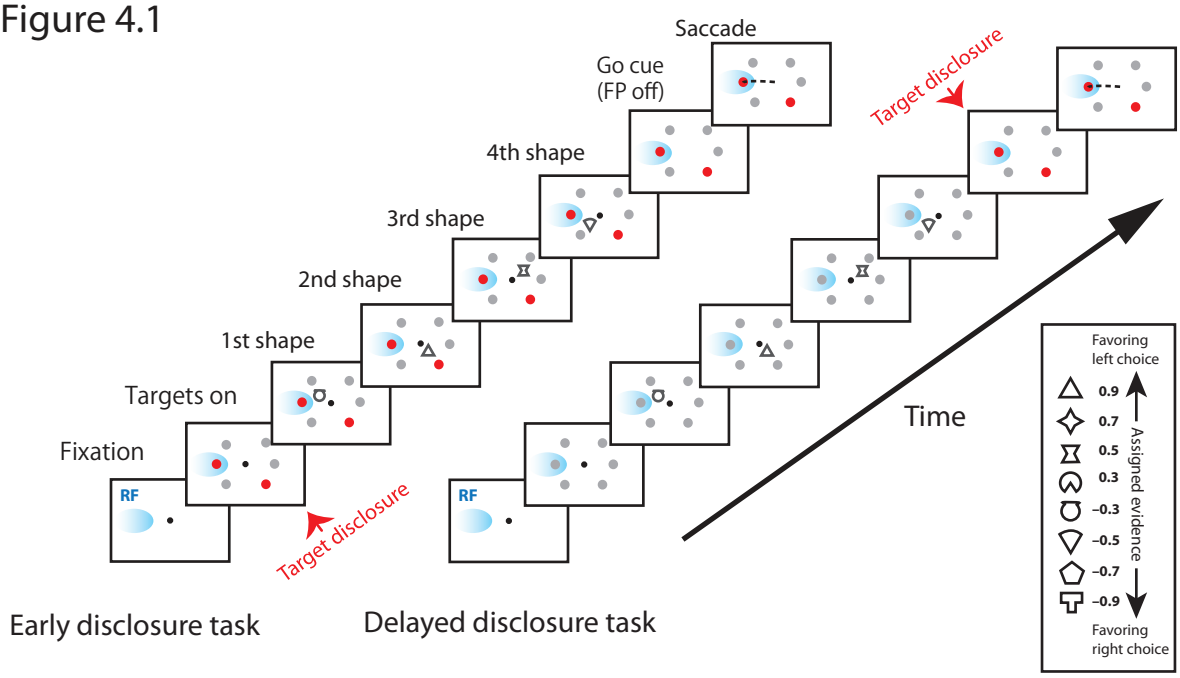
Saccade
Go cue
(FP off)
4th shape
3rd shape
2nd shape
1st shape
Targets on
Fixation
RF
Target disclosure
RF
Target disclosure
Time
Favoring
left choice
0.9
0.7
0.5
0.3
–0.3
–0.5
–0.7
–0.9
Assigned evidence
Favoring
right choice
Early disclosure task
Delayed disclosure task

Figure 4.2
Behavioral performance
a, **b**, The proportion of left choice is plotted against the total evidence from four presented shapes in a trial. A red curve shows a logistic curve fitted to each choice function (Eq. 2, 3) and its slope value ($\beta_1 \pm$ 95% c.i.) is indicated in the upper left corner.
c, **d**, The leverage of the 1st thought the 4th shape on choice ($\gamma_{1\text{-}4}$ in Eq. 4). Error bars show standard errors.
e, **f**, The leverage of each unique shape on choice ($\hat{w}_{1-8}$ in Eq. 5) is plotted against the assigned evidence. Error bars show standard errors.

Figure 4.2

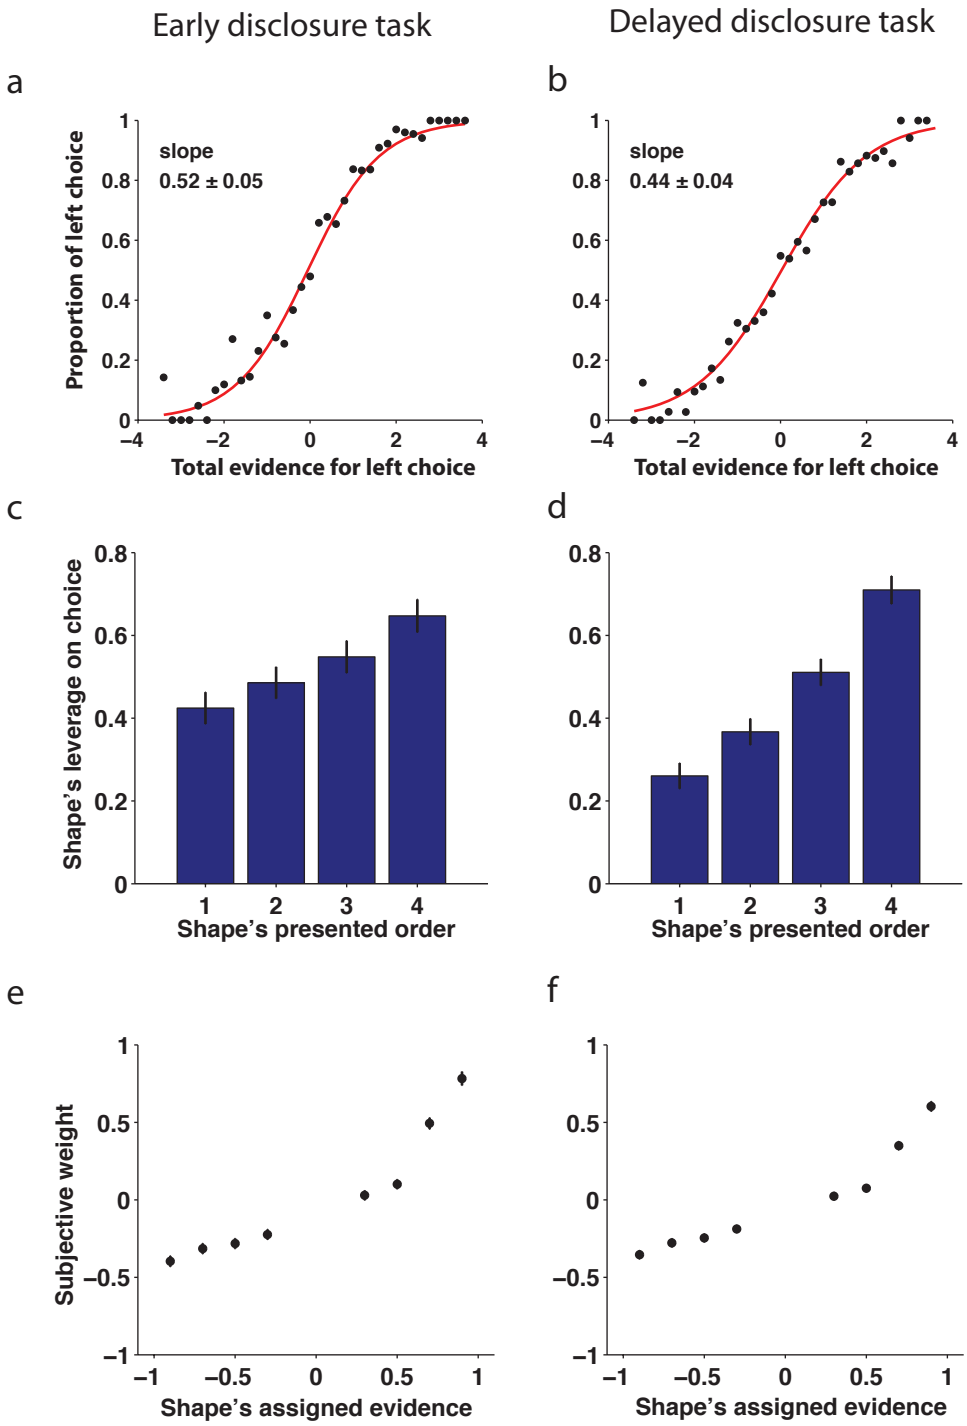
Early disclosure task
Delayed disclosure task
a
b
slope
0.52 ± 0.05
slope
0.44 ± 0.04
Proportion of left choice
Total evidence for left choice
Total evidence for left choice
c
d
Shape's leverage on choice
Shape's presented order
Shape's presented order
e
f
Subjective weight
Shape's assigned evidence
Shape's assigned evidence

Figure 4.3
Evolution of LIP population firing rate
a, LIP response for the early disclosure task. The response curves are aligned to the onset of an event (the 1st through the 4th shape or the go cue) in each column, and plotted until the onset of the next successive event. The response curves are sorted into 5 quintiles in order of total logLR in each epoch (column) and smoothed by running mean (50ms time window shifted by 10 ms). The warmer color indicates that the total evidence in that quintile favors a choice target on the RF side of the visual field (i.e. left side). The Inset above each column shows the mean firing rate during each shape epoch as a function of total evidence. The mean firing rate in each shape epoch is computed from a time period marked by gray shading on the abscissa. The mean firing rate after the go cue is computed from the period between the time of go cue and 100ms after the go cue or the time of saccade, whichever happened first. Color-filled circles represent the quintiles in each epoch, and a dashed line shows a regression line. The slope (± 95% c.i.) of the regression line is indicated in the upper left corner.
b, LIP responses are plotted in the same manner for the delayed disclosure task. The color of the abscissa indicates before (black) and after (red) the target disclosure.

Figure 4.3

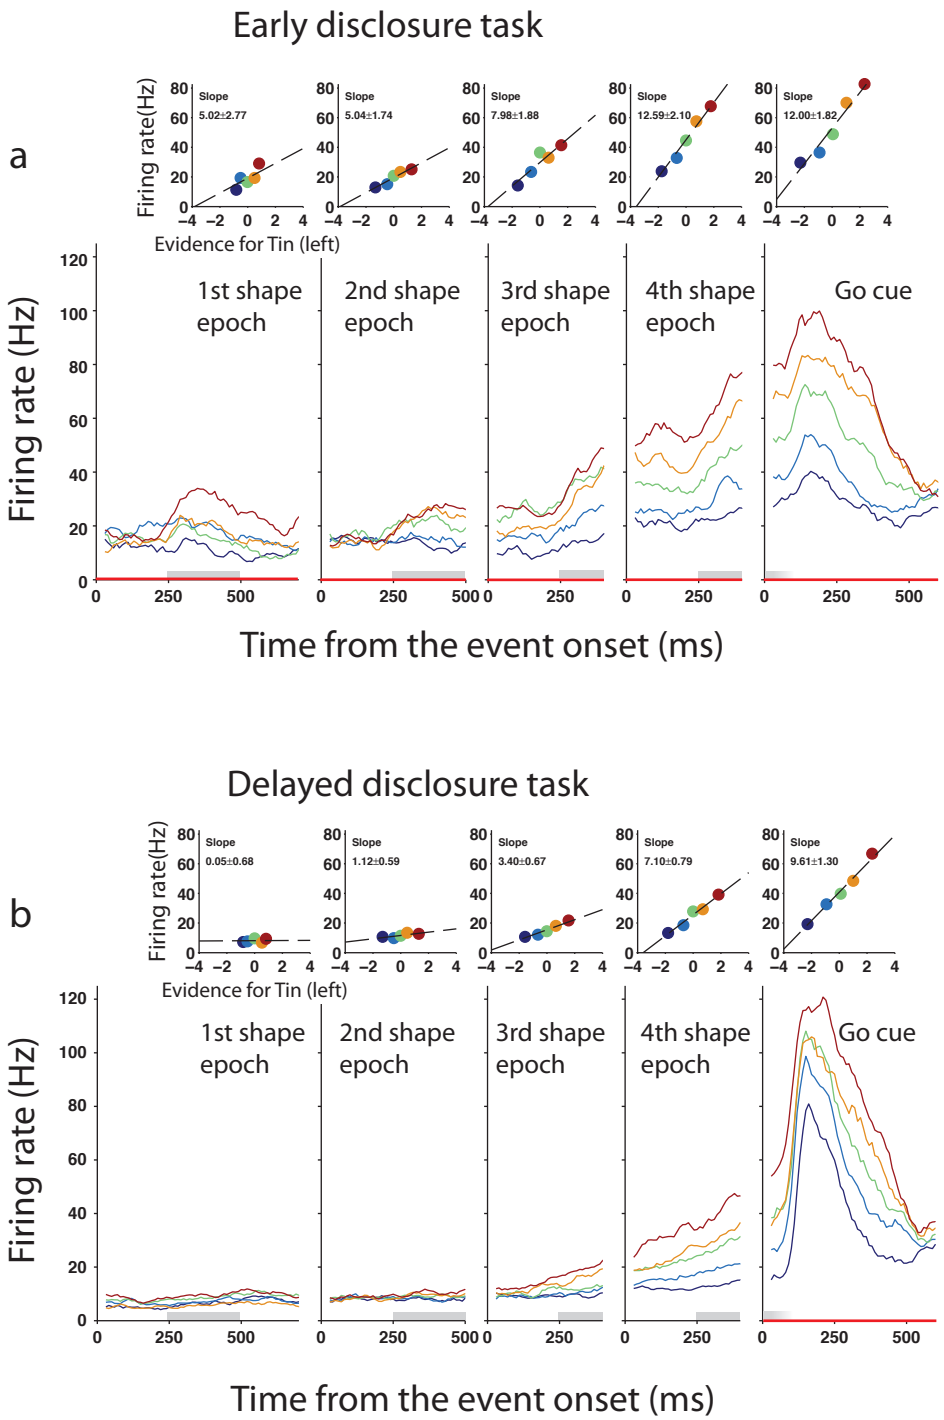
Early disclosure task
a
Firing rate(Hz)
Slope
5.02±2.77
Slope
5.04±1.74
Slope
7.98±1.88
Slope
12.59±2.10
Slope
12.00±1.82
Evidence for Tin (left)
1st shape epoch
2nd shape epoch
3rd shape epoch
4th shape epoch
Go cue
Firing rate (Hz)
Time from the event onset (ms)
Delayed disclosure task
b
Slope
0.05±0.68
Slope
1.12±0.59
Slope
3.40±0.67
Slope
7.10±0.79
Slope
9.61±1.30
Evidence for Tin (left)
1st shape epoch
2nd shape epoch
3rd shape epoch
4th shape epoch
Go cue
Firing rate (Hz)
Time from the event onset (ms)

Figure 4.4
Leverage of shape evidence on LIP firing rate. The leverage of shapes are shown for the early disclosure task (a-e) and for the delayed disclosure task(f-j). The LIP firing rate in each epoch is modeled as a weighted sum of evidence conferred by shapes presented in that epoch and its preceding epochs (Eq. 6). The estimated weights ($r_{1\text{-}4}$ in Eq. 6) quantify the magnitude of a response modulation induced by a unit evidence (logLR=1). Error bars show standard errors.

Figure 4.4

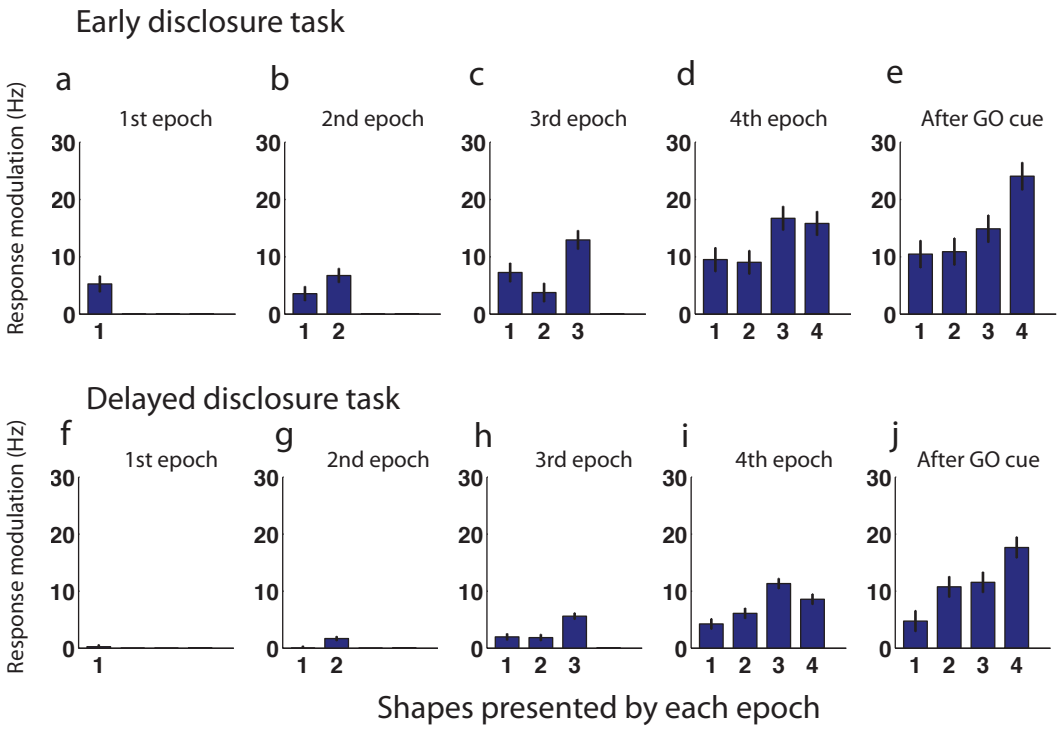
Early disclosure task
a
1st epoch
b
2nd epoch
c
3rd epoch
d
4th epoch
e
After GO cue
Response modulation (Hz)
30
20
10
0
1
1 2
1 2 3
1 2 3 4
Delayed disclosure task
f
1st epoch
g
2nd epoch
h
3rd epoch
i
4th epoch
j
After GO cue
Response modulation (Hz)
Shapes presented by each epoch

Figure 4.S1
Variation in the appearance of targets across recording sessions. Although the event sequence was consistent with the one shown in Figure 4.1 (a), some sessions had a different implementation of the target disclosure (b-c) or a larger number of putative targets (d). The time length of events is indicated above each panel.
a, Same as in Figure 4.1. Three putative targets on each side are 45° apart from each other.
b, Same as Figure 4.1 except that after the early target disclosure, the choice targets turned back to gray upon the onset of the first shape.
c, Same as Figure 4.1 except that gray putative targets were not displayed throughout a trial.
d, Twelve targets were displayed at the same eccentricity and equally spaced around the fixation point.

Figure 4.S1

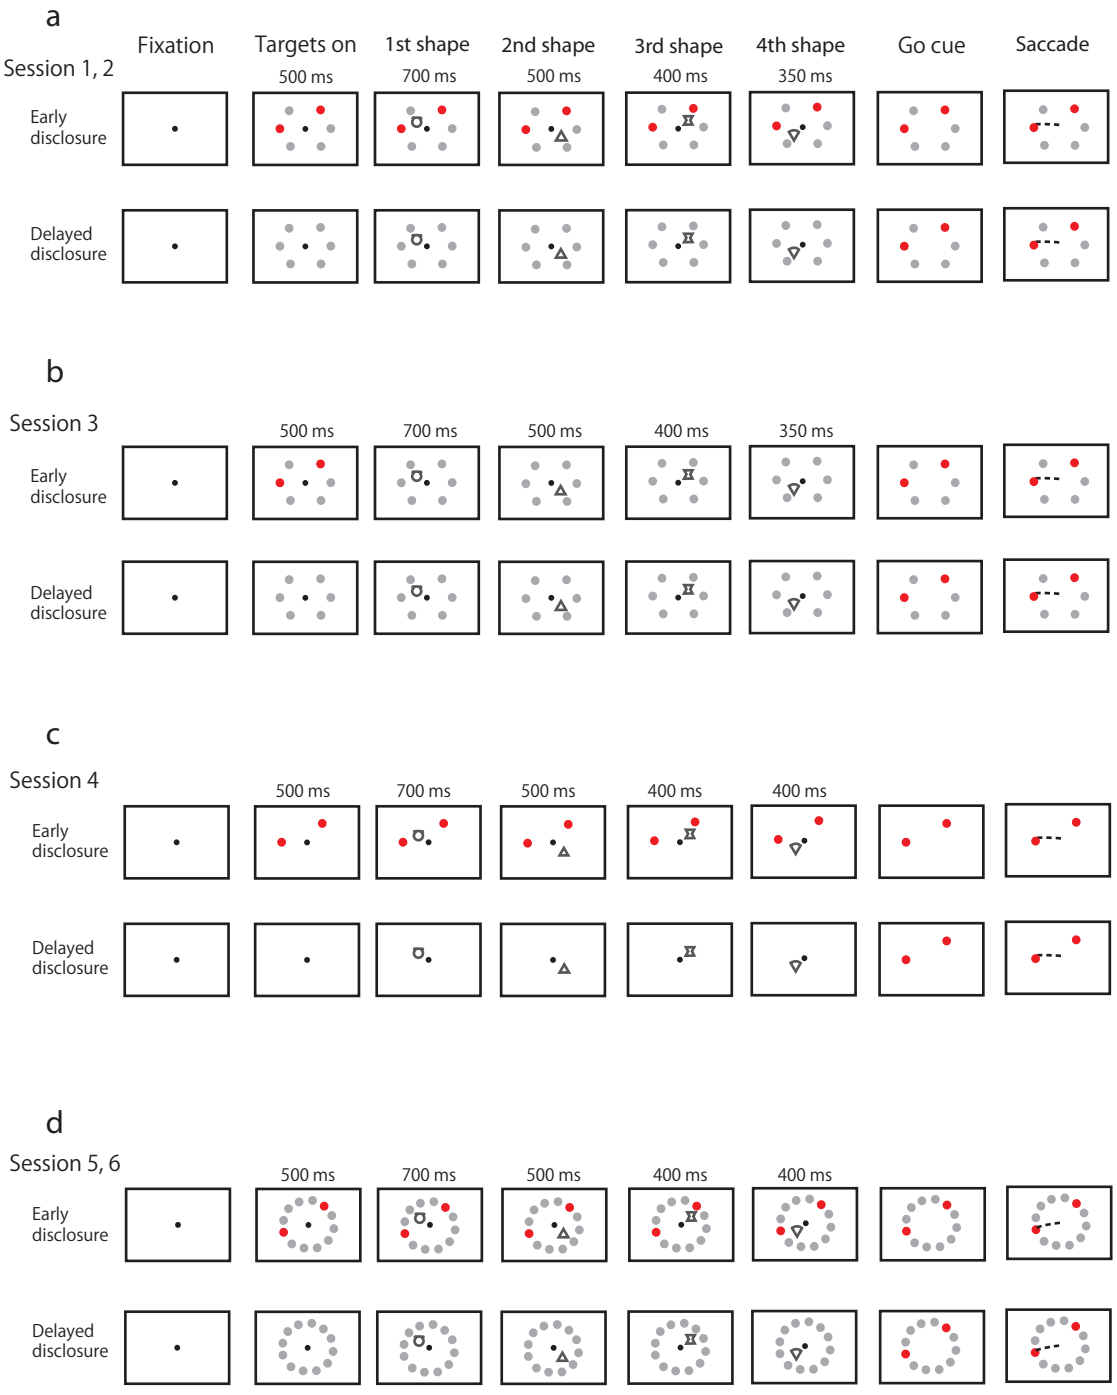
a
Session 1, 2
Fixation
Targets on
1st shape
2nd shape
3rd shape
4th shape
Go cue
Saccade
500 ms
700 ms
500 ms
400 ms
350 ms
Early disclosure
Delayed disclosure
b
Session 3
500 ms
700 ms
500 ms
400 ms
350 ms
Early disclosure
Delayed disclosure
c
Session 4
500 ms
700 ms
500 ms
400 ms
400 ms
Early disclosure
Delayed disclosure
d
Session 5, 6
500 ms
700 ms
500 ms
400 ms
400 ms
Early disclosure
Delayed disclosure

References

Barash S, Bracewell RM, Fogassi L, Gnadt JW, Andersen RA (1991) Saccade-related activity in the lateral intraparietal area. II. Spatial properties. J Neurophysiol 66:1109-1124.

Barlow HB, Levick WR, Yoon M (1971) Responses to single quanta of light in retinal ganglion cells of the cat. Vision research Suppl 3:87-101.

Basso MA, Wurtz RH (1998) Modulation of neuronal activity in superior colliculus by changes in target probability. The Journal of neuroscience : the official journal of the Society for Neuroscience 18:7519-7534.

Behrens TE, Hunt LT, Rushworth MF (2009) The computation of social behavior. Science 324:1160-1164.

Bennur S, Gold JI (2011) Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:913-921.

Berman RA, Wurtz RH (2010) Functional identification of a pulvinar path from superior colliculus to cortical area MT. The Journal of neuroscience : the official journal of the Society for Neuroscience 30:6342-6354.

Berman RA, Wurtz RH (2011) Signals conveyed in the pulvinar pathway from superior colliculus to cortical area MT. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:373-384.

Brainard DH (1997) The Psychophysics Toolbox. Spatial vision 10:433-436.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. The Journal of neuroscience : the official journal of the Society for Neuroscience 12:4745-4765.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. Visual neuroscience 10:1157-1169.

Burr DC, Santoro L (2001) Temporal integration of optic flow, measured by contrast and coherence thresholds. Vision research 41:1891-1899.

Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. Nature Neuroscience 11:693-702.

Cisek P, Puskas GA, El-Murr S (2009) Decisions in changing conditions: the urgency-gating model. The Journal of neuroscience : the official journal of the Society for Neuroscience 29:11560-11571.

DiCarlo JJ, Zoccolan D, Rust NC (2012) How does the brain solve visual object recognition? Neuron 73:415-434.

Ding L, Gold JI (2012) Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. Cereb Cortex 22:1052-1067.

Ditterich J (2006) Evidence for time-variant decision making. The European journal of neuroscience 24:3628-3641.

Ditterich J, Mazurek ME, Shadlen MN (2003) Microstimulation of visual cortex affects the speed of perceptual decisions. Nature Neuroscience 6:891-898.

Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A (2012) The cost of accumulating evidence in perceptual decision making. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:3612-3628.

Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415:429-433.

Fechner GT (1860) Elemente der Psychophysik: Breitkopf und Härtel.

Felleman DJ, Van Essen DC (1991) Distributed hierarchical processing in the primate cerebral cortex. Cerebral Cortex (New York, NY: 1991) 1:1-47.

Fetsch CR, DeAngelis GC, Angelaki DE (2013) Bridging the gap between theories of sensory cue integration and the physiology of multisensory neurons. Nature reviews Neuroscience 14:429-442.

Fetsch CR, Pouget A, DeAngelis GC, Angelaki DE (2012) Neural correlates of reliability-based cue weighting during multisensory integration. Nat Neurosci 15:146-154.

Franzier P, Yu AJ (2008) Sequential hypothesis testing under stochastic deadlines. Advances in Neural Information Processing Systems 20:465-472.

Freedman DJ, Assad JA (2011) A proposed common neural mechanism for categorization and perceptual decisions. Nat Neurosci 14:143-146.

Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. Science 291:312-316.

Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. Experimental brain research Experimentelle Hirnforschung Experimentation cerebrale 70:216-220.

Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. Trends in cognitive sciences 5:10-16.

Gold JI, Shadlen MN (2002) Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. Neuron 36:299-308.

Gold JI, Shadlen MN (2007) The neural basis of decision making. Annual Review of Neuroscience 30:535-574.

Green DM, Swets JA (1966) Signal Detection Theory and Psychophysics: John Wiley and Sons.

Hanks TD, Ditterich J, Shadlen MN (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nature Neuroscience 9:682-689.

Harvey CD, Coen P, Tank DW (2012) Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature 484:62-68.

Hikosaka O, Wurtz RH (1983) Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. J Neurophysiol 49:1268-1284.

Horwitz GD, Batista AP, Newsome WT (2004) Representation of an abstract perceptual decision in macaque superior colliculus. J Neurophysiol 91:2281-2296.

Huang Y, Rao RP (2013) Reward optimization in the primate brain: a probabilistic model of decision making under uncertainty. PloS one 8:e53344.

Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. The Journal of neuroscience : the official journal of the Society for Neuroscience 25:10420-10436.

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

Kable JW, Glimcher PW (2009) The neurobiology of decision: consensus and controversy. Neuron 63:733-745.

Kiani R, Shadlen MN (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324:759-764.

Kiani R, Hanks TD, Shadlen MN (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 28:3017-3029.

Klaes C, Westendorff S, Chakrabarti S, Gail A (2011) Choosing goals, not rules: deciding among rule-based action plans. Neuron 70:536-548.

Knill DC (2007) Robust cue integration: a Bayesian model and evidence from cue-conflict studies with stereoscopic and figure cues to slant. J Vis 7:5 1-24.

Komiyama T, Sato TR, O'Connor DH, Zhang YX, Huber D, Hooks BM, Gabitto M, Svoboda K (2010) Learning-related fine-scale specificity imaged in motor cortex circuits of behaving mice. Nature 464:1182-1186.

LaMotte RH, Mountcastle VB (1975) Capacities of humans and monkeys to discriminate vibratory stimuli of different frequency and amplitude: a correlation between neural events and psychological measurements. J Neurophysiol 38:539-559.

Lewis JW, Van Essen DC (2000) Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. The Journal of comparative neurology 428:112-137.

Lynch JC, Graybiel AM, Lobeck LJ (1985) The differential projection of two cytoarchitectonic subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus. The Journal of comparative neurology 235:241-254.

Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. Nat Neurosci 9:1432-1438.

Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. Annu Rev Neurosci 24:167-202.

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

O'Connell RG, Dockree PM, Kelly SP (2012) A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. Nat Neurosci 15:1729-1735.

Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. Journal of Vision 5:376-404.

Parker A, Hawken M (1985) Capabilities of monkey cortical cells in spatial-resolution tasks. Journal of the Optical Society of America A, Optics and image science 2:1101-1114.

Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. Annu Rev Neurosci 21:227-277.

Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spatial vision 10:437-442.

Pouget A, Beck JM, Ma WJ, Latham PE (2013) Probabilistic brains: knowns and unknowns. Nat Neurosci 16:1170-1178.

Ratcliff R (1978a) Theory of Memory Retrieval. Psychological Review 85:59-108.

Ratcliff R (1978b) A theory of memory retrieval. Psychological Review 85:59-108.

Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. Neural computation 20:873-922.

Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 22:9475-9489.

Romo R, Hernandez A, Zainos A (2004) Neuronal correlates of a perceptual decision in ventral premotor cortex. Neuron 41:165-173.

Schall JD, Morel A, King DJ, Bullier J (1995) Topography of visual cortex connections with frontal eye field in macaque: convergence and segregation of processing streams. The Journal of neuroscience : the official journal of the Society for Neuroscience 15:4464-4487.

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

Shadlen MN, Kiani R (2013) Decision making as a window on cognition. Neuron 80:791-806.

Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. The Journal of neuroscience : the official journal of the Society for Neuroscience 16:1486-1510.

Simen P, Contreras D, Buck C, Hu P, Holmes P, Cohen JD (2009) Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. Journal of experimental psychology Human perception and performance 35:1865-1897.

Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. Trends in neurosciences 27:161-168.

Tanaka K (1996) Inferotemporal cortex and object vision. Annu Rev Neurosci 19:109-139.

Tolhurst DJ, Movshon JA, Dean AF (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. Vision research 23:775-785.

Tsetsos K, Gao J, McClelland JL, Usher M (2012) Using Time-Varying Evidence to Test Models of Decision Dynamics: Bounded Diffusion vs. the Leaky Competing Accumulator Model. Frontiers in neuroscience 6:79.

Van Essen DC (2002) Windows on the brain: the emerging role of atlases and databases in neuroscience. Current opinion in neurobiology 12:574-579.

Wald A (1947) Sequential Analysis: Wiley.

Wald A, Wolfowitz J (1948) Optimum Character of the Sequential Probability Ratio Test. The Annals of Mathematical Statistics 19:326-339.

Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. Nature 411:953-956.

Watamaniuk SN, Sekuler R (1992) Temporal and spatial integration in dynamic random-dot stimuli. Vision research 32:2341-2347.

Wimmer GE, Shohamy D (2012) Preference by association: how memory mechanisms in the hippocampus bias decisions. Science 338:270-273.

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

Znamenskiy P, Zador AM (2013) Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. Nature 497:482-485.

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature 370:140-143.

Shinichiro Kira

August, 2014

Lab Address: 630 West 168th St., P&S Building Room 16-409, New York, NY 10032, USA
E-mail: shinkira@uw.edu, sk3752@columbia.edu

EDUCATION

December 2012 – August 2014 **Columbia University**, New York, New York, USA.
- Shadlen Lab (thesis projects continued)

September 2008 – August 2014 **University of Washington**, Seattle, Washington, USA.
Ph.D.
- Shadlen Lab and Neurobiology and Behavior Program.

April 2002 – March 2008 **Tokyo Medical and Dental University**, Tokyo, Japan.
M.D.

October 2005 – February 2006 **Imperial College of Science, Technology and Medicine**, London, UK.
- Biophysics Department, The Blackett Laboratory.

CURRENT PROJECTS (Advisor: Dr. Michael N. Shadlen, Columbia University)

2011 – 2014 ***Thesis Project: "Neural mechanism of routing information to a decision circuit."***
- Testing whether neurons in the lateral intra parietal area (LIP) encode a commitment to an option prior to subsequent oculomotor planning, or just a specific eye movement.
- Recording from LIP neurons while preventing monkeys from planning a specific eye movement during decision-making.

2008 – 2014 ***Thesis Project: "Probabilistic reasoning by neurons."***
- Behaviorally testing whether monkeys accumulate a stream of discrete pieces of evidence up to a decision criterion when they can keep observing the evidence until they terminate decisions.
- Recording from LIP neurons to test whether their firing rate encodes accumulated evidence and a decision criterion.

2008 – 2014 ***"An effect of time pressure on perceptual decision-making."***
- Manipulated the time cost of decision by introducing cancellations of trials while human subjects performed the random dot motion discrimination task.
- Under the increased time cost, subjects lowered decision criterion to make faster decisions.

MANUSCRIPTS IN PREPARATION

Kira, S., Yang, T., Shadlen, M.N. Neural mechanism for decisions based on sequential samples of evidence that differ in reliability.

Kira S., Shadlen, M.N. The effect of time pressure can be explained by a dynamic change in the evidence required for a decision.

ABSTRACTS / PRESENTATIONS

October 2012	Kira S, Yang T, Shadlen MN. "A neural mechanism for reasoning from sequential samples of evidence that differ in reliability" Program No. 730.05, Annual Meeting of Society for Neuroscience, New Orleans, LA, USA
November 2010	Kira S, Shadlen MN. "An effect of time pressure on perceptual decision-making" Program No. 503.7 Annual Meeting of Society for Neuroscience, San Diego, California, USA
February 2010	Kira S, Shadlen MN. "The effect of time pressure on decision-making" Computational and Systems Neuroscience, Salt Lake City, Utah, USA.
November 2007	Hanks TD, Kiani R, Kira S, and Shadlen MN. "Evidence for the incorporation of psychological urgency in decision making." Program No. 507.4 Annual Meeting of Society for Neuroscience, San Diego, California, USA.

HONORS / AWARDS

February 2010	**2010 Cosyne Travel Fellowships**
September 2008 – August 2012	**Overseas Research Scholarship from the Nakajima Foundation**
March 2007	**Overseas Research Scholarship from Tokyo Medical and Dental University**