

Elapsed Decision Time Affects the Weighting of Prior Probability in a Perceptual Decision Task

Timothy D. Hanks, Mark E. Mazurek, Roozbeh Kiani, Elisabeth Hopp, and Michael N. Shadlen

Howard Hughes Medical Institute, Department of Physiology and Biophysics, and National Primate Research Center, University of Washington, Seattle, Washington 98195-7290

Decisions are often based on a combination of new evidence with prior knowledge of the probable best choice. Optimal combination requires knowledge about the reliability of evidence, but in many realistic situations, this is unknown. Here we propose and test a novel theory: the brain exploits elapsed time during decision formation to combine sensory evidence with prior probability. Elapsed time is useful because (1) decisions that linger tend to arise from less reliable evidence, and (2) the expected accuracy at a given decision time depends on the reliability of the evidence gathered up to that point. These regularities allow the brain to combine prior information with sensory evidence by weighting the latter in accordance with reliability. To test this theory, we manipulated the prior probability of the rewarded choice while subjects performed a reaction-time discrimination of motion direction using a range of stimulus reliabilities that varied from trial to trial. The theory explains the effect of prior probability on choice and reaction time over a wide range of stimulus strengths. We found that prior probability was incorporated into the decision process as a dynamic bias signal that increases as a function of decision time. This bias signal depends on the speed–accuracy setting of human subjects, and it is reflected in the firing rates of neurons in the lateral intraparietal area (LIP) of rhesus monkeys performing this task.

Introduction

Decision making is an integrative process that combines evidence, prior knowledge, and expected rewards and costs. Probability, or some other measure of belief, furnishes a common framework for combining these factors (Jaynes, 2003). Thus, it has been suggested that there is a probabilistic evaluation of sensory signals in the brain (Barlow, 1969; Carpenter and Williams, 1995; Zemel et al., 1998; Rao, 2004; Jazayeri and Movshon, 2006; Ma et al., 2006; Gold and Shadlen, 2007; Beck et al., 2008). Ultimately, the combined information that bears on a decision appears to converge at the level of single neurons (Platt and Glimcher, 1999; Schall and Thompson, 1999; Glimcher, 2003; Romo et al., 2004; Sugrue et al., 2004; Gold and Shadlen, 2007; Platt and Huettel, 2008; Rorie et al., 2010). Together, these ideas suggest that neurons in the brain combine probabilistic signals into a decision variable (DV). This raises the question of how the brain converts a spike rate into units of probability or degree of belief. In its most basic form, this can be addressed by examining how a neural representation of a DV incorporates probability associated with prior knowledge—that is, prior probability.

The dominant theories for the incorporation of prior probability into a decision process involve fixed changes of either the DV or decision rule (Edwards, 1965; Link and Heath, 1975; Carpenter and Williams, 1995; Gold et al., 2008; Ratcliff and McKoon, 2008; Simen et al., 2009). Consistent with these theories, prior probability has been shown to have a static representation in the brain (Basso and Wurtz, 1998; Platt and Glimcher, 1999). However, it is unclear how the brain combines this type of representation of prior probability with other sources of information, which may confer more or less leverage on the decision. Some information sources are more reliable than others. For example, sampling two red balls in a row from one of two urns that contain a 90:10 mixture of red/white balls is a reliable cue that the urn in question contains more red balls. In contrast, the same evidence would be less reliable if the urns contained 55:45 mixtures. This is a common situation for a decision maker: in addition to uncertainty about the hypothesis, there is uncertainty about the reliability of the evidence.

We propose that, when a subject is deliberating based on a stream of evidence of unknown reliability, the brain can exploit the elapsed time of the decision process to combine prior probability with the representation of the evidence in a way that gives less weight to the evidence when it is less reliable. Many decisions can be explained by the accumulation of evidence to a criterion level or “bound.” If the evidence and bound share units of probability, then the latter would establish the probability of a correct decision, regardless of decision time or stimulus strength. Stronger and weaker stimuli would lead to faster and slower accumulations of probability to the same level, and decision accuracy would not depend on elapsed time.

More typically, however, stronger stimuli result in both faster and more accurate choices than weaker stimuli (Link, 1992; Rat-

Received Oct. 26, 2010; revised Jan. 29, 2011; accepted Feb. 24, 2011.

This work was supported by the Howard Hughes Medical Institute (HHMI), National Eye Institute Grant EY11378, and National Center for Research Resources Grant RR00166. T.D.H. was supported by an HHMI predoctoral fellowship. M.E.M. was supported by National Institutes of Health Training Grant GM07108, a Poncin grant, and an Achievement Rewards for College Scientists fellowship. E.H. was supported by the German Academic Exchange Service. We thank John Palmer, Ruben Moreno-Bote, Alex Pouget, Jan Drugowitsch, and Tianming Yang for thoughtful discussions and Melissa Mihali and Lori Jasinski for technical assistance.

Correspondence should be addressed to Michael N. Shadlen, University of Washington Medical School, Box 357290, Seattle, WA 98195-7290. E-mail: shadlen@u.washington.edu.

DOI:10.1523/JNEUROSCI.5613-10.2011

Copyright © 2011 the authors 0270-6474/11/316339-14\$15.00/0

cliff and Rouder, 1998; Ratcliff and Smith, 2004; Palmer et al., 2005), as depicted in psychometric functions. Thus, the bound does not represent a fixed amount of probability but a threshold on some other quantity, such as neural activity. Depending on the stimulus, the bound represents a different probability of a correct choice. The relationship between the accumulated evidence and probability of a correct choice changes with elapsed decision time. Thus, when considering stimuli of mixed reliability, more rapid decisions are associated on average with more reliable evidence and slower decisions are associated with less reliable evidence.

In principle, the brain could exploit this relationship when combining prior probability with accumulated sensory evidence. At shorter elapsed decision times, when the evidence on average corresponds to higher reliability, prior knowledge would be incorporated as a smaller bias signal relative to a fixed amount of accumulated evidence. At the longer elapsed decision times, when the evidence on average corresponds to lesser reliability, prior knowledge would be incorporated as a larger bias signal relative to that same amount of accumulated sensory evidence. Because the relationship between the accumulated evidence and probability changes with time, this would allow the influence of prior knowledge to be constant in terms of probability. We show that the brain incorporates this regularity when weighting prior probability in a direction discrimination task.

Materials and Methods

Behavioral task. Two human (one female, one male) and four rhesus monkey (two female, two male) subjects were trained to perform a motion discrimination task. A chin rest and forehead bar were used to stabilize the heads of the human subjects for the duration of each experimental session; a head post was used for the same purpose for the monkeys. Stimuli were presented on a computer monitor (75 Hz frame rate) using the Psychophysics Toolbox for Matlab (Brainard, 1997). Eye position was monitored in humans using the EyeLink infrared video tracking system (sampling rate, 250 Hz) and in monkeys using a scleral search coil (1 kHz). Trials began with the appearance of a single dot that the subject was required to fixate. After a variable delay (at least 500 ms for the monkeys), two bright red choice targets appeared at an equal distance from the fixation point and 180° apart. After a variable delay, the random-dot motion stimulus appeared in an aperture 5–10° in diameter that was either centered at the fixation point or located up to 10° eccentric, along a direction orthogonal to the axis formed by the choice targets. The motion stimulus consists of sets of randomly positioned dots, which are shown for one video frame and then updated 40 ms later (e.g., dots in frame 1 were updated in frame 4; dots in frame 2 were updated in frame 5, and so forth). To update, each dot is either displaced by 14.4 min arc (i.e., 6.0°/s) with probability C , termed the motion coherence (coh), or replaced at a random location. The dot density was $16.7 \text{ dots} \cdot \text{deg}^{-2} \cdot \text{s}^{-1}$.

For each trial, there were two possible directions of motion, differing by 180° and corresponding to the directional axis formed by the choice targets. Motion strength (the percentage of coherently moving dots) was chosen randomly from a set, $C \in \{0, 3.2, 6.4, 12.8, 25.6, 51.2\}$. For one of the human subjects, an alternative set was used, $C \in \{1.6, 3.2, 6.4, 12.8, 25.6, 51.2\}$. The subject's task was to determine the direction of coherent motion, which it indicated by making a saccade to the appropriate choice target (e.g., right target for rightward motion, left target for leftward motion). The subjects could indicate a decision at any time after motion onset. Subjects received positive feedback for all correct choices and on a random percentage of the trials determined by the prior probability (see below) when there was no net motion (0% coherence). The positive feedback for monkeys was a liquid reward; for humans, it was a morale-enhancing auditory tone.

We changed the prior probability of the direction of motion in different blocks of trials. In neutral prior probability blocks, each of the two directions of motion was shown with 50% probability. In unequal prior

probability blocks, one of the directions was shown with 80% probability and the other with 20% probability. By convention in the unequal priors blocks, we define positive motion coherence as the direction of higher probability and negative motion coherence as the direction of lower probability. In the neutral priors blocks, positive is defined to correspond to the rightward direction. In a small subset of neurons for one monkey, we used a pair of intermediate priors (67 and 33%; four cells for 67% and five cells for 33%). We observed intermediate effects, consistent with the conclusions drawn from the larger dataset. These data are not included in the analyses.

For the humans, we provided verbal instructions specifying the prior probability. Blocks with different prior probabilities were collected in different sessions. Because we could not usefully provide verbal instructions for the monkeys, we developed a routine to help them recognize the change in prior probability. Sessions always started with a neutral prior probability block consisting of 200–400 trials. Next, we presented a block of trials with unequal prior probability. The prior was chosen to favor the direction that the monkey chose less often in the preceding neutral prior block; those biases were generally small (equivalent to $-1.7 \pm 0.2\%$ coh, on average). The unequal priors block was signaled with 20 cue trials consisting of 100% coherence motion in the direction of the higher prior. These cue trials were followed by 300–600 trials in the new condition, or until the session had to be stopped as a result of the monkey's failure to perform the task or the loss of adequate neural isolation. Finally, in some experiments, a third block of trials was shown with the prior probability favoring the opposite direction of motion. This block was also signaled with cue trials as above. Even with this structured technique, the monkeys generally did not adjust their bias immediately. Logistic regression applied to a moving window of ± 30 trials indicates that this adjustment took ~ 200 trials to reach stability for the first change and ~ 300 trials for the second, so we analyzed only the trials after these adjustment periods.

The monkeys were trained to achieve a similar, stable speed–accuracy regime, suitable for neural recording. We used a combination of rewards and penalties to counter the natural tendency of monkeys to respond with short latencies. To minimize anticipation, a random interval drawn from a truncated exponential distribution (range, 500–3000 ms; mean, 1200 ms) separated presentation of choice targets and random dot motion. For all monkeys, a timeout penalty was added to the standard intertrial interval. This penalty was graded to punish faster errors (duration range, 400–4000 ms). For one monkey, reward was conditionally delayed so that, although reward was always given for correct choices, it was delivered no sooner than 1000 ms after motion onset. For another monkey, in some training sessions, reward size was adjusted to encourage fast and slow reaction time (RT) at easy and difficult trials. This protocol was used only for a subset of the training and was not used when recording behavioral or physiological data. For most data presentation, behavior was combined across all monkeys; similar results were found in each individually.

The human subjects were trained to work in two different speed–accuracy regimes. This was achieved using verbal instructions and extensive practice. First, the humans were trained to perform the task at a single speed–accuracy set point with neutral priors for the two directions of motion until they reached a stable psychometric threshold and RTs were stable for multiple sessions (at least six sessions for each subject). When switching speed–accuracy regimes for the neutral priors case, subjects were told to operate in a different speed regime by responding faster or slower. We monitored RTs on a session-by-session basis to assess stability in the new regime, requiring at least two sessions of adjustment with the switch. Unequal priors sessions for a given speed–accuracy regime always followed neutral priors sessions for the same regime. For the unequal priors sessions, subjects were instructed to maintain their general speed regime; they were not given additional instructions about a specific target RT. The subjects were naive to the theory being tested and its predictions.

All training, surgery, and experimental procedures were in accordance with the National Institutes of Health *Guide for Care and Use of Laboratory Animals* and were approved by the University of Washington Animal Care Committee or Human Subjects Committee.

Behavioral analyses. Behavioral data were fit using a bounded accumulation model. The model works by accumulating momentary evidence to an upper bound (+ A) or a lower bound ($-A$), corresponding to the two direction choices. We refer to the accumulated evidence as the decision variable in this model. Positive evidence favors one choice and negative evidence favors the other choice. As described above, when the priors are unequal, we define the positive direction to be the one with higher prior probability. The momentary evidence gathered in each time step is drawn from a Gaussian distribution with unit variance for 1 s and mean μ determined by a linear transform of the motion strength: $\mu = kC$, where C is the motion strength, and k is a free parameter that scales the motion strength appropriately. This relationship is reasonable because the expected difference in firing rates between direction-selective neurons in the middle temporal area (MT; also known as area V5) is known to vary linearly, on average, as a function of motion strength (Britten et al., 1993). Both the momentary evidence and the decision variable of this model can be related to neural responses (Mazurek et al., 2003; Shadlen et al., 2006). The bound reached first by the accumulated evidence determines the choice, and the decision time is determined by how long it takes to reach that bound. RT is a combination of this decision time with an additional stimulus-independent nondecision time, t_{nd} . This basic model with three free parameters (k , A , and t_{nd}) must be modified to explain the effect of prior probability on choice and RT (strategies 2 and 3, below).

We use three strategies to analyze the effect of prior probability on behavior. The first, logistic regression, offers a descriptive (i.e., atheoretical) measure of the magnitude of the choice bias:

$$P_+ = \{1 + \exp[-(\beta_0 + \beta_1 C + \beta_2 I)]\}^{-1}, \quad (1)$$

where P_+ is the probability that the subject chooses the direction favored by the prior, C is motion strength (with sign indicating motion direction according to the convention described above), and I is an indicator variable (1 for the block with unequal prior and 0 otherwise). The β_i are fitted coefficients. β_0 represents the behavioral bias under neutral priors, which was negligible in these experiments, β_1 represents the effect of motion on the log odds of a “positive direction” choice, and β_2 quantifies the bias attributable to the prior. It is convenient to represent the prior-induced bias as an equivalent change in motion strength, β_2/β_1 . Statistical tests of the null hypothesis (no effect of prior) are based solely on β_2 ($H_0: \beta_2 = 0$).

The second strategy is to alter the parameters of the bounded accumulation model by permitting an offset of the starting point of the DV. This essentially implements a static bias signal. It introduces one additional free parameter (V_0) in addition to the three described above for the bounded accumulation model.

The third strategy instantiates our theory of a time-dependent bias signal, $\lambda(t)$, that is determined by a subject’s accuracy at a given elapsed decision time. This falls within a class of models with a dynamic bias signal (DBS). Unless otherwise indicated, we use DBS to refer to the particular formulation of $\lambda(t)$, derived in accordance with our theory, as follows. We first calculate a time-dependent accuracy (TDA) function from the neutral prior condition, which is simply the subject’s accuracy as a function of elapsed decision time across the ensemble of stimuli, that is, $P_c(t)$. A fine-grained estimate of the TDA function is obtained directly from the bounded accumulation model fit to the neutral prior data.

We next postulate a mapping between the value of the decision variable $V(t)$ at the bounds ($\pm A$) and the log posterior odds (LPO) of a correct choice as a function of terminating the process at time t , which is simply the log odds of the TDA function at that time. Because we consider the mapping in log space, it simplifies the combination of probabilities. Under this idea, the log odds of the prior probability simply adds to the LPO represented by the DV at that time. The implementation rests on an approximation that the LPO at any particular time is a linear function of the DV at that time. In fact, the relationship is nonlinear, but the approximation is reasonable. It simplifies the representation of the prior to be a fraction of the bound height. Under this implementation, the prior effectively pushes the DV closer to the positive bound (by our sign convention) by the

fraction of the LPO corresponding to the bound at the given elapsed decision time. Thus, the dynamic bias signal is as follows:

$$\lambda(t) = \frac{\log\left(\frac{\phi P}{1 - \phi P}\right)}{\log\left(\frac{P_c(t)}{1 - P_c(t)}\right)} A(t), \quad (2)$$

where P is the prior probability that motion will be in one direction (the one we call positive by convention). The expression contains an additional parameter, ϕ , that scales the prior probability, P , to account for the possibility that a subject misestimates the prior.

Notice that the log term in the numerator of Equation 2 is a constant. This makes sense because the prior is fixed for the duration of the decision—actually for the duration of a block of hundreds of decisions. However, the “bias signal” $\lambda(t)$ increases monotonically as a function of time because the LPO of a correct choice diminishes with longer decision times. An intuitive rationale for the nature of this relationship stems from three considerations. (1) Each sample of momentary evidence is proportional to the log of the likelihood ratio (LLR) of drawing that sample under the two possible directions (Gold and Shadlen, 2001). (2) For a fixed motion strength, the constant of proportionality between sample and LLR is also fixed. This second consideration implies that, under neutral priors, the bound height corresponds to the LPO of making a correct choice for any one motion strength at any time. A flat bound therefore implies that the LPO is constant as a function of time. (3) The LPO of making a correct choice varies with motion strength. This is clearly illustrated by the logistic psychometric function. If accuracy changes with motion strength, by definition the LPO changes as well. Together, these three considerations imply that the same bound represents a different LPO for each coherence (Link, 1992; Shadlen et al., 2006). Nevertheless, each stopping time is associated with a different mixture of coherences. In particular, longer stopping times are associated with a larger fraction of low coherence stimuli. Thus, if the decision has not terminated yet, it is increasingly likely that the evidence is less reliable. The equation above for $\lambda(t)$ was derived precisely to take this relationship into account in determining the bias signal that is added to the DV for a given prior probability.

The full computation for the propagation of the decision variable, $V(t)$, can be expressed by the following equations:

$$d\theta = kCdt + dW, \quad (3)$$

$$V(t) = \theta(t) + \lambda(t). \quad (4)$$

Equation 3 is a stochastic differential equation that describes drift diffusion: C is motion strength, k is the fitted drift coefficient, and W is a standard Wiener process. Equation 4 furnishes the DV as the sum of drift diffusion [the solution to Eq. 3, subject to $\theta(0) = 0$] and the dynamic bias signal (Eq. 2). The process terminates when the DV reaches one of the decision bounds, $\pm A(t)$. Note that the decision bounds are constant for the standard model; they change with time only for the alternative model with urgency described below.

We fit the choice and RT data from both neutral and unequal priors conditions with our DBS model using just four degrees of freedom (k , A , t_{nd} , and ϕ). These parameters were fit simultaneously to the full dataset. We calculated the probability of the DV crossing the upper or lower bound at each time using a numerical solution for the Fokker–Planck equation (Chang and Cooper, 1970; Press et al., 1988; Kiani and Shadlen, 2009). All fits were performed using the method of maximum likelihood on the choices and mean RTs. We compared the fits of the model with those obtained using an alternative “static bias signal” model in which $\lambda(t)$ is replaced by a static offset to one of the decision bounds. This model also has four degrees of freedom (k , A , t_{nd} , and the offset V_0). Note that this parameterization is equivalent to a model with two independent bounds. Comparisons between the DBS and static bias signal models were based on the Bayes information criterion (BIC), $-2\Lambda + k \log(n)$, where Λ is the maximized log likelihood of the model, k is the number of free parameters, and n is the number of data points. This final term is a

correction that penalizes extra free parameters, so it is constant when comparing models that have the same number of free parameters. When comparing two models, the one with a lower BIC criterion is preferable. The difference in the BIC scaled by 0.5 approximates the log of the Bayes factor, the likelihood that one model is better than another (Kass and Raftery, 1995). A Bayes factor larger than 10 indicates strong evidence in favor of a model, and a value larger than 100 is considered decisive (Jeffreys, 1961). When comparing models, we report the approximate Bayes factor. For all comparisons performed, the value corresponds to a Bayes factor that should be considered decisive.

When we altered the speed–accuracy regime (see Fig. 5), we established a new estimate of parameter A for our DBS model under neutral priors in the new regime. Importantly, we fit only the neutral priors data with this single parameter change. We used the TDA function from these data to derive $\lambda'(t)$ in this new regime. Then holding all other parameters to their original values, we predicted the choices and RTs under the unequal priors. In this way, the unequal priors data in the new speed–accuracy regime were not used in any way to determine the values of the free parameters of the model. These predictions are shown by the dashed curves in Figure 5. The fraction of the variance reported in Results (R^2) compares the predicted curves to the mean RTs and choice frequencies.

The alternative static bias signal model does not give a prescription for how much to change the offset when altering the speed–accuracy set point. Therefore, to examine how well it can explain the data in the new speed–accuracy regime, we performed an entirely new fit. To gauge the quality of this fit, we also performed new fits to the data in the new speed–accuracy regime with our DBS model (instead of the prediction-based approach described in the preceding paragraph). This provides a fair comparison because both models are allowed four degrees of freedom in each speed–accuracy regime.

We further tested whether the pattern of behavior we observed could be explained with a static bias signal model combined with a nonbiased time-dependent bound collapse. A symmetric bound collapse could implement an “urgency” signal that acts as a soft deadline for deciding (Churchland et al., 2008; Cisek et al., 2009). The time-dependent bound collapse was modeled as a hyperbolic function, so that the magnitude of the bound (before addition of the bias signal) was described by the following equation: $|A(t)| = a - u_{\infty} \frac{t}{t + \tau_{1/2}}$. Other parameterizations

may work as well, but this particular choice has been found to be reasonable for monkeys performing the motion discrimination task (Churchland et al., 2008). Because the bound collapse brings both bounds closer to zero in a symmetric manner, it alone cannot account for the asymmetric effects of prior probability on choice and RT. Nevertheless, we wanted to test whether such an elaboration could improve the fit of a static bias signal model to the behavioral data compared with our DBS model. We compared our DBS model with collapsing bounds to the static bias signal model with collapsing bounds using the BIC, as described above.

Neural recordings. Fifty-two neurons were recorded in the lateral intraparietal area (LIP) of two rhesus monkeys while they performed the reaction time motion discrimination task. These were two of the four monkeys that were trained to perform the behavioral task. We used standard methods for extracellular recording of action potentials from single neurons (Roitman and Shadlen, 2002). Neurons were selected using anatomical and physiological criteria. Stereotaxic coordinates combined with structural magnetic resonance imaging (MRI) scans were used to identify LIP and to direct the placement of recording electrodes. We believe the majority of neurons recorded in this study were located in the ventral portion of LIP (Lewis and Van Essen, 2000). This assessment is based on (1) registration of recording locations to each monkey’s MRI and comparison to a standard flat map [Caret software (Van Essen et al., 2001)] and (2) the response properties of adjacent regions.

Once the proper anatomic location was identified, neurons were selected if they exhibited spatially selective persistent activity during memory-guided saccades (Hikosaka and Wurtz, 1983; Gnadt and Andersen, 1988). In this screening task, monkeys fixated a central fixation point while a target was flashed briefly (200 ms) in the periphery and then

extinguished. This was followed by a delay of 500–2000 ms during which the monkey was required to remember the location of the flashed target. After this delay, the fixation point was extinguished and the monkey was required to make a saccade to within 3° of the location of the flashed target. By varying the location of the target from trial to trial, we identified the spatial location that caused the largest response in the neuron during the memory period; this spatial location is termed the response field (RF). Once the RF was determined, we performed repeated trials of the memory saccade task in which the target was placed at one of two locations, one inside the RF and one outside the RF. Neurons were included if the average spike rate during the memory period was significantly greater ($p < 0.05$, two-tailed t test) for inside-RF trials than for outside-RF trials.

The stability and selectivity of the neuron was reassessed throughout the recording session using two tasks: the memory saccade task and a visual saccade task that was identical to the memory saccade task except that the target remained illuminated throughout the trial. Visual or memory saccade trials were randomly interleaved with trials of the motion discrimination task. If a neuron appeared to have changed or lost its spatial selectivity or memory activity, the experiment was discontinued.

Neural analyses. Peristimulus time histograms (PSTHs) were generated by combining responses across all cells using 10 ms bins. PSTHs were aligned to two different events: motion onset and saccade initiation. When aligning to motion onset, the time period within 100 ms of saccade initiation was excluded to avoid contamination of the averages with perisaccadic activity. Similarly, when aligning to the saccade, the time period within 210 ms of motion onset was excluded to avoid contamination from the dip in firing rate that follows motion onset. Trials were grouped by direction and strength of motion with strong motion consisting of 51.2% coh, medium motion consisting of 6.4, 12.8, and 25.6% coh, and weak motion consisting of 0 and 3.2% coh. For physiological analyses, T_{in} motion refers to the direction associated with the target in the RF of the recorded cell. T_{out} motion refers to the opposite direction.

To measure a bias signal from the neural recordings, we compared responses on the neutral and unequal priors blocks. For all neurons, data were collected in the neutral prior probability condition. However, for many neurons, data were collected in just one of the two possible unequal prior probability conditions: either the prior favoring T_{in} or the prior favoring T_{out} . We estimated both a static and dynamic component of the neural bias signal. The static component is that part of the bias signal that does not depend on elapsed decision time. To analyze this, we examined whether there was a change in the average excursion of the LIP responses, that is, the difference between the level of activity at the start of motion integration and at the coalescence of responses at the end of decision formation. This difference was estimated directly from mean firing rates ~ 200 ms after motion onset and ~ 60 ms before saccade initiation for T_{in} choices. These times correspond to the neural signatures of these processes for this dataset (see Fig. 6) and are consistent with previous studies (Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Churchland et al., 2008).

We used three methods to analyze the dynamic component of the neural bias signal, which corresponds to the part that depends on elapsed decision time. All three are based on estimates of the effect of prior probability on the buildup rate of the neural responses, and all begin with estimates of a buildup rate on a trial-by-trial basis. By estimating a buildup rate for each trial individually, we are able to sidestep an “attrition” artifact that arises when estimating firing rates from averages over many trials. Because RT is broadly distributed even for the same motion strength and prior condition, such averages comprise different trials at different times. For example, as time elapses and some trials complete, attrition of these trials from the average tends to bias the average rate to lower values. Our single-trial approach avoids this bias by forming estimates of the buildup rate from individual responses and then averaging these.

For each trial, we estimated the firing rate function by convolving the spike train with an α -like function: $(1 - e^{-(t/g)})e^{-(t/d)}$, where $d = 20$ ms and $g = 1$ ms. We extracted an estimate of buildup rate for each trial by

applying linear regression to these smoothed rate functions in the epoch from 200 ms after motion onset to 60 ms before saccade initiation for T_{out} choices and 100 ms before saccade initiation for T_{in} choices. This time period was chosen based on the epoch of decision formation in which LIP reflects the integration of motion evidence (coherence and direction). Our data would place the end of this epoch ~ 60 ms before saccade initiation, consistent with previous studies (Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Churchland et al., 2008). However, to mitigate concern about including any part of a possible stereotyped perisaccadic burst in the estimates of buildup rate, we excluded a full 100 ms before the saccade for T_{in} choices.

The first analysis is a simple test of the hypothesis that a dynamic bias signal is present under unequal priors, without stipulating its exact form. For every neuron, we calculated the average buildup rate for each motion strength, direction, and prior probability condition. To analyze the influence of prior probability on buildup rate, we compared these values averaged across the population of cells for each motion strength and direction (see Fig. 7*a,b*). When comparing the neutral with each unequal priors condition, we included only the cells with data collected in the corresponding unequal priors condition ($n = 37$ neurons for neutral compared with prior favoring T_{in} ; $n = 38$ neurons for neutral compared with prior T_{out}). In addition, each cell was required to have at least five trials in the particular condition of interest (i.e., direction and strength of motion) to be included in the average for that condition. No cells were entirely excluded based on that requirement. Statistical hypotheses concerning the effects of coherence and priors on buildup rates were evaluated using two-way ANOVA, with coherence and priors as the two factors. We also included an interaction term to determine whether the effect of priors on the buildup rate depended on motion strength. Below, we describe more stringent statistical analyses of the full neurally derived bias signals.

The remaining two analyses achieve a piecewise linear approximation to the first derivative of the time-dependent bias signal. The first approach borrows from the analysis described in the preceding paragraph. In particular, we estimated the slopes from individual trials, but we restricted this to a 150 ms sliding window. The slope was estimated every 25 ms for all trials that contained the full sliding window for that particular time point. For this analysis, we used the 0% coherence motion condition, because it afforded the longest RT. Similar trends were observed for other motion strengths. We calculated the difference in the average slope measured in this way between the neutral and unequal priors trials for the same time points (see Fig. 7*c,d*). The difference provides a piecewise approximation of the rate of change of the neural bias signal.

The second approach is based on the buildup rates calculated for each motion strength (see Fig. 7*a,b*). Because the different motion strengths are associated with different RTs, we can use this as leverage to estimate the time dependence of the neural bias signal. For each motion strength, we calculated the difference in the average buildup rate under neutral and unequal priors using only neurons that were recorded in both conditions. This was done separately for the unequal prior that favored the T_{in} choice and the unequal prior that favored the T_{out} choice. This difference in buildup rate caused by the prior at each motion strength can be viewed as a linear approximation to the rate of change of the bias signal for the RTs associated with that motion strength. To combine these linear approximations at each time during decision formation, we calculated a weighted average based on the percentage of trials that had not yet terminated at that time for each motion strength. This provides an approximation of the rate of change of the neural bias signal. The integral of this rate of change converts it into an estimate of the dynamic component of the bias signal itself, in units of spike rate. Adding the static component described above yields a neurally derived estimate of the full bias signal (see Fig. 7*e,f*).

We tested the statistical significance of this bias signal using a nonparametric bootstrap analysis. We determined the sampling distribution of the neurally derived bias signal by calculating it in the same way as described above while resampling with replacement from the set of trials in the neutral and unequal priors conditions. The number of resampled trials was selected to match the number of trials used in each condition for the non-bootstrap analysis. This procedure was repeated 1000 times

to generate an estimate of the sampling distribution. We could then test whether the bias signal was significant based on examining the 95% confidence interval (CI) determined from the bootstrap method. We used the same analysis to confirm that the neurally derived bias signal remained significant when controlling for a possible confound of subjects' choices. For this, we derived an estimate of the neural bias signal from T_{in} and T_{out} choices separately and combined the estimated bias signals from each. Using the same bootstrap procedure, we found that this signal was also significant.

We also examined the effect of prior probability on the neural response in the 200 ms before motion onset. In addition to determining the effect on mean firing rate, we tested whether variations in the responses in this epoch were correlated on a trial-by-trial basis with responses at the start of decision formation (~ 200 ms after motion onset). For each neuron, we standardized the responses for each prior probability condition in the analyzed epochs. We then tested whether the standardized responses between the two epochs were correlated, using Pearson's correlation coefficient.

To predict the change in behavior from the neurally derived bias signal (see Fig. 8*b*), we used the k , A , and t_{nd} parameters of the bounded accumulation model fit to the behavior data obtained under neutral priors. However, instead of using the dynamic bias signal of this model [$\lambda(t)$], we substituted the neurally derived bias signal. As described above, this signal consists of the static component of the neural bias signal—that is, the measured change in LIP excursion—and the dynamic component of the neural bias signal derived from the change in buildup rates for each motion strength. Because we have an estimated bias signal for both the high and low unequal priors conditions, we averaged the magnitude (i.e., the absolute value) of these signals to make the behavioral prediction. To map this to the model, we expressed the neurally derived bias signal as a percentage change of the bound height. This was calculated using the ratio of estimated neural bias signal to the measured LIP excursion in the neutral priors condition. The measured LIP excursion is the difference in firing rate at the start and end of the decision process—effectively, the neural bound.

We used logistic regression to test for a possible trial-by-trial association between the neural signatures of the dynamic bias signal and the monkeys' choices. The logistic model is as follows:

$$P_+ = \{1 + \exp[-(\beta_0 + \beta_1 C + \beta_2 \text{sgn}(L) + \beta_3 z_0 + \beta_4 z_{\text{bu}})]\}^{-1}, \quad (5)$$

where P_+ is the probability of a T_{in} choice, C is the signed motion strength, L is the log of the prior odds [$\text{sgn}(L)$ is its sign: -1 , 0 , or $+1$], and the z terms are transformed neural responses computed from spike rate at the beginning of decision formation (z_0) and buildup rate during decision formation (z_{bu}). Because the buildup rate is affected by the motion strength, we calculated its z -scores within signed motion strength conditions in addition to within cell and prior probability condition. The β_i are fitted coefficients (maximum likelihood). Equation 5 permits tests of whether variation in the neural parameters (z) exert an influence on choice in the presence of the explanatory variables (C and L). We also performed an expanded version of the regression that included an additional z -score term based on the firing rate in the 200 ms epoch before motion onset.

Results

We trained two humans and four monkeys (*Macaca mulatta*) to perform a reaction time version of a two-choice direction discrimination task (Fig. 1). Task difficulty was controlled by manipulating the percentage of coherently moving dots on a trial-by-trial basis. Because a range of motion strengths were randomly interleaved, subjects lacked knowledge of the reliability of evidence they would receive on each trial. Subjects were instructed to determine the net direction of coherent motion and to indicate a choice by making an eye movement to the corresponding target. The motion stimulus remained visible until saccade initiation. We manipulated the prior probability of the direc-

tion of motion in blocks of trials (range, 200–1000 trials per block).

We begin by describing the effects of prior probability on subjects' behavior. Model fits to the subjects' choices and RTs support the theory that the influence of prior probability depends on elapsed time during decision formation. We then present two additional tests of the predictions of this theory. The first is that the effect of prior probability depends on the speed–accuracy regime that a subject applies. The second is that the effect of prior probability on neurons involved in decision making depends on elapsed time during decision formation. We tested the first prediction in human subjects by altering their speed–accuracy tradeoff. We tested the second prediction in rhesus monkeys by measuring neural activity from area LIP, an area shown previously to represent prior probability in a saccade task (Platt and Glimcher, 1999) and to represent a DV for the motion discrimination task (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Hanks et al., 2006; Gold and Shadlen, 2007; Kiani and Shadlen, 2009; Rorie et al., 2010).

The behavioral effects of prior probability

For all subjects and under all conditions, the strength and direction of stimulus motion governed the pattern of choices and RTs. Here and throughout, we refer to the strength of motion as a percentage of coherently moving dots and use the sign to indicate direction. When the prior probability was equal for the two directions of motion, stronger positive motion led to a higher proportion of positive choices, whereas stronger negative motion led to a lower proportion of positive choices (Fig. 2, filled symbols). Stronger motion in either direction also led to faster RTs (Fig. 2, filled symbols). These observations are consistent with previous studies using this task in monkey and human subjects (Roitman and Shadlen, 2002; Palmer et al., 2005; Cohen and Newsome, 2009).

When the prior probability favored one direction over the other, there was a shift in the choice and RT functions (Fig. 2, open symbols). By convention, we defined the positive direction of motion to be the one favored by the prior. Thus, at a given stimulus strength, subjects made more choices corresponding to the positive direction of motion. In addition, RTs decreased for these choices, and RTs increased for the choices corresponding to the negative direction of motion.

We applied a model of bounded evidence accumulation to understand the pattern of choices and RT (Fig. 3*a*). This framework explains a large range of psychophysical experiments in both humans and monkeys (Laming, 1968; Link, 1992; Ratcliff and Rouder, 1998; Gold and Shadlen, 2002; Ratcliff and Smith, 2004; Palmer et al., 2005; Bogacz et al., 2006; Kiani et al., 2008). In the model, noisy sensory evidence is accumulated over time into a DV that undergoes diffusion with drift until reaching an upper or lower bound that terminates the process and establishes the choice. The drift rate is proportional to motion strength, thus biasing the diffusion toward the positive or negative bound. In

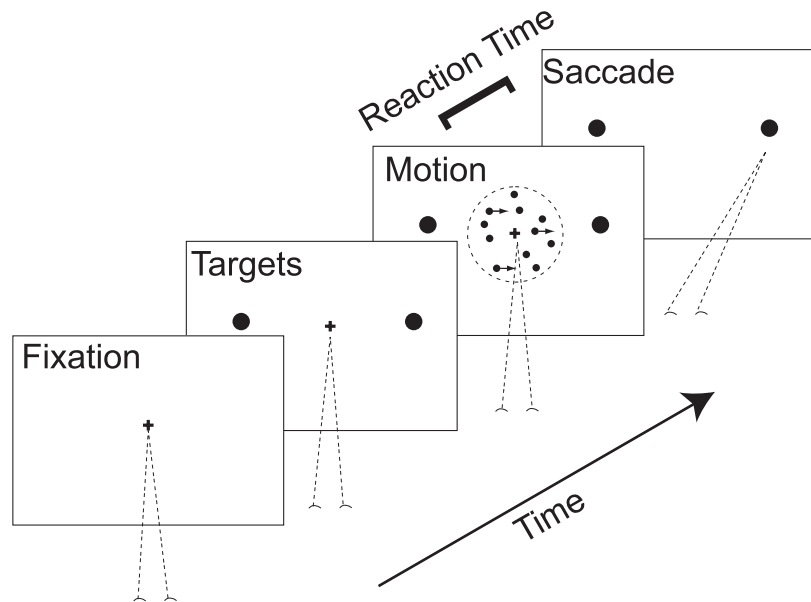


Figure 1. Task design. The subjects discriminated between two possible directions of motion in dynamic random dot displays. The subjects could respond at any time after onset of the random dot motion by making a saccadic eye movement to one of the choice targets. The elapsed time from motion onset to saccade initiation is the reaction time. The proportion of coherently moving random dots determines the motion strength and hence the reliability of the evidence bearing on the direction decision. This was fixed for a given trial but drawn randomly across different trials from a set of motion strengths spanning a range of difficulties. In different blocks of trials, the prior probability of the direction of motion was either neutral (50:50) or favored one of the directions (80:20). For neural recordings in the monkeys, one choice target was placed in the response field of a single LIP neuron, and the other target was placed in the opposite hemifield.

its simplest form, considered here, the model involves just three parameters (see Materials and Methods). When the prior probability was neutral—that is, equal for the two directions of motion—this simple model explained the choices and mean RTs (Fig. 2, solid curves), consistent with previous studies (Mazurek et al., 2003; Palmer et al., 2005).

When the prior probability favors one direction, how does this knowledge affect the decision process? As shown by the open symbols in Figure 2, it causes a shift in the RT and choice functions. Our goal is to understand why the shift has the particular magnitude observed and what the pattern of changes in both the RT and choice functions tells us about the mechanism. In particular, we aim to identify a signal that is added to the DV to cause this bias in behavior. To develop our argument, we need to share an intuition about why a bias signal added to the DV might change as a function of decision time although the prior probability is fixed throughout the duration of the decision—indeed, throughout the block of trials.

A dynamic model of bias is motivated by the fact that, under neutral priors, accuracy decreases as a function of decision time across the ensemble of stimuli presented (Fig. 4*a,b*). This is a straightforward consequence of the longer RTs associated with more difficult trials. As can be appreciated from the choice and RT functions, weaker stimuli, which yield lower accuracy, also result in longer RTs on average. Because this full relationship, termed the time-dependent accuracy (TDA) function, describes how accuracy changes as a function of decision time, it links elapsed time during the course of a decision to the average reliability of the sensory evidence gathered up to that point. When the criterion or bound on the DV for choice commitment remains unchanged between weaker and stronger stimuli, a difference in accuracy implies different levels of evidence reliability for the same magnitude of the DV. In particular, it shows that the

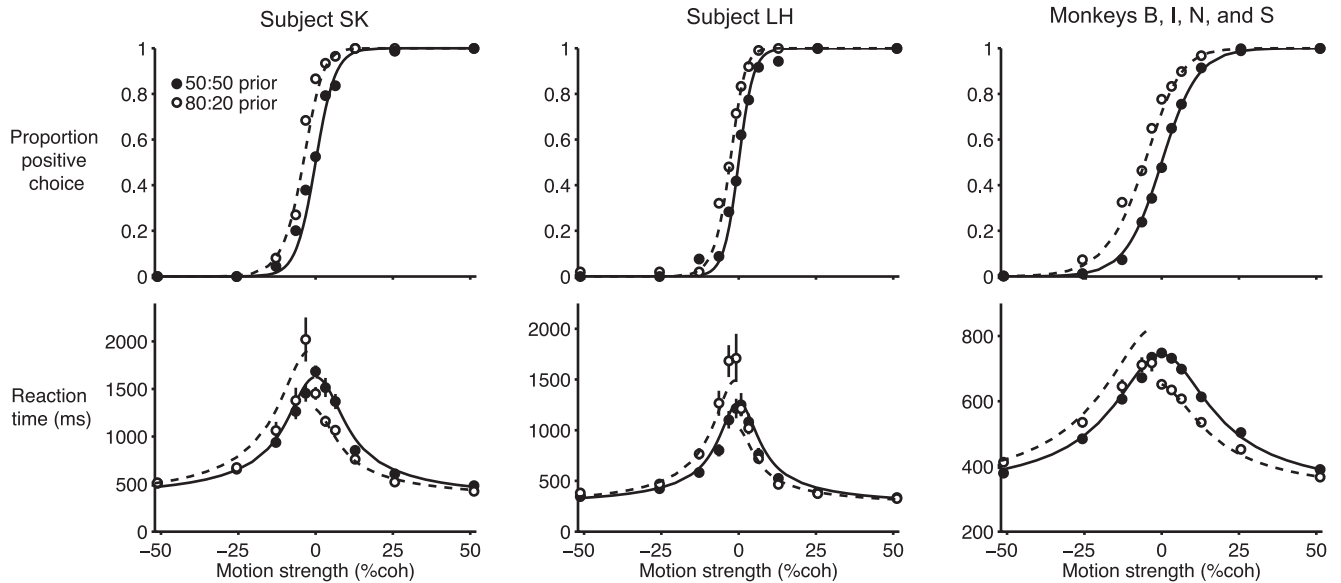


Figure 2. A DBS model explains choices and RTs when the prior probability of motion direction is altered. The top row shows the effect of motion strength and prior probability on choices. The proportion of positive choices is plotted as a function of motion strength. The sign of motion strength connotes direction. For the blocks with unequal prior, the positive direction is the one favored by the 80:20 or 20:80 prior (see Materials and Methods). Filled and open symbols depict the choice proportions for the neutral and unequal prior conditions, respectively. The bottom row shows the effect of motion strength and prior probability on mean reaction time. Filled and open symbols follow the same convention as the choice functions. Error bars are SEM; many are smaller than the symbols. For all the choice and RT functions, the solid and dashed curves show the DBS model fits to the data in the neutral and unequal priors conditions, respectively. The left two columns show data from individual human subjects (2400 and 2785 trials, respectively). The right column shows aggregate data from four rhesus monkeys (36,687 trials).




Figure 3. Bounded accumulation of evidence model. Momentary evidence for one of the alternatives and against the other is distributed as a unit-variance Gaussian with mean proportional to motion strength. This momentary evidence is accumulated over time. The accumulated evidence, termed the decision variable, resembles a random walk or drift-diffusion process. We have shown a smoothed version of the DV for this schematic. The process terminates with a positive or negative choice when the DV reaches the upper or lower bound, respectively, at $+A$ or $-A$. In the motion discrimination task, we believe the momentary evidence is represented predominantly in MT and other motion-selective areas, whereas LIP carries a representation of the accumulated evidence that forms the DV. **a**, Neutral priors. The DV begins midway between the bounds, and the DV is biased only by the momentary evidence, which induces a drift proportional to the motion coherence. **b**, Unequal priors. Prior probability adds an additional bias to the DV. This added bias signal is not known. Here, it is depicted as an increasing function of decision time. In the example, it causes an earlier decision for the positive direction.

same magnitude DV corresponds on average to a greater level of evidence reliability at shorter elapsed decision times and a lower level of evidence reliability at longer elapsed decision times (Kiani and Shadlen, 2009). To exploit this temporal structure, subjects could incorporate prior probability into the decision process as a dynamic bias signal added to the DV that depends on the TDA function (Fig. 3b).

We computed the dynamic bias signal as the probabilistic combination of the prior probability with the odds of a correct choice. These odds are implied by the TDA function for a given

elapsed decision time. The prior adds an amount to the DV that corresponds to a fraction of the log odds of a correct choice at that given decision time (see Materials and Methods). Although theoretically this alone should provide a prescription for the dynamic bias signal, we also introduced an additional parameter that scales the prior probability to account for the possibility that a subject fails to adequately acquire or incorporate this knowledge perfectly. Later, we will discuss the use of this parameter. We found that this model describes the pattern of behavior observed when prior probability was manipulated (Fig. 2, dashed curves).




Figure 4. TDA functions. The probability of a correct choice is plotted as a function of reaction time for the neutral priors condition. Graphs are from two human subjects who were tested in two different speed–accuracy regimes. The solid and dashed traces show the higher accuracy and faster speed regimes, respectively. Faster speed led to reduced accuracy at all decision times and a change in the shape of the TDA functions.

Moreover, for both human subjects, the model assigned a bias estimate that was within 85% of the actual prior ($\phi = 0.88$ and 0.85 for subjects SK and LH, respectively); for all four monkeys, the model assigned a level of bias within 75% of the actual prior (mean $\phi = 0.80 \pm 0.01$).

We also performed fits to the data that allowed a bias signal to be added as a static offset with the magnitude of that offset as a free parameter. This is the standard assumption for the incorporation of prior probability in a decision (Edwards, 1965; Link and Heath, 1975; Carpenter and Williams, 1995; Gold et al., 2008; Ratcliff and McKoon, 2008; Simen et al., 2009). This static bias signal model can be compared directly with our DBS model because the models have the same number of free parameters (see Materials and Methods). The DBS model provided a better fit to the data than the static bias signal model in all cases (BIC, Bayes factor $>10^8$ for both human subjects and $>10^6$ for all four monkeys; see Materials and Methods). We take this as evidence in favor of the idea that the brain exploits elapsed time to combine probabilistic representations.

To some readers, it might seem counterintuitive to use a dynamic signal to represent a stationary prior. One possibility we have yet to consider is that a stationary signal representing prior probability is added to another dynamic signal that has nothing to do with bias. For example, it has been suggested that the terminating bounds are themselves dynamic, collapsing symmetrically, to cut short less informative trials (Ditterich, 2006). Such collapsing bounds [equivalent to a time-dependent urgency signal added to competing accumulators in (Churchland et al., 2008)] would not induce a choice bias, but perhaps it would allow

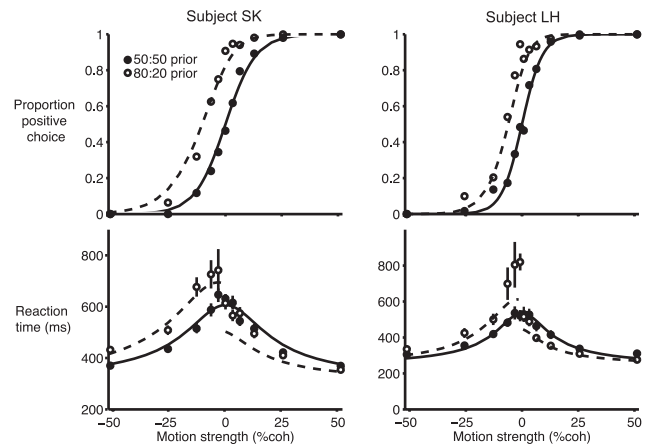


Figure 5. The DBS model explains the effects of prior probability in different speed–accuracy regimes. Symbols show the observed choice proportions and mean RT in a faster speed regime for two human subjects (2400 and 3125 trials). Data from blocks with neutral and unequal priors are shown by the filled and open symbols, respectively. Same plotting conventions as Figure 2. Solid curves are model fits to the data in the neutral prior probability condition. All parameters except the decision bound were fixed to the values used to fit data in the higher-accuracy regime. So, there is just one degree of freedom in the fit. Dashed curves are predicted choice and accuracy functions for the unequal prior condition (data shown by open symbols). These theoretical predictions are obtained from the DBS model with parameters established by fits to the data from the high-accuracy regime and the neutral priors data from the faster-speed regime. For details, see Results.

a static bias signal to explain the data. We therefore compared our DBS model with one with a static bias signal. Both models were allowed two additional degrees of freedom to incorporate symmetric collapsing bounds (see Materials and Methods). We found that, when bounds were allowed to collapse, our DBS model remained superior to the static bias signal model for both the humans and the monkeys (BIC, Bayes factors $>10^6$ in all cases; see Materials and Methods).

Based on these observations, we conclude that subjects were aware of the change of prior probability and that they incorporated this information into the decision by adding (or subtracting) a time-varying quantity to the decision variable. The added signal appears to be indexed to the log odds of a correct choice as a function of decision time under neutral priors.

A psychophysical test: change of the speed–accuracy regime

Because our DBS model depends critically on the relationship between evidence reliability and elapsed decision time that is implied by the TDA function, altering this relationship provides a strong test of the theory. This was accomplished by instructing our human subjects to perform the direction discrimination task in a faster speed–accuracy regime. After several practice sessions using neutral priors, both subjects achieved a new set point in both speed and accuracy (Fig. 5, filled symbols). Decisions were both faster and less accurate, on average, compared with the slower regime. The reduction in accuracy suggests that decisions made under the speed instruction were based on less accumulated evidence. Indeed, this tradeoff between speed and accuracy was consistent with a change in the decision bound combined with the same model parameters that were derived from the slower speed–accuracy regime (Fig. 5, solid curves), consistent with previous studies (Reddi et al., 2003; Palmer et al., 2005). Importantly, trials with matching RTs in the two speed–accuracy regimes were associated with different levels of accuracy, as shown by the TDA functions in Figure 4 (both under neutral priors). The change in the TDA reflects a combination of

Table 1. Behavioral bias measured for the choice behavior of both human subjects for each speed–accuracy regime

Subject	High-accuracy regime		Fast-speed regime	
	Bias	CI	Bias	CI
SK	5.1% coh	3.1–6.3%	10.8% coh	9.1–12.6%
LH	4.1% coh	3.1–5.0%	8.3% coh	6.9–9.8%

The bias was in favor of the choice associated with 80% prior probability.

factors: a change in the level of accumulated evidence required to terminate decisions and the different mixture of stimulus strengths contributing to each RT in the two regimes. The new TDA, ascertained under neutral priors, leads to a new prescription for incorporation of unequal priors in the decision process.

According to our hypothesis, the new relationship between elapsed decision time and accuracy ought to affect the dynamic incorporation of prior information into the decision process. The change in decision bound invoked to explain the neutral priors data in the high-speed regime is symmetric for the two choices, so it causes no bias itself. However, it induces a change in the mapping between accuracy and the DV (i.e., accumulated evidence) as a function of elapsed time. In the high-speed regime, the prior should exert more leverage on the decision because the accumulated evidence is less reliable at all times. The prior does not change, of course. It is the evidence, in this new regime, that supports a smaller probability of a correct answer. The TDA function thus establishes the shape and magnitude of the dynamic bias signal that should be added to the accumulated evidence as a function of elapsed decision time. This leads to a quantitative prediction for choice and RT under unequal priors (Fig. 5, dashed curves).

The observed pattern of choices and RT for the two human observers supports these predictions. Similar to the other speed–accuracy regime, there was a shift in the choice and RT functions when the prior probability favored the positive direction of motion (Fig. 5, open symbols). However, the effect of the prior on choices was greater in the high-speed regime. These differences were statistically reliable for both subjects (Table 1). Most importantly for our theory, the magnitude and pattern of changes in choice proportions and RTs conformed reasonably well to the predictions of the theory (subject 1, $R^2 = 0.80$; subject 2, $R^2 = 0.78$). Note that the dashed curves in Figure 5 are not fits to the data (open symbols); instead, the prediction comes entirely from model parameters determined by the high-accuracy fits and the bound change associated with the neutral priors data for the high-speed regime.

Although we have already compared our DBS model with a static bias signal model for the high-accuracy regime, it is also illuminating to perform this comparison using data from the high-speed regime. It is important to note that the standard static bias signal model does not provide a prescription for how much to change the bias offset in different speed–accuracy regimes. So, from the start it does not have the same explanatory power as our DBS model. We therefore compared models by allowing each to fit the data from the high-speed regime with all four parameters in each model free. This is the same analysis we described above for the high-accuracy data. For both human subjects, our DBS model provided a better fit to the data than a static bias signal model (BIC, Bayes factors $>10^{15}$; see Materials and Methods). Following similar logic as described for the high-accuracy fits, we also compared models with symmetric (unbiased) time-dependent bound changes. Again, our DBS model remained superior to a static bias signal model (BIC, Bayes factors $>10^{15}$).

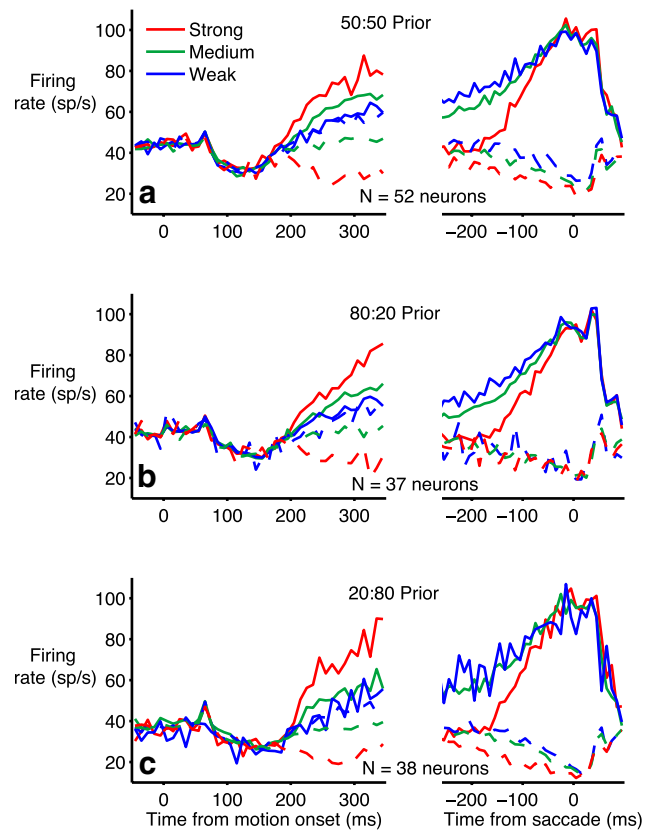


Figure 6. A neural correlate of a decision variable in LIP. Average firing rates of LIP neurons are shown as a function of time during decisions for neutral (**a**) and unequal (**b**, **c**) priors. Each neuron was recorded in the neutral and one or both of the unequal priors conditions. One of the choice targets (T_{in}) was in the response field of the neuron, and the direction of coherent motion either supported this or the opposite choice (T_{out}). On the left side of the graphs, response averages are aligned to the onset of random dot motion. On the right side of the graphs, response averages are aligned to saccade initiation. Trials are grouped by motion strength (color) and direction (solid and dashed for motion toward T_{in} and T_{out} , respectively) with only correct choices contributing to these averages. Response averages are calculated in 10 ms bins. For all conditions, motion evidence (strength and direction) affects the response averages from ~ 210 ms after motion onset until ~ 60 ms before saccade initiation. This defines the epoch in which decision formation is reflected in the LIP firing rate. **a**, Neutral priors block. **b**, Prior probability favors the direction rewarded by choosing T_{in} (80:20 prior). **c**, Prior probability favors the direction rewarded by choosing T_{out} (20:80 prior).

Measurement of the bias signal in neurons

A neural correlate of a DV has been demonstrated previously in LIP of monkeys trained to indicate their direction decision with a saccadic eye movement (Roitman and Shadlen, 2002). We therefore tried to discern the changes in this neural correlate when the prior probability favored one of the directions over the other. We recorded from 52 LIP neurons in two of the monkeys whose behavioral results were described above.

On the blocks of trials using the neutral prior, we observed the pattern of changes in LIP firing rate similar to previous reports (Fig. 6a) (Roitman and Shadlen, 2002; Churchland et al., 2008). Just after motion onset, there was a dip in activity, followed by a rise or fall in the responses that reflected the direction and strength of motion as well as the decision for T_{in} or T_{out} . This decision-related activity is apparent from ~ 200 ms after motion onset. As shown previously, this is the latency from a change in stimulus motion to a change in the response of LIP neurons whose RF overlaps the choice target (Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Kiani et al., 2008). It is much longer than the latency of a simple visual response to a light in the RF (Bisley

et al., 2004). We are interested in the neural responses accompanying evidence accumulation, beginning at this time and ending ~ 60 ms before saccade initiation. For T_{in} choices, this decision end point is marked by a common level of firing rate, regardless of motion strength and RT (Roitman and Shadlen, 2002). These features, which are consistent with the bounded accumulation mechanism, were also apparent in blocks of trials using unequal priors (Fig. 6*b*).

To ascertain a bias signal from the neural recordings, we compared responses on the neutral and unequal priors blocks. We sought evidence for two aspects of a bias signal: (1) a change in the average excursion of firing rate from beginning to end of the decision and (2) a change in buildup rates during decision formation. Note that a change in excursion is mathematically equivalent to a starting point offset and also to a static component of the bias signal. The higher prior caused a subtle decrease of 3.7 [95% CI, 0.1–7.3 spikes per second (sp/s)] in the excursion, which was borderline significant ($p < 0.05$). Based on the excursion of LIP firing rate measured in the neutral prior condition, this static change is equivalent to a reduction of that excursion by $\sim 9\%$. The lower prior caused a larger but less reliable increase in the excursion (5.6 sp/s; 95% CI, -1.6 to 12.8 sp/s; $p = 0.13$).

Interestingly, the level of activity at the beginning of the decision appears to be established in part by the level of firing rate before the motion stimulus was shown. The firing rate in the epoch 200 ms preceding motion was 4.6 ± 1.1 sp/s greater when the prior favored T_{in} than when it favored T_{out} ($p < 0.01$), although only the latter was significantly different from the firing rate under neutral priors ($p < 0.01$ for neutral prior vs prior favoring T_{out} ; $p = 0.42$ for neutral prior vs prior favoring T_{in}). In addition, variation in the firing rate in this epoch was correlated on a trial-to-trial basis with the firing rate at the beginning of decision formation ($r = 0.26$; $p < 0.001$; data not shown; see Materials and Methods). This is consistent with previous findings that have demonstrated a neural correlate of prior probability in LIP when a monkey is simply presented two choice targets, before a decision is formed (Platt and Glimcher, 1999).

More importantly for ascertaining a neural signature of the dynamic bias signal predicted by our theory, we also examined whether there was a change in the rate of response buildup in this epoch. This is the slope of the coherence- and time-dependent ramping activity evident in Figure 6. A change in buildup rate would effectively implement a dynamic component of the bias signal because the influence of the prior on spike rate would change as time elapses during the trial.

The higher prior caused an increase in the buildup rate (Fig. 7*a*) ($p < 0.01$, ANOVA). This implies that the change in firing rate attributable to the prior is not constant but increases as a function of time; in other words, there is an increasing bias signal. The lower prior (i.e., bias against the T_{in} direction) caused a decrease in the buildup rate (Fig. 7*b*) ($p < 0.01$), although the effect is not as pronounced as for the high prior. This regression-based analysis of buildup rates establishes a crude neural correlate of prior probability, without recourse to the theory we are advancing. We note, however, that the change in buildup rate caused by the prior was more prominent at lower motion strengths ($p < 0.01$). This observation is consistent with a time-dependent increase in the magnitude of the bias signal because the decisions with longer duration tend to be those made on the lower coherence trials.

To achieve a better sense of the time dependence of this effect, we used a sliding window to estimate the difference in buildup rates between neutral and unequal priors trials as a function of

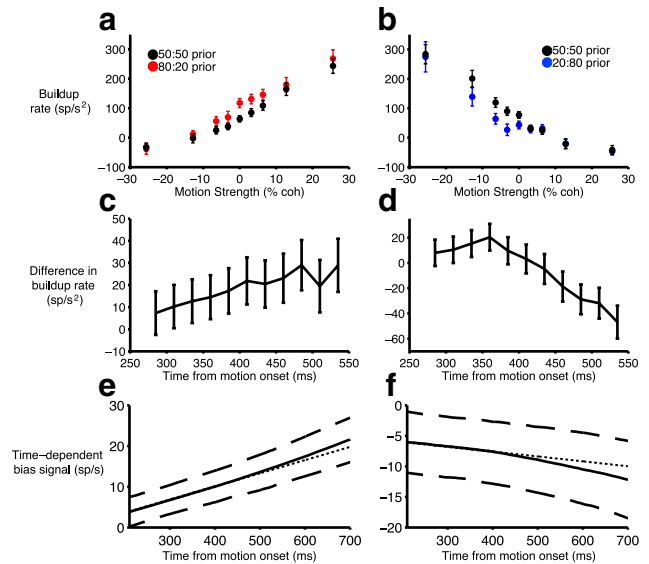


Figure 7. LIP responses exhibit a dynamic bias signal when prior probability is unequal. *a*, Effect of motion strength and prior probability in favor of T_{in} on buildup rates. The buildup rate is the slope of the firing rate versus time function from the time corresponding to the start of motion integration (see Materials and Methods). Positive motion strengths correspond to motion in the T_{in} direction favored by the prior (i.e., toward). For each neuron, the buildup rate was calculated at each motion strength, direction, and prior condition. The points are averages \pm SE across the population of neurons tested in both blocks ($n = 37$ neurons). Black and red symbols are for neutral and 80:20 prior, respectively. The prior caused an increase in the buildup rates, mainly for the weaker motion strengths. *b*, Effect of motion strength and prior probability in favor of T_{out} on buildup rates. Positive motion strengths correspond to motion in the direction favored by the prior (i.e., toward T_{out}). Same conventions as in *a* ($n = 38$ neurons studied in both blocks). The prior caused a decrease in the buildup rates, mainly for the weaker motion strengths. The strongest coherences (51.2%) are not shown in *a* and *b*; there were no reliable effects of priors on buildup rates at that motion strength. *c*, *d*, Time-dependent effect of prior probability on buildup rate. This analysis uses only the trials with 0% coherence motion. The buildup rate was estimated in a sliding time window (width, 150 ms; time on abscissa indicated window midpoint). The difference in buildup rates for neutral and unequal prior are shown as a function of time (window starting points separated by 25 ms). The graph begins at the start of the motion-dependent buildup, and the graph stops when there are less than half of trials remaining in the averages. *c*, When the prior favored T_{in} , there was an increase in the buildup rate as a function of time compared with the neutral prior. *d*, When the prior favored T_{out} , there was a decrease in the buildup rate as a function of time compared with the neutral prior. *e*, *f*, Neurally derived estimate of the bias signal. This is based on a weighted average of the effect of priors on the buildup rates for each motion strength, as shown in *a* and *b*. The weight of each motion strength is based on the percentage of trials that have not yet terminated, as a function of time. The integral of this piecewise rate of change function provides an estimate of the dynamic component of the neural bias signal. The neurally derived bias signal (solid curve) is the sum of this dynamic component and a static component estimated from the change in LIP excursion (see Materials and Methods). The thin gray line shows a linear bias signal extrapolated from the unweighted average of the effect of priors on the buildup rates for each motion strength. The dashed lines show the boundaries of the 95% CI on the neurally derived bias signal (bootstrap). *e*, When the prior favored T_{in} , the estimated bias signal is positive and increasing as a function of time. *f*, When the prior favored T_{out} , the estimated bias signal is negative and decreasing.

time during the trial (see Materials and Methods). We focused on the 0% coherence motion strength for this analysis because these trials have the longest decision times. The results of this analysis are depicted in Figure 7, *c* and *d*. When the prior favored T_{in} , the difference in buildup rates was more prominent as time elapsed during decision formation. This indicates that the bias signal is not just increasing but accelerating as a function of elapsed decision time. When the prior favored T_{out} , as expected the difference in buildup rates was of opposite sign, and the magnitude of the difference in buildup rates also increased at later times during the

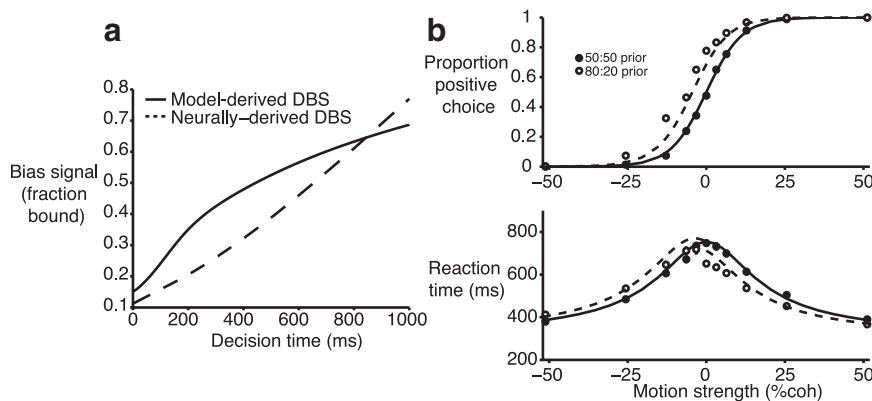


Figure 8. The neurally derived bias signal predicts the effect of the unequal prior. *a*, Comparison of neurally derived bias signal with the one predicted from our DBS model. Both signals are plotted as a fraction of the bound height. The signals are of comparable magnitude but differ to some degree in shape. *b*, Predicted choice and RT functions derived from the neurally derived bias signal. The symbols show the same choice proportions and RTs for the monkeys as shown in Figure 2 (same conventions). The solid traces are fits of the bounded accumulation model to the data obtained in the neutral prior probability condition. Behavioral predictions for the unequal prior probability condition were made by substituting the neurally derived bias signal for the model-derived bias signal in the bounded accumulation framework (see Materials and Methods). These predictions are shown by the dashed curves.

trial. This analysis confirms that the prior probability adds (or subtracts) a time-dependent signal to the spike rate of LIP, and it provides a direct estimate of the shape of this signal.

We used a second technique to derive the dynamic component of the bias signal using the responses at all motion strengths. The difference in buildup rate caused by the prior at each motion strength is a linear approximation to the rate of change of the bias signal. Because the different motion strengths are associated with different RTs, at each time we calculated an “average” of these linear approximations, weighting them by the percentage of trials that had not yet terminated for each coherence (see Materials and Methods). This provides a piecewise function for the rate of change of the bias signal. To derive the bias signal, one simply integrates the piecewise rate of change function and adds the static offset component derived from the analysis of excursions (above; see Materials and Methods). This yields the neurally derived dynamic bias signal shown in Figures 7, *e* and *f*. To test whether this bias signal was significant, we used a bootstrap procedure to establish confidence intervals for the full bias signal and its dynamic component alone. Both were significant at all times ($p < 0.05$). A similar analysis also showed that the neurally derived bias signal remains significant when restricting the analysis to trials ending in the same saccadic choice (see Materials and Methods; $p < 0.05$).

The neurally derived bias signal resembles the bias signals proposed on theoretical grounds to explain the shifts of choice and RT functions. Both increase with elapsed decision time (Fig. 8*a*). A more informative analysis is to incorporate the neurally derived bias signal into the bounded diffusion model that was used to explain the monkeys’ choices and RTs. To achieve this, we expressed the neural bias signal as a fraction of the excursion to bound under neutral priors. We then attempted to predict the behavior under non-neutral priors using this neurally derived bias signal. This is the same strategy we used to make predictions for our human subjects in the fast speed–accuracy regime. We derive a model for the bounds and decision variable under neutral priors and attempt to predict the behavior by adding a bias signal. Here, we use the neurally derived bias signal instead of one derived from the TDA function. As shown in Figure 8*b*, the shift in choice and RT predicted from the neurally derived bias signal are in reasonable agreement with the monkey’s behavior ($R^2 =$

0.87). The dashed curves corresponding to unequal priors data are predictions, not fits. Only data from neutral priors were used for the fit (solid curves).

As a final test of our idea, we examined whether the neurally derived bias signal influenced the monkey’s choices on a trial-by-trial basis. To assess this, we used single-trial measures of the firing rate at the beginning of decision formation and the rate of response buildup during decision formation. For both measures, we standardized responses for each neuron to remove the effect of experimentally controlled variables: prior probability, stimulus motion strength, and direction (see Materials and Methods). What remains are the residual variations in initial firing rate (z_0) and the buildup rate (z_{bu}) during decision formation expressed in units of SD. Using logistic regression, we determined whether trial-by-trial variation in these measures alters the log odds of a

T_{in} choice beyond the explanatory power of the stimulus and prior probability (Eq. 3). We found that both the initial firing rate and buildup rate during decision formation had significant leverage on choices (z_0 , $\beta_3 = 0.30$, $p < 0.001$; z_{bu} , $\beta_4 = 0.97$, $p < 0.001$).

Because other studies have shown a neural correlate of choice bias in LIP before decision formation (Seidemann, 1998; Platt and Glimcher, 1999; Shadlen and Newsome, 2001; but see Gold et al., 2008), we wanted to test whether variability in the firing rate before motion onset had additional explanatory power on choices. We therefore included an additional term, $\beta_5 z_{pre}$, in the logistic regression (Eq. 5), where z_{pre} is the standardized response in the 200 ms epoch before motion onset. We found that this parameter did not have significant leverage on choices ($\beta_5 = -0.02$, $p = 0.51$), whereas the effects of z_0 and z_{bu} did not change appreciably ($\beta_3 = 0.31$, $\beta_4 = 0.97$, $p < 0.001$ for both). This suggests that any impact of the pre-motion responses on choices is mediated via the dynamic bias signal.

Discussion

We have found that human and nonhuman primates incorporate information about prior probability into a perceptual decision by adding a dynamic bias signal to the DV. By dynamic, we mean that the bias signal offsets the DV by a different amount as a function of time. This seems peculiar at first blush because prior probability is a constant function of time. The strategy is sensible, however, when the decision time depends on the reliability of the evidence and the DV is not in units of probability. In this case, the mapping between the DV and probability changes with time. A dynamic bias signal allows a decision maker to assign appropriately greater leverage to high-quality evidence while letting the prior exert more leverage on the decision when evidence is less reliable.

We found that a bounded accumulation model that incorporates this strategy describes the effect of prior probability on choice and RT in humans and monkeys. The idea provides a principled theory for how prior probability should affect choices and RTs when the mapping between decision time and the reliability of evidence is altered. We found that these predictions were satisfied when human subjects were instructed to change

their speed–accuracy set point. Finally, we found that a neural correlate of a dynamic bias signal is added to LIP responses in monkeys performing the task. The source of this signal is unknown, but it may be related to other time-dependent signals reported in LIP associated with timing and urgency (Leon and Shadlen, 2000; Janssen and Shadlen, 2005; Maimon and Assad, 2006; Churchland et al., 2008). However, unlike these dynamical signals, what is added to one accumulator (e.g., for rightward choices) must be subtracted from the other.

Our theory amends the dominant view that prior probability is incorporated into the decision process as a static change of either the DV or the decision bound (Edwards, 1965; Link and Heath, 1975; Carpenter and Williams, 1995; Platt and Glimcher, 1999; Gold et al., 2008; Ratcliff and McKoon, 2008; Simen et al., 2009). An important inspiration for a static bias signal comes from the sequential probability ratio test (SPRT) (Wald, 1947; Wald and Wolfowitz, 1947). In SPRT, a bound on the accumulated LLR as represented by the DV entails that the accuracy is constant as a function of decision time. Thus, the TDA function is flat. In that situation, the prior should exert the same influence on the DV regardless of decision time. This is a special case of the theory we propose here, and there is experimental support for it in decision tasks similar to ours using a fixed motion strength for all trials (Simen et al., 2009).

An advantage of our proposal is that probability is implicitly represented in a learned association between brain state and behavioral consequence. Unlike SPRT, in which the DV is in units of LLR, the DV of the brain is naturally expressed in units of spike rate, which arise from computations on quantities such as motion energy (Adelson and Bergen, 1985; Born and Bradley, 2005). We are suggesting that the brain forms an association between the state of the DV at a given time and the probability that it will lead to a correct choice. Absent information about reliability itself, the brain can exploit knowledge learned by the association of its own DV and the probability of a correct decision.

These are the conditions that apply in our experiments and in many other situations in which a subject makes a decision based on a stream of evidence that is presumed to arise from one source but the reliability of the source is unknown. If the subject were to know the reliability of that stream of evidence, then the decision process could exploit this information to weight the prior appropriately. For example, if the subject knows she will be shown a 0% coherence motion stimulus, an appropriate strategy is to choose the direction favored by the prior and to answer immediately.

In contrast, our strategy is useful when the reliability is unknown, what might be called uncertainty about uncertainty. When subjects face a stream of evidence with uncertain reliability, a decision rule induces an association between the state of the DV at a given decision time and the probability of a correct choice. In that case, elapsed time confers information about the reliability of the evidence (Kiani and Shadlen, 2009). A dynamic bias signal may be viewed as a method to marginalize over the distribution of evidence reliability at a given stopping time for a bounded accumulation process. Thus, our theory is consistent with a Bayesian approach under the idea that the subject's knowledge of evidence reliability comes from elapsed decision time.

The method by which the brain determines the association between the state of the DV and the probability of a correct choice remains an open question. Although we propose our idea as a step toward a normative theory, it is perhaps more appropriate to view it as a principled heuristic. Our specific proposal, as presented above, assumes that the association is formed with the expected performance for a given elapsed decision time under

neutral priors. However, the addition of a bias signal could potentially affect this computation as well. For example, prior knowledge might promote the adoption of a different speed regime or the desire for a different TDA than the one that results from our proposal. Our theory does not yield a full normative solution to this problem. Nevertheless, by incorporating the new information about prior probability into an existing, acceptable framework that permits calibration of that information, our theory provides a sensible and straightforward approach.

This theory is an advance beyond less principled heuristics. Within the context of bounded accumulation or drift-diffusion mechanisms, one might consider whether a bias changes the starting point or drift rate of the process. Although a starting point offset provides the normative solution when the conversion of the DV to units of probability remains constant for the entire trial duration, this solution is not appropriate when accuracy is time dependent. Thus, it is not surprising that starting point offset is inconsistent with the pattern of choice and RT in our subjects. A change in drift rate introduces a bias signal that changes linearly with time (Ashby, 1983; Diederich and Busemeyer, 2006). Such an increase in drift rate is a reasonable approximation to the dynamic component of the bias signal we propose. However, our proposal provides a principled rationale for how much influence the prior has on the DV. Unlike the other proposals, which identify parameters to fit, our theory specifies the magnitude of the change in behavioral bias in the face of a prior and predicts the change in this magnitude under different conditions, such as a change in speed–accuracy regime. This *bravado* is mitigated somewhat by our need to fit the subjects' internal estimate of the prior. However, it is not surprising that subjects should underestimate the prior ($\phi < 1$), because there is good reason to bias a prior toward neutral [a prior on prior probability (Good, 1983)]. Whatever the reason, it appears that our human subjects held the same biased prior across different speed–accuracy regimes.

The idea we propose could be extended to more complex decisions and to other changes in reward structure affecting costs and risk. The central idea is that the state of a DV at the time of decision termination is associated with a set of benchmark outcomes. If a neural firing rate (e.g., in LIP) represents a DV, then the association specifies the scaling of firing rate to the units of this benchmark, as a function of decision time. In the formulation pursued here, the benchmark outcome is the odds of a correct choice when the DV is at the termination bound, so the scaling is from the DV relative to a bound into units of log odds. The benchmark could also include the incentives and costs associated with different choices. Thus, asymmetric rewards should also give rise to a dynamic bias signal when there is uncertainty about the reliability of the evidence.

Recent studies examined the effects of just such a manipulation using a fixed stimulus duration version of the motion discrimination task in monkeys (Feng et al., 2009). In this fixed duration design, behavioral data cannot distinguish static bias offsets from dynamic bias signals. However, recordings from LIP provided evidence for a static offset without providing evidence for a dynamic bias signal (Rorie et al., 2010). We think that performing the same manipulation with an RT version of the task may allow a more sensitive test of our theory.

There is ample reason to believe that elapsed time might influence the computations that underlie decision processes. Prolonged deliberation comes with the cost of time. To mitigate this cost, time-dependent signals that implement “decision urgency” have been suggested as a way to impose a soft deadline for decid-

ing (Churchland et al., 2008; Cisek et al., 2009). In addition to being used to reign in overextended deliberation, elapsed time also seems to play a role in evaluation of sensory evidence. Subjects have been found to use a combination of accumulated evidence and viewing time to determine their level of confidence; that is, the same amount of accumulated evidence seems to afford different levels of confidence depending on the time taken to accumulate that evidence (Kiani and Shadlen, 2009). Based on our results, it appears that a similar strategy is used when combining accumulated evidence with prior information. This supports the more general conclusion that the brain can exploit elapsed decision time to calibrate neural representations of probabilistic quantities. Thus, an inherent malleability exists in the mapping between spike rates and probabilities, and this malleability depends, at least partly, on time.

References

- Adelson EH, Bergen JR (1985) Spatio-temporal energy models for the perception of motion. *J Opt Soc Am A* 2:284–299.
- Ashby FG (1983) A biased random walk model for two choice reaction times. *J Math Psychol* 27:277–297.
- Barlow HB (1969) Pattern recognition and the responses of sensory neurons. *Ann N Y Acad Sci* 156:872–881.
- Basso MA, Wurtz RH (1998) Modulation of neuronal activity in superior colliculus by changes in target probability. *J Neurosci* 18:7519–7534.
- Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A (2008) Probabilistic population codes for Bayesian decision making. *Neuron* 60:1142–1152.
- Bisley JW, Krishna BS, Goldberg ME (2004) A rapid and precise on-response in posterior parietal cortex. *J Neurosci* 24:1833–1838.
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. *Psychol Rev* 113:700–765.
- Born RT, Bradley DC (2005) Structure and function of visual area MT. *Annu Rev Neurosci* 28:157–189.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10:1157–1169.
- Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62.
- Chang JS, Cooper G (1970) A practical difference scheme for Fokker-Planck equations. *J Comp Physiol* 6:1–16.
- Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11:693–702.
- Cisek P, Puskas GA, El-Murr S (2009) Decisions in changing conditions: the urgency-gating model. *J Neurosci* 29:11560–11571.
- Cohen MR, Newsome WT (2009) Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. *J Neurosci* 29:6635–6648.
- Diederich A, Busemeyer JR (2006) Modeling the effects of payoff on response bias in a perceptual discrimination task: bound-change, drift-rate-change, or two-stage-processing hypothesis. *Percept Psychophys* 68:194–207.
- Ditterich J (2006) Evidence for time-variant decision making. *Eur J Neurosci* 24:3628–3641.
- Edwards W (1965) Optimal strategies for seeking information: models for statistics, choice reaction-times, and human information-processing. *J Math Psychol* 2:312–329.
- Feng S, Holmes P, Rorie A, Newsome WT (2009) Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Comput Biol* 5:e1000284.
- Glimcher PW (2003) The neurobiology of visual-saccadic decision making. *Annu Rev Neurosci* 26:133–179.
- Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70:216–220.
- Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 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. *Annu Rev Neurosci* 30:535–574.
- Gold JI, Law CT, Connolly P, Benucci S (2008) The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning. *J Neurophysiol* 100:2653–2668.
- Good IJ (1983) Good thinking: the foundations of probability and its applications. Minneapolis: University of Minnesota.
- Hanks TD, Ditterich J, Shadlen MN (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat Neurosci* 9:682–689.
- 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.
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25:10420–10436.
- Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 8:234–241.
- Jaynes ET (2003) Probability theory: the logic of science. Cambridge, UK: Cambridge UP.
- Jazayeri M, Movshon JA (2006) Optimal representation of sensory information by neural populations. *Nat Neurosci* 9:690–696.
- Jeffreys H (1961) Theory of probability, Ed 3. Oxford: Clarendon.
- Kass RE, Raftery A (1995) Bayes factors. *J Am Stat Assoc* 90:773–795.
- 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. *J Neurosci* 28:3017–3029.
- Laming DRJ (1968) Information theory of choice-reaction times. London: Academic.
- Leon MI, Shadlen MN (2000) Representation of elapsed time by neurons in area LIP of the macaque. *Soc Neurosci Abstr* 26:249.213.
- Lewis JW, Van Essen DC (2000) Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *J Comp Neurol* 428:112–137.
- Link SW (1992) The wave theory of difference and similarity. Hillsdale, NJ: Erlbaum.
- Link SW, Heath RA (1975) A sequential theory of psychological discrimination. *Psychometrika* 40:77–105.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Maimon G, Assad JA (2006) A cognitive signal for the proactive timing of action in macaque LIP. *Nat Neurosci* 9:948–955.
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cereb Cortex* 13:1257–1269.
- Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J Vis* 5:376–404.
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238.
- Platt ML, Huettel SA (2008) Risky business: the neuroeconomics of decision making under uncertainty. *Nat Neurosci* 11:398–403.
- Press WH, Flannery BP, Teukolsky SA, Vetterling WT (1988) Numerical recipes in C. Cambridge, UK: Cambridge UP.
- Rao RP (2004) Bayesian computation in recurrent neural circuits. *Neural Comput* 16:1–38.
- Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput* 20:873–922.
- Ratcliff R, Rouder JN (1998) Modeling response times for two-choice decisions. *Psychol Sci* 9:347–356.
- Ratcliff R, Smith PL (2004) A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111:333–367.
- Reddi BA, Asrress KN, Carpenter RH (2003) Accuracy, information, and response time in a saccadic decision task. *J Neurophysiol* 90:3538–3546.
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489.
- Romo R, Hernández A, Zainos A (2004) Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron* 41:165–173.
- Rorie AE, Gao J, McClelland JL, Newsome WT (2010) Integration of sen-

- sory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS One* 5:e9308.
- Schall JD, Thompson KG (1999) Neural selection and control of visually guided eye movements. *Annu Rev Neurosci* 22:241–259.
- Seidemann E (1998) Neural mechanisms mediating conversion of visual signals into perceptual decisions in a direction discrimination task. PhD thesis, Stanford University Neuroscience Program.
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916–1936.
- Shadlen MN, Hanks TD, Churchland AK, Kiani R, Yang T (2006) The speed and accuracy of a simple perceptual decision: a mathematical primer. In: *Bayesian brain: probabilistic approaches to neural coding* (Doya K, Ishii S, Rao R, Pouget A, eds), pp 209–237. Cambridge, MA: Massachusetts Institute of Technology.
- 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. *J Exp Psychol Hum Percept Perform* 35:1865–1897.
- Sugrue LP, Corrado GS, Newsome WT (2004) Matching behavior and the representation of value in the parietal cortex. *Science* 304:1782–1787.
- Van Essen DC, Drury HA, Dickson J, Harwell J, Hanlon D, Anderson CH (2001) An integrated software suite for surface-based analyses of cerebral cortex. *J Am Med Inform Assoc* 8:443–459.
- Wald A (1947) *Sequential analysis*. New York: Wiley.
- Wald A, Wolfowitz J (1947) Optimum character of the sequential probability ratio test. *Ann Math Statist* 19:326–339.
- Zemel RS, Dayan P, Pouget A (1998) Probabilistic interpretation of population codes. *Neural Comput* 10:403–430.