

# Neural Dynamics of Choice: Single-Trial Analysis of Decision-Related Activity in Parietal Cortex

Anil Bollimunta,<sup>1</sup> Douglas Totten,<sup>1</sup> and Jochen Ditterich<sup>1,2</sup>

<sup>1</sup>Center for Neuroscience and <sup>2</sup>Department of Neurobiology, Physiology, and Behavior, University of California, Davis, Davis, California 95618

Previous neurophysiological studies of perceptual decision-making have focused on single-unit activity, providing insufficient information about how individual decisions are accomplished. For the first time, we recorded simultaneously from multiple decision-related neurons in parietal cortex of monkeys performing a perceptual decision task and used these recordings to analyze the neural dynamics during single trials. We demonstrate that decision-related lateral intraparietal area neurons typically undergo gradual changes in firing rate during individual decisions, as predicted by mechanisms based on continuous integration of sensory evidence. Furthermore, we identify individual decisions that can be described as a change of mind: the decision circuitry was transiently in a state associated with a different choice before transitioning into a state associated with the final choice. These changes of mind reflected in monkey neural activity share similarities with previously reported changes of mind reflected in human behavior.

## Introduction

Sequential sampling models, such as the drift-diffusion model (for review, see Ratcliff and McKoon, 2008) or the leaky competing accumulator model (Usher and McClelland, 2001), are a popular class of models for perceptual decision-making. In this framework, every decision is characterized by a gradual accumulation of sensory evidence until a decision criterion is reached. If a neural circuit were to implement such a mechanism, pools of decision-related neurons should gradually change their activity on a single-trial basis. In contrast, other authors have suggested that arriving at a decision might be better characterized by a more abrupt, jump-like transition from an uncommitted to a committed state (Durstewitz and Deco, 2008; Miller and Katz, 2010). As will be discussed in more detail below, either type of dynamics can be realized with neural attractor networks.

Previous neurophysiological studies of the random-dot motion direction discrimination task have focused on single-unit (SU) recordings in nonhuman primates. Decision-related activity has been observed in a variety of brain areas, including the lateral intraparietal area (LIP) in parietal cortex (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Churchland et al., 2008), areas in (pre)frontal cortex (Kim and Shadlen, 1999), and the superior colliculus (SC; Horwitz and Newsome, 2001). Although the trial-averaged firing

rate of decision-related LIP neurons has been demonstrated to be consistent with an integration-to-threshold mechanism (Mazurek et al., 2003; Ditterich, 2006; Churchland et al., 2008), the neural dynamics of individual decisions cannot be reliably assessed with the limited number of spikes that are emitted by a single neuron. Horwitz and Newsome (2001) have noticed streakiness in the spike trains of SC neurons and have raised the question whether this is attributable to relatively fast transitions between a low and a high firing rate state. In contrast, Roitman and Shadlen (2002) and Churchland et al. (2011) have provided indirect arguments for LIP neurons undergoing a gradual firing rate change.

To directly assess the neural dynamics in LIP during individual decisions, we, for the first time, take advantage of simultaneous recordings from multiple decision-related neurons during a three-choice version of the motion discrimination task. First, using simultaneous recordings from multiple neurons belonging to the same decision pool, we focused on the question of how the winning pool makes its transition from an uncommitted to a committed state. Using a number of different analysis techniques, we provide evidence for single-trial LIP spike trains being more consistent with a gradual change in activity rather than more abrupt state transitions. To analyze simultaneous recordings from neurons with different choice selectivity, we performed a hidden Markov model (HMM) analysis (Radons et al., 1994; Abeles et al., 1995; Seidemann et al., 1996; Gat et al., 1997; Jones et al., 2007; Miller and Katz, 2010). This analysis allowed us to further characterize the single-trial dynamics and to detect “changes of mind” during individual decisions. We characterize the properties of these changes of mind reflected in monkey neural activity and compare them with previously reported properties of changes of mind reflected in human behavior (Resulaj et al., 2009).

## Materials and Methods

**Experimental design and data collection.** Details of the experimental approach can be found in the study by Bollimunta and Ditterich (2012). Briefly, three rhesus macaques (*Macaca mulatta*) of either sex were im-

Received Nov. 17, 2011; revised June 18, 2012; accepted July 18, 2012.

Author contributions: J.D. designed research; A.B. performed research; A.B., D.T., and J.D. analyzed data; J.D. wrote the paper.

This material is based on research sponsored by the Air Force Research Laboratory under Agreements FA9550-07-1-0205 and FA9550-07-1-0537. We thank Monica Chau, Ryan Herbert, Kelly McFarland, and Michelle Stuart for animal care and training. A special thank you goes to Paul Miller for providing some MATLAB code and helping us with getting started with the HMM analysis. We are also thankful to QNX Software Systems for supporting our research through providing their operating system free of charge as part of their QNX-in-Education program.

Correspondence should be addressed to Jochen Ditterich, University of California, Davis, Center for Neuroscience, 1544 Newton Court, Davis, CA 95618. E-mail: jditterich@ucdavis.edu.

A. Bollimunta's present address: Laboratory of Sensorimotor Research, National Eye Institute, Bethesda, MD 20892.

DOI:10.1523/JNEUROSCI.5752-11.2012

Copyright © 2012 the authors 0270-6474/12/3212684-18\$15.00/0

planted with a recording chamber giving access to LIP in a sterile surgery under general anesthesia. A multielectrode drive with five independently movable electrodes (Mini Matrix; Thomas Recording) was used for isolating multiple SUs with the help of Sort Client (Plexon). The response fields (RFs) of these SUs were then mapped simultaneously using a memory-guided saccade task. Only SUs with spatially selective activity during the memory period were selected for recording during the perceptual decision task. All procedures were approved by the Institutional Animal Care and Use Committee and were in accordance with the *National Institutes of Health Guide for the Care and Use of Laboratory Animals*.

The monkeys had to keep fixation, monitored by an infrared video eye tracker (Applied Science Laboratories), while watching a three-component random-dot motion stimulus. On each trial, the coherences of the three embedded motion components, ranging from 0 to 40%, were randomly drawn from 51 possible combinations. The three directions of the coherent motion components were always separated by 120°. Each motion direction was associated with one of three choice targets that appeared on the computer screen before motion stimulus onset. These targets were placed such that as many of the isolated SUs as possible had exactly one of the targets in their RFs. This could result in multiple neurons with overlapping RFs having the same target in their RFs, but sometimes we had neurons with non-overlapping RFs and were able to place two different targets in the RFs. The monkeys watched the stimulus until they were ready to respond and then made a goal-directed eye movement to one of the targets. They received a fluid reward for choosing the target associated with the strongest motion component in the stimulus. When all three motion components had identical coherences, a situation referred to as zero net motion strength, reward was given randomly with a probability of 1/3. The monkeys' choices and response times (RTs) were recorded.

*Generation of artificial datasets and firing rate plots.* All analyses described here were not only applied to the real LIP data but also to artificial datasets that were designed to test the sensitivity and reliability of our analysis techniques. Artificial data matched critical properties of the real datasets, including number of trials and initial and final firing rates. For the streak index (SI) analysis and for the maximum likelihood fits, the initial and final firing rates were obtained from the real data using linear regression. Spikes in 400 successive (1 ms) time bins were counted across all trials, multiplied by 1000, and divided by the number of trials to obtain rate estimates. Linear regression was applied, and the resulting rates associated with the first and the last time bins were taken as the initial and final firing rates ( $\lambda_{\text{initial}}$  and  $\lambda_{\text{final}}$ ). For the global likelihood ratio analysis, we simply used the average firing rates during the first and last 40 ms. For each real LIP dataset, we created one matching artificial ramping dataset (with the same number of trials as in the real dataset) by generating Poisson spike trains following the rate function

$$\lambda_{\text{ramping}}(t) = \lambda_{\text{initial}} + (\lambda_{\text{final}} - \lambda_{\text{initial}}) \times \frac{t}{T}$$

with  $T$  being the duration of the analyzed time interval (400 ms for the SI analysis and the maximum likelihood fits and 300 ms for the global likelihood ratio analysis) and  $t$  ranging from 0 to  $T$ . The second artificial dataset (artificial jumps) was created by generating Poisson spike trains following the family of rate functions:

$$\lambda_{\text{jumping}}(t) = \begin{cases} \lambda_{\text{initial}}; & t < t_{\text{step}} \\ \lambda_{\text{final}}; & t \geq t_{\text{step}} \end{cases}$$

For each generated trial,  $t_{\text{step}}$  was randomly drawn from a uniform distribution ranging from 0 to  $T$ .

For the HMM analysis, we also generated two artificial datasets (artificial ramps and artificial jumps) for each real LIP dataset. The artificial datasets had the same number of trials and the trials had the same lengths as in the real dataset, but the spike trains were replaced by nonstationary Poisson spike trains that were generated according to the following rules. Because the average response of the LIP neurons during the first 200 ms of a trial was not sensitive to the sensory stimulus (Bollimunta and Ditterich, 2012; see also Roitman and Shadlen, 2002; Churchland et al.,

2008), the first 200 ms of each artificial spike train were generated by using the average firing rate time course of the particular neuron, determined by counting spikes in 50 ms windows. After this, in the case of the artificial ramps, the remaining duration of the trial was used to have the firing rate ramp to average final firing rate of this neuron for this particular type of choice, as determined by the average spike count during the last 50 ms of each trial ending with this choice. In the case of the artificial jumps, a random time was drawn from a uniform distribution ranging from 200 ms to the trial duration. Each neuron remained at its current firing rate up to this random time and then jumped to its average final firing rate for the particular choice.

All firing rates shown in Figures 2*a*, 5*a*, 7*b*, and 8, *b* and *c*, have been smoothed with a Gaussian kernel with an SD of 20 ms.

*SI analysis.* For obtaining the SIs, the spike counts in 16 successive 25 ms bins (or, alternatively, 10 successive 40 ms bins) were obtained for each trial. We then calculated the median spike count in each of the 16 (10) bins across all trials. On a trial-by-trial basis, the spike counts in each bin were compared with the corresponding median spike counts. If the spike count was larger than the median, the symbol "1" was assigned to the corresponding bin. If the spike count was smaller than the median, the symbol "0" was assigned to the corresponding bin. If the spike count was identical to the median, either a 0 or 1 were randomly assigned. We then counted how often switches either from 0 to 1 or from 1 to 0 occurred in the resulting 16-digit (10-digit) string. If there are  $N$  switches, one says that the string contains  $N + 1$  runs. If the 0s and 1s are randomly arranged, the number of runs is expected to be drawn from a distribution with mean  $\mu$  and SD  $\sigma$  with

$$\mu = 1 + \frac{2mn}{m+n}$$

and

$$\sigma = \sqrt{\frac{2mn \times (2mn - m - n)}{(m+n)^2 \times (m+n-1)}}$$

with  $m$  and  $n$  being the number of 0s and 1s in the string (Horwitz and Newsome, 2001). The SI is defined as  $(\text{number of runs} - \mu)/\sigma$  and is therefore expected to follow a distribution with 0 mean and an SD of 1 if the 0s and 1s are randomly arranged. Such an arrangement is, for example, expected when each spike train is an instantiation of the same (potentially non-homogeneous) Poisson process. Conversely, if the trial-averaged response shows a gradual increase in firing rate, but individual trials are characterized by a sudden jump from a low firing rate to a high firing rate at some random time, we expect fewer runs than in the case of a random arrangement of 0s and 1s and therefore a negative SI. A  $t$  test was used to determine whether the mean SI across trials was significantly different from 0. For testing distributions of mean SIs across sessions or neurons, a sign test was used to determine whether the median of the distribution was significantly different from 0, and a Wilcoxon's test was used for comparing distributions because the data were no longer expected to follow normal distributions. To determine whether the real data SIs were significantly closer to the results for a particular type of artificial data than for the other type, we performed another Wilcoxon's test to compare the real data SIs with the midpoints between the results for the two artificial datasets.

In contrast to Horwitz and Newsome (2001), who reported a single SI based on all trials, we calculated an SI for each individual trial and analyzed its distribution across trials. Besides providing a single-trial measure, this has the advantage that the mean of the SI distribution across trials is not expected to scale with the number of trials, whereas the absolute value of the number reported by Horwitz and Newsome tends to increase with the number of analyzed trials. Thus, the numbers reported by Horwitz and Newsome cannot be directly compared with our mean SIs, but what can be compared is how often the estimated SI significantly deviates from 0.

*Global likelihood ratios between a linear model and a step model.* To test whether the single-trial population spike trains are more consistent with

a linear model or with a step model, we estimated the global likelihood ratio for these two models. The linear model assumes that each trial is an instantiation of the same non-homogeneous Poisson process with a linear rate change over time. The step model assumes that each trial is an instantiation of non-homogeneous Poisson processes with common initial rate and common final rate and a rate jump at some random time, which can be different across trials.

Let  $\{t_i^r\}$  be the set of spike times recorded from neurons with the same choice selectivity, where  $t_i^r$  is the time of the  $i$ th spike in the  $r$ th trial. The global likelihood ratio between a linear model and a step model can then be written as

$$\frac{P(\{t_i^r\}|\text{Linear})}{P(\{t_i^r\}|\text{Step})},$$

where  $P(\{t_i^r\}|\text{Linear})$  is the global likelihood of observing the recorded spike trains given the linear model, and  $P(\{t_i^r\}|\text{Step})$  is the global likelihood of observing the recorded spike trains given the step model.

Assuming that each trial is an instantiation of the same non-homogeneous Poisson process with a rate that changes linearly over time, starting at a rate of  $\lambda_{\text{initial}}$  and finishing at a rate of  $\lambda_{\text{final}}$ , the global likelihood for the linear model based on  $N$  trials can be written as (Gregory, 2005; Bollimunta et al., 2007)

$$P(\{t_i^r\}|\text{Linear}) = \int_{\lambda_{\text{initial}}} \int_{\lambda_{\text{final}}} \left( \prod_{r=1}^N L_{\text{Linear}}^r(\lambda_{\text{initial}}, \lambda_{\text{final}}) \right) \times p(\lambda_{\text{initial}}, \lambda_{\text{final}}) d\lambda_{\text{initial}} d\lambda_{\text{final}},$$

where  $L_{\text{Linear}}^r(\lambda_{\text{initial}}, \lambda_{\text{final}})$  is the likelihood of observing the spike train recorded in the  $r$ th trial given a linear model with parameters  $\lambda_{\text{initial}}$  and  $\lambda_{\text{final}}$  and  $p(\lambda_{\text{initial}}, \lambda_{\text{final}})$  is the prior probability density of model parameters (see below). The likelihood was calculated by binning the spike train with a resolution ( $\Delta t$ ) of 1 ms. Given a rate function  $\lambda(t)$  of a non-homogeneous Poisson process, the probability of observing a spike in a given time bin at time  $t$  is given by  $p_{\text{spike}}(t) = \lambda(t) \times \Delta t$ , and the probability of not observing a spike in the same time bin is therefore  $p_{\text{no spike}}(t) = 1 - p_{\text{spike}}(t)$ . The likelihood of observing a particular spike train given rate function  $\lambda(t)$  is therefore

$$L = \prod_{\text{all bins with spikes}} (\lambda(t_j) \times \Delta t) \times \prod_{\text{all bins without spikes}} (1 - \lambda(t_j) \times \Delta t),$$

with  $t_j$  being the time corresponding to a particular bin.

Assuming that each trial is an instantiation of a non-homogeneous Poisson process with initial rate  $\lambda_{\text{initial}}$  and final rate  $\lambda_{\text{final}}$  but rate functions that step from  $\lambda_{\text{initial}}$  to  $\lambda_{\text{final}}$  at different times  $t_{\text{step}}^r$ , the global likelihood for the step model can be written as

$$P(\{t_i^r\}|\text{Step}) = \int_{\lambda_{\text{initial}}} \int_{\lambda_{\text{final}}} \int_{t_{\text{step}}^1} \int_{t_{\text{step}}^2} \dots \int_{t_{\text{step}}^N} \left( \prod_{r=1}^N L_{\text{Step}}^r(\lambda_{\text{initial}}, \lambda_{\text{final}}, t_{\text{step}}^r) \right) \times p(\lambda_{\text{initial}}, \lambda_{\text{final}}, \{t_{\text{step}}^r\}) d\lambda_{\text{initial}} d\lambda_{\text{final}} dt_{\text{step}}^1 \dots dt_{\text{step}}^N$$

where  $L_{\text{Step}}^r(\lambda_{\text{initial}}, \lambda_{\text{final}}, t_{\text{step}}^r)$  is the likelihood of observing the spike train recorded in the  $r$ th trial given a step model with parameters  $\lambda_{\text{initial}}$ ,  $\lambda_{\text{final}}$ , and a step time of  $t_{\text{step}}^r$  for the  $r$ th trial, which can be calculated according to the likelihood formula given above, and  $p(\lambda_{\text{initial}}, \lambda_{\text{final}}, \{t_{\text{step}}^r\})$  is the prior probability density of model parameters (see below).

Although the computation of the global likelihood for the linear model is relatively straightforward, the step model requires integration over  $N + 2$  parameters and becomes computationally intractable even for modest numbers of trials. We therefore estimated the global likelihood ratio for only four trials at a time. Furthermore, we used a shorter 300 ms time window ranging from 350 to 50 ms before saccade onset and sampled the parameter space with a resolution of 2 Hz for the rates and a

resolution of 10 ms for the step times. All model parameters were assumed to be statistically independent, and the priors were set to uniform distributions, excluding negative values, ranging from  $\hat{\lambda}_{\text{initial}} - 50$  Hz to  $\hat{\lambda}_{\text{initial}} + 50$  Hz for the initial firing rate, from  $\hat{\lambda}_{\text{final}} - 50$  Hz to  $\hat{\lambda}_{\text{final}} + 50$  Hz for the final firing rate, and from 0 to 300 ms for each step time (relative to the beginning of the analyzed interval), with  $\hat{\lambda}_{\text{initial}}$  being the average firing rate during the first 40 ms of the analyzed interval and  $\hat{\lambda}_{\text{final}}$  being the average firing rate during the last 40 ms. For each session, we repeated this analysis 100 times, each time based on a different random set of four trials. The analysis was always performed on the real LIP dataset as well as two matched artificial datasets, one ramping and one jumping, with initial rate  $\hat{\lambda}_{\text{initial}}$  and final rate  $\hat{\lambda}_{\text{final}}$ .

A sign test was used to determine whether the median global log likelihood ratio (LLR) for a given session or across sessions was significantly different from 0. A Wilcoxon's test was used for comparing the distributions of median global LLRs across sessions. Another Wilcoxon's test was used to determine whether the global LLR for the real data was significantly different from the midpoint between the global LLRs for the two artificial datasets and therefore closer to one of them than to the other.

*Maximum likelihood fits of non-homogeneous Poisson processes.* To allow rate variations across trials and only test for the shape of the rate function underlying individual spike trains, we also obtained maximum likelihood estimates of both previously mentioned non-homogeneous Poisson models (linear and step) on an individual trial basis. Likelihoods were calculated according to the formula presented in the previous section. If the rate function  $\lambda(t)$  has a set of parameters, the combination of parameters maximizing the likelihood is called the maximum likelihood solution. We considered two different rate functions:

$$\lambda_{\text{linear}}(t) = \lambda_{\text{initial}} + \Delta\lambda \times t$$

and

$$\lambda_{\text{step}}(t) = \begin{cases} \lambda_{\text{initial}}; & t < t_{\text{step}} \\ \lambda_{\text{final}}; & t \geq t_{\text{step}} \end{cases},$$

with  $\lambda_{\text{linear}}(t)$  having two free parameters, the initial firing rate  $\lambda_{\text{initial}}$  and the firing rate slope  $\Delta\lambda$ , and  $\lambda_{\text{step}}(t)$  having three free parameters, the initial firing rate  $\lambda_{\text{initial}}$ , the final firing rate  $\lambda_{\text{final}}$ , and the step time  $t_{\text{step}}$ . The maximum likelihoods associated with both models are  $L_{\text{max,linear}}$  and  $L_{\text{max,step}}$ .

Because the two models have unequal numbers of parameters, a penalty term has to be introduced when comparing the maximum likelihoods. We used the Hannan–Quinn information criterion (HQIC), which, based on classifying the artificial datasets, proved less biased than the more commonly used Akaike information criterion (AIC) or Bayesian information criterion (BIC) (Hannan and Quinn, 1979). The HQIC values reported in this paper were calculated as

$$HQIC = -2 \times \ln L_{\text{max}}/n + 2 \times k \times \ln(\ln(n))/n,$$

with  $k$  being the number of free parameters (two for the linear model and three for the step model) and  $n$  being the number of data points (400 in our case because each trial had 400 time bins). A smaller HQIC value means that the associated model provides a better explanation for the data.

For obtaining the maximum likelihood solutions, the maximum of the likelihood function had to be found, which was done using a multidimensional simplex algorithm, and was relatively straightforward for the linear model when reasonable initial estimates were provided for the model parameters. The fit of the step function is more problematic because it can easily get stuck in local maxima. To avoid this problem, rather than directly fitting a step function, we applied an iterative fit of a logistic function with an increasingly steeper slope:

$$\lambda_{\text{logistic}}(t) = \lambda_{\text{initial}} + \frac{\Delta\lambda}{1 + \exp(-\alpha \times (t - t_{\text{step}}))},$$

with  $\alpha$  being the slope parameter, which was initialized with a value of 0.025 and multiplied by a factor of 10 after each iteration. The last iteration was run when  $\alpha$  reached a value of at least 10, in which case it was set

to 10 for the final iteration. The best estimates of  $\lambda_{\text{initial}}$ ,  $\Delta\lambda$ , and  $t_{\text{step}}$  after each iteration were used as the initial parameter guesses for the next iteration, ensuring a stable convergence of the fit. The only imposed parameter constraints were that  $\lambda(t)$  was not allowed to exceed 1000 Hz and  $t_{\text{step}}$  had to remain inside the analyzed time interval.

A sign test was used to determine whether the median HQIC difference (between the linear and the step model) across trials or across sessions was significantly different from 0. A Wilcoxon's test was used to compare distributions of HQIC differences across sessions. Another Wilcoxon's test was used to determine whether the HQIC difference for the real data was significantly different from the midpoint between the HQIC differences for the two artificial datasets and therefore closer to one of them than to the other.

The MATLAB functions that were used for obtaining the maximum likelihood estimates can be obtained from <http://master.peractionlab.org/software>.

**HMM analysis.** For obtaining the HMM for a recording session, the spike trains recorded during each trial were converted into an emission sequence as explained in the Results, HMM analysis and changes of mind, below. Because RT varied across trials, the emission sequences had various lengths. The whole set of emission sequences was then used for fitting the model using the Baum–Welch algorithm in MATLAB (MathWorks; note that one has to be added to each of the emitted symbols in this case because MATLAB does not accept 0 as a valid emission). Ten different initial conditions for the emission matrix were used for a given model order (number of states) to avoid getting stuck in a local optimum, and the model with the maximum likelihood was chosen as the best model of this order. The transition matrix was always initialized with 0.99 along the main diagonal and all identical off-diagonal entries. The emission matrix was initialized such that the spiking probability of each neuron in each state was randomly chosen to reflect a uniform distribution of firing rates between 0 and 50 Hz.

Specifying the appropriate order of an HMM is a difficult problem. Previous applications to the analysis of spike trains have circumvented this problem by choosing an arbitrary number of states (Seidemann et al., 1996; Jones et al., 2007; Miller and Katz, 2010). We took a more data-driven approach. Note that the maximum likelihood cannot be used for finding the optimal model order because it never decreases as the model order increases. An attempt to use a strongly consistent estimator, like the one proposed by Baras and Finesso (1992), failed because it almost always limited the number of required states to two. We therefore used the AIC, which applies a conservative penalty term for increasing model order, for obtaining an upper bound for the model order and a series of tests to further reduce the model order as necessary. The first two tests were based on diagnostic tools proposed by McCane and Caelli (2004). A sudden drop in the inverse condition number (ICN) of the augmented matrix with increasing model order for an order that was not larger than the one determined based on the AIC was taken as an indicator that the order should not be chosen larger than the one just before the drop in the ICN. The second test was based on examining the values of  $s_i$  as defined by McCane and Caelli (2004, their Eq. 19). Any  $s_i$  of  $<0.9$  for a given model order was taken as an indicator of redundancy among the states and therefore as a hint that a smaller model order should be chosen. Our final test was designed to avoid overfitting. To this end, we randomly split the emission sequences into two halves. We used one half for fitting the model and the other half for calculating the likelihood. In contrast to the maximum likelihood calculated from the same data that were also used for fitting the model, this likelihood has a maximum for some model order, indicating that models with more states start reflecting peculiarities in the part of the data used for model fitting, which are not present in the part of the data used for calculating the likelihood. If the peak in the likelihood was observed for a model order that was smaller than the one determined based on the AIC, the model order was reduced accordingly.

The resulting transition and emission matrices were used for constructing the HMM diagrams as shown in Figures 7a and 8a. Each entry in the transition matrix reflects the probability of making a particular state transition in each time step (2 ms in our case). Transition probabilities of at least 0.001 (when rounded to three decimals) are indicated by arrows. Entries in the emission matrix reflect the probabilities of a par-

ticular neuron emitting a spike in each time step (2 ms) when the system is in a particular state. Thus, multiplying these numbers by 500 results in the firing rates shown in the diagrams. The determined transition and emission matrices can then be used to decode the emission sequences of individual trials (see Figs. 7e–i, 8d–i). The decoding results in a set of posterior state probabilities as a function of time (shown as colored lines in the figures). Consistent with previous work (Seidemann et al., 1996; Jones et al., 2007; Miller and Katz, 2010), a posterior probability exceeding 0.8 was taken as an indicator that the system was in one of the identified discrete states (marked by colored shaded areas in the figures).

**Motion energy filters.** To extract the time course of motion energy from our visual stimuli, we used motion energy filters similar to the ones used by Kiani et al. (2008) but with a slightly faster temporal response to shift the preferred speed closer to the speeds used in our stimuli. The spatial filters were

$$f_1(x', y') = \cos^4(\alpha) \times \cos(4\alpha) \times \exp\left(-\frac{y'^2}{2\sigma_g^2}\right)$$

$$f_2(x', y') = \cos^4(\alpha) \times \sin(4\alpha) \times \exp\left(-\frac{y'^2}{2\sigma_g^2}\right)$$

$$\alpha = \tan^{-1}\left(\frac{x'}{\sigma_c}\right).$$

$\sigma_g$  was set to  $0.05^\circ$  to approximate the direction selectivity of MT neurons. As in the study by Kiani et al. (2008),  $\sigma_c$  was chosen as  $0.35^\circ$ .  $(x', y')$  is a rotated coordinate system able to extract motion energy in different preferred directions  $\theta$ :

$$x' = \cos(\theta) \times x - \sin(\theta) \times y,$$

$$y' = \sin(\theta) \times x + \cos(\theta) \times y.$$

We always used three sets of filters whose preferred directions were aligned with the three directions of coherent motion in the random-dot stimulus. The temporal impulse responses were

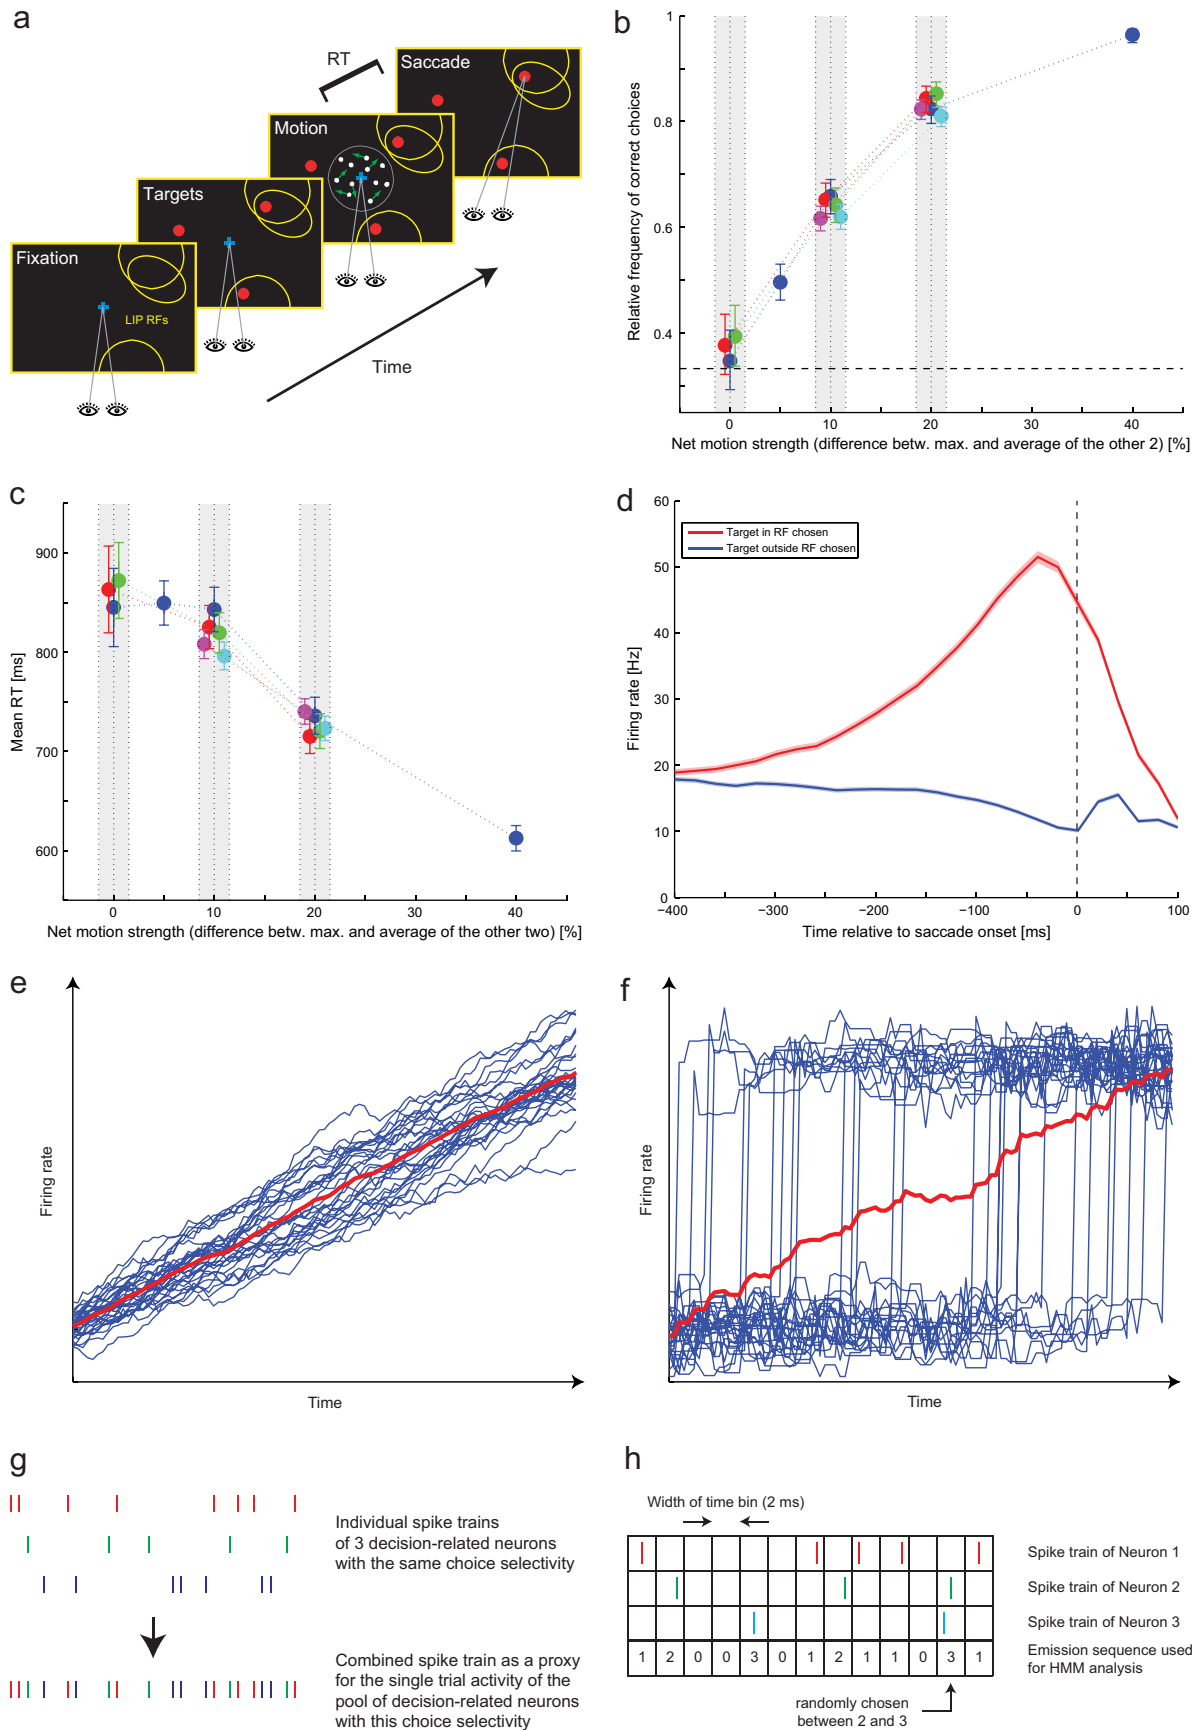
$$g_1(t) = (100 \times t)^3 \times \exp(-100 \times t) \times \left[ \frac{1}{3!} - \frac{(100 \times t)^2}{5!} \right]$$

$$g_2(t) = (100 \times t)^5 \times \exp(-100 \times t) \times \left[ \frac{1}{5!} - \frac{(100 \times t)^2}{7!} \right].$$

For each preferred direction, we constructed a pair of filters in space–time quadrature ( $f_1g_1 + f_2g_2$  and  $f_2g_1 - f_1g_2$ ) and convolved the stimulus with the resulting kernels, extending from  $-0.7^\circ$  to  $+0.7^\circ$  in space and from 0 to 0.2 s in time. The motion energy at a particular time point is therefore determined by the preceding 200 ms of the stimulus. The results were squared and summed and then collapsed across space to obtain the motion energy in a particular direction ( $\theta$ ) as a function of time. The time course of the net motion energy supporting a particular choice was calculated as the difference between the motion energy of the stimulus component supporting this choice and the average motion energy of the other two stimulus components.

## Results

We trained three rhesus monkeys to perform a three-choice version of the random-dot motion direction discrimination task, which is illustrated in Figure 1a. While looking at a fixation spot, the monkeys watched a random-dot motion stimulus that had three embedded components of coherent motion with directions that were all separated by  $120^\circ$ . The strengths of the three motion components were randomly varied from trial to trial. The monkeys were trained to identify the strongest component of motion and to report their choice by making a saccade to the target closest to the identified direction (of three available targets). The viewing duration was under the animal's control, and the monkey received a fluid reward for choosing the correct target. For the



**Figure 1.** Experimental task, behavior, average firing rates, and spike train preparation. **a**, Three-alternative forced-choice motion direction discrimination task: monkeys had to maintain fixation while watching a three-component random-dot motion stimulus. The task was to identify the strongest component of motion in the stimulus and to make a goal-directed eye movement to the target (red) closest to the identified direction. Choices and RTs were measured. We recorded simultaneously from multiple LIP neurons with exactly one of the targets in (*Figure legend continues.*)

analyses presented here, it is not particularly important that we used a three-component stimulus. In fact, the monkeys' behavior suggests that they would have made similar choices if just presented with a single component motion stimulus with a net direction identical to the direction of the strongest component in our stimulus and a coherence that is given by the difference between the coherence of the strongest component in our stimulus and the average coherence of the other two components. We call this difference the "net motion strength." Analyses that take advantage of the simultaneous control of sensory evidence for and against particular choices, as provided by the three-component stimulus, have been presented previously (Bollimunta and Ditterich, 2012).

The proportion of correct choices as a function of net motion strength, ranging from chance performance for zero net motion strength to almost perfect choice behavior for 40% net motion strength, is shown in Figure 1*b*. Different colors represent different combinations of the strengths of the three components in our stimulus that could give rise to the same net motion strength. The mean RTs, ranging from ~600 ms for 40% net motion strength to ~850 ms for non-informative stimuli (zero net motion strength), are shown in Figure 1*c*. While the monkeys performed this task, we recorded extracellularly from multiple LIP neurons that had exactly one of the choice targets in their RF, using a multielectrode drive. An example of how these RFs might have been arranged in a particular experimental session is shown in Figure 1*a*. The RFs were mapped using a memory-guided saccade task, and only neurons with spatially selective activity during the memory period were recorded. We are therefore focusing on the same subpopulation of LIP neurons that have been studied in previous work on perceptual decision-making (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Churchland et al., 2008; Law and Gold, 2008; Rorie et al., 2010; Bennur and Gold, 2011). The average firing rate of the recorded LIP neurons as a function of whether the monkey made a saccade to the target inside the RF (red trace) or to one of the targets outside the RF (blue trace) is shown in Figure 1*d*. Responses are aligned with the time of saccade onset, indicated by the dashed vertical line. The red trace indicates that, on average, LIP neurons showed a gradual increase in their firing rate during the last 400 ms up to ~40 ms before a saccade to the target inside

the RF of the recorded neuron. These non-single-trial-related results have previously been reported by Bollimunta and Ditterich (2012).

However, the shown trace is the result of averaging >9000 individual decisions. It is easy to demonstrate that very different behavior at the single-trial level can lead to very similar results after averaging across trials. The thin blue lines in Figure 1*e* show hypothetical single-trial firing rate functions. In this case, all of the functions are noisy versions of a linear ramp. The average firing rate across trials, as indicated by the thick red line, is again a ramp-like function. A quite different situation is illustrated in Figure 1*f*. In this case, the thin blue lines all reflect noisy versions of a sudden jump from a low to a high firing rate state at random times. The average firing rate across trials is again a ramp-like function (thick red line). Thus, the average firing rate across many trials does not reveal the neural dynamics of individual decisions. To address the question of how the activity of pools of decision-related LIP neurons changes during individual decisions, we will now focus on the analysis of spike trains that have been recorded during individual trials.

### SI analysis

We first focus on the question of how the winning pool of decision-related LIP neurons makes its transition from an uncommitted (low firing rate) to a committed (high firing rate) state. To this end, we analyzed the data from sessions in which we had recorded from at least two neurons with the same choice selectivity (same target in the RF). Averaged across trials, these neurons had to gradually increase their firing rate during the last 400 ms before a saccade to the target in the RF to ensure that we were looking at subpopulations of neurons with a trial-averaged response that was representative of the population. A few neurons, which did not show a ramping response until later in the trial (e.g., 200 ms before the saccade), were thus excluded from this analysis. Only trials that resulted in choosing the target inside the RF were analyzed, meaning that the recorded neurons were part of the winning pool, and all spikes that were emitted during the last 400 ms before saccade onset were extracted for analysis. Eighteen datasets had the required properties. The number of simultaneously recorded neurons with the same choice selectivity was either two or three, and the number of analyzed trials per session ranged from 79 to 360. Because we were interested in the dynamics of the population response and because the first analysis techniques that will be described here operate on a single spike train, we first generated a combined spike train for each trial by collapsing across individual neurons as illustrated in Figure 1*g*. This combined spike train is our best single-trial estimate of the response of the pool of LIP neurons with the same choice selectivity as the recorded exemplars. On average, the firing rate of the combined spike trains increased by 71 spikes/s during the analyzed interval, ranging from 41 to 108 spikes/s across datasets. This corresponds to an average increase of 34 spikes/s per recorded neuron, ranging from 20 to 54 spikes/s across datasets.

The idea of the analysis that will be presented in the following was to obtain the SI for each single-trial LIP spike train and to compare the distribution of SIs with those obtained from two artificially created datasets. These were tailored to match the real LIP dataset in three aspects: number of trials, average firing rate 400 ms before saccade onset, and average firing rate at the time of saccade onset. The first artificial dataset was generated under the assumption that each trial was an instantiation of the same non-homogeneous Poisson process with a linear change in firing rate from the initial to the final rate as determined from the real data.

←  
(Figure legend continued.) each RF (yellow ellipses). Different neurons could either have the same target or two different targets in their RFs. **b**, Accuracy: relative frequency of correct choices as a function of net motion strength (difference between the coherence of the strongest component and the average of the other two coherences). Different colors reflect different combinations of the two weaker coherences (different combinations of coherences could result in the same net motion strength). All points within a particular gray bar reflect the exact same net motion strength and have only been shifted horizontally for visualization purposes. The error bars reflect 95% confidence intervals, and the dashed horizontal line indicates chance performance. **b–d** are adapted from Bollimunta and Ditterich (2012) with permission. **c**, Mean RTs as a function of net motion strength; same conventions as in **b**. **d**, Average LIP firing rate as a function of whether the target in the RF (red) or one of the targets outside the RF (blue) was chosen. Responses are time-locked to saccade onset (dashed vertical line). The shaded areas indicate 95% confidence bands. **e**, Hypothetical trial-averaged ramping response (red) resulting from individual trials with ramping responses (blue). **f**, Hypothetical trial-averaged ramping response (red) resulting from individual trials with jumping responses (blue; fast transitions from a low to a high firing rate state at different times). **g**, A combined spike train is constructed from three simultaneously recorded single-trial spike trains for obtaining a single-trial estimate of the population response. This approach is used for the SI analysis at the population level and for comparisons of non-homogeneous Poisson models. **h**, Example of constructing an emission sequence from three simultaneously recorded spike trains. These emission sequences are used for the HMM analysis.

The second artificial dataset was generated under the assumption that each trial was an instantiation of a non-homogeneous Poisson process whose rate stepped from the initial to the final rate at some random time, with a uniform distribution of step times. The SI was used by Horwitz and Newsome (2001) to analyze spike trains recorded from SC. It analyzes local spike count fluctuations around the median across trials and assesses whether these fluctuations are completely random or whether they are systematically grouped (for details, see Materials and Methods). If the fluctuations are random, the SI is expected to follow a distribution with zero mean. This would, for example, be the case if each trial were an instantiation of the same non-homogeneous Poisson process. In contrast, if individual trials were characterized by a step-like transition from a low to a high firing rate, we would expect a distribution of SIs with a negative mean. In contrast to Horwitz and Newsome (2001), who calculated a single SI across all trials, we calculated an SI for each trial and analyzed the distribution across trials. This has the advantage of having access to a single-trial measure, and the mean of the SI distribution across trials is not expected to change with the number of available trials, which makes it easier to compare recording sessions with different numbers of trials.

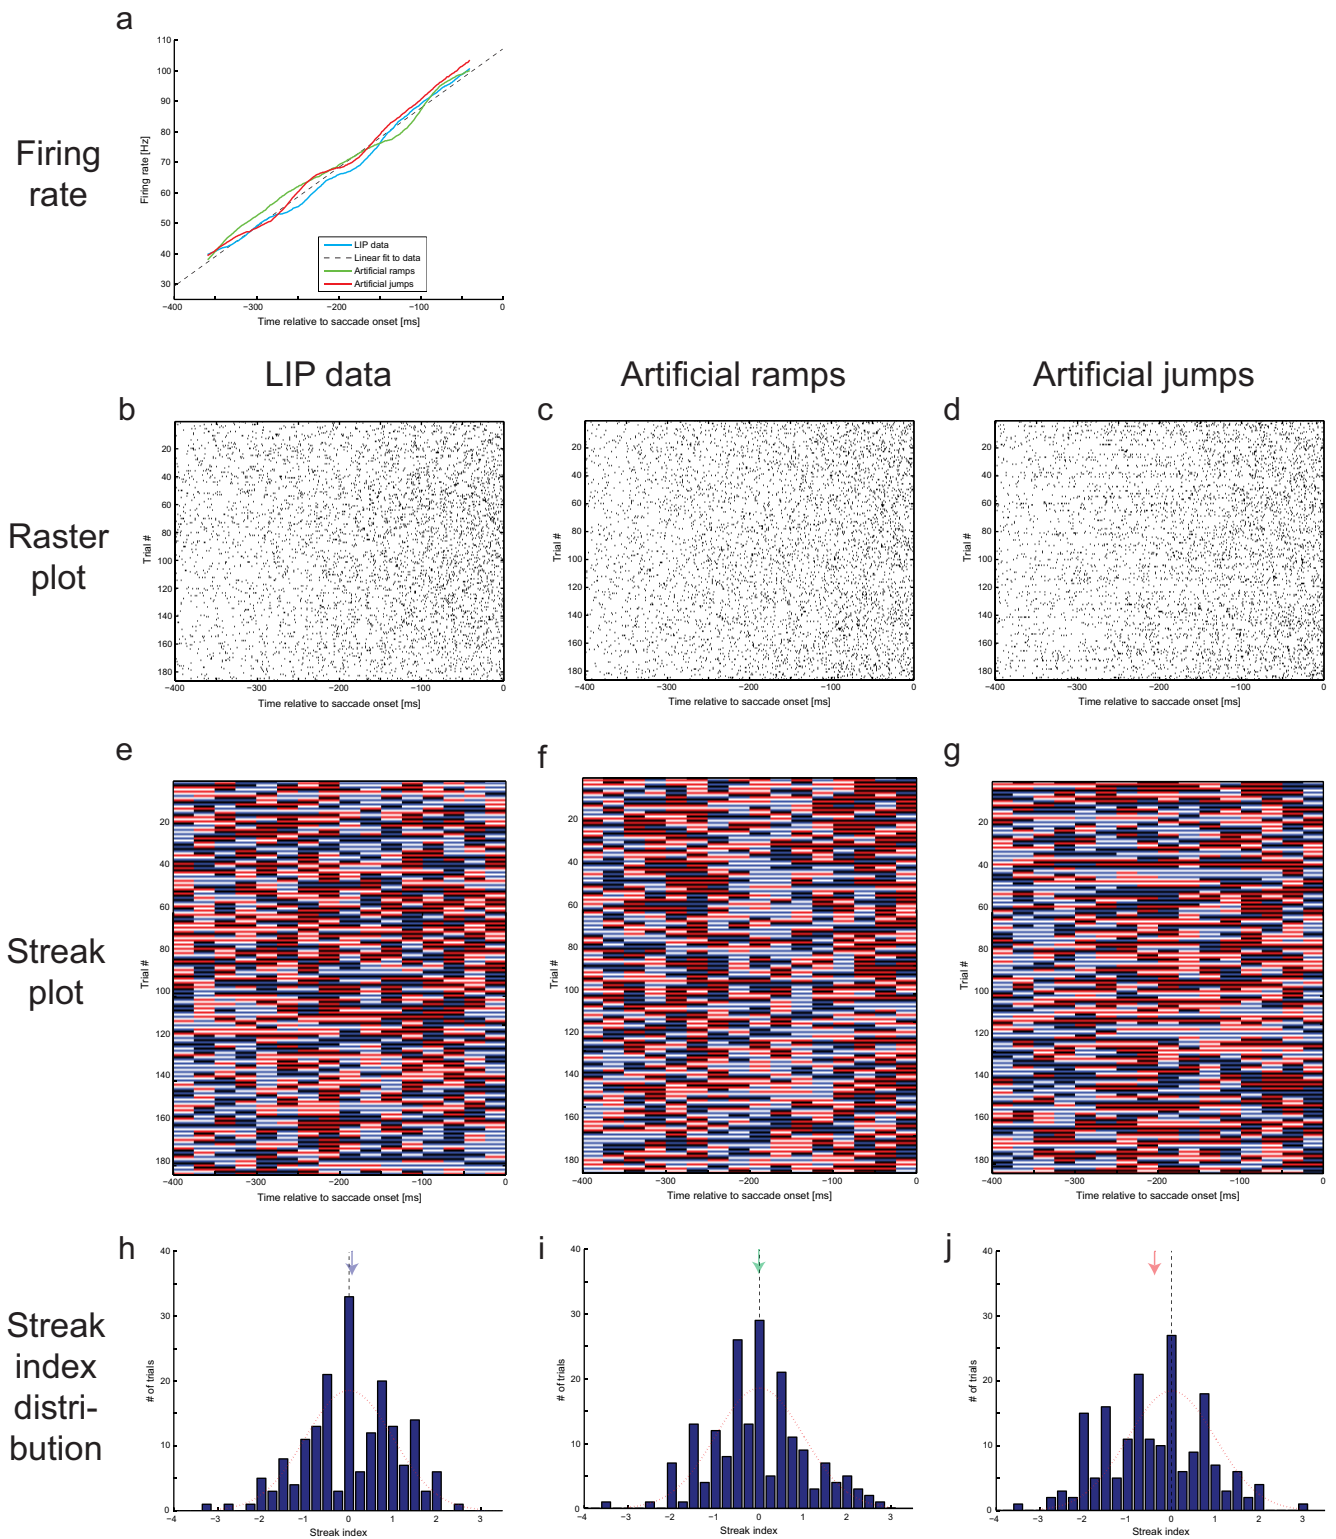
The SI analysis of an example session with three simultaneously recorded neurons is illustrated in Figure 2. The blue line in Figure 2*a* shows the average firing rate of the combined spike trains across trials. A linear fit (black dashed line) was used to estimate the initial (29 Hz) and the final (107 Hz) firing rate. Please note that these rates have not been divided by the number of contributing neurons. The raster plot of the real LIP data is shown in Figure 2*b*. The raster plots of the artificial ramping data and the artificial jumping data are shown in Figure 2, *c* and *d*. Note the somewhat streaky appearance of the artificial jump data. The green and red lines in Figure 2*a* verify that the average firing rates of both artificial datasets match the average firing rate of the real data. The streak plots in Figure 2*e–g* illustrate how the SI is obtained. Each trial is subdivided into 16 successive bins with a width of 25 ms, and the number of spikes in each of these bins is counted. The spike count is then compared with the median spike count in this particular bin across trials. A spike count that is smaller than the median is marked white or blue (on alternating trials), and a spike count that is larger than the median is marked black or red. If the spike count happened to be identical to the median, a color was randomly assigned. The SI is based on the number of color changes during a trial. Figure 2*i* shows the distribution of SIs across trials for the artificial ramp data. The distribution is expected to be similar to a normal distribution with zero mean and an SD of 1, as represented by the dotted red line. The actual mean of the distribution (indicated by the green arrow) was  $-0.02$ , which was not significantly different from 0 ( $p = 0.76$ ). Figure 2*j* shows the distribution of SIs for the artificial jump data with a mean of  $-0.39$  (red arrow), which was significantly different from 0 ( $p = 9 \times 10^{-6}$ ). The distribution of SIs for the real LIP data is shown in Figure 2*h* with a mean of  $+0.08$  (blue arrow; not significantly different from 0;  $p = 0.32$ ).

Before assessing the SIs of the real LIP data, we made sure that three different instantiations of both artificial datasets were classified appropriately, meaning that the mean SI for the artificial ramp data was not significantly different from 0 ( $p > 0.05$ ), whereas the artificial jump data had a negative mean SI that was significantly different from 0 ( $p < 0.05$ ). This criterion was met for 15 of the 18 analyzed sessions. The distributions of mean SIs across sessions for real LIP data (blue), artificial ramp data (green), and artificial jump data (red), using an analysis bin width

of 25 ms, as used previously by Horwitz and Newsome (2001), are shown in Figure 3*a*. Stars mark mean SIs that were significantly different from 0 ( $p < 0.05$ ). By virtue of our control criterion, this had to be the case for all of the artificial jump datasets and for none of the artificial ramp datasets. Nine of the real LIP datasets had a mean SI that was not significantly different from 0, but six of them had a positive mean SI that was significantly different from 0. Thus, a number of sessions had significantly more median crossings than would be expected from a random arrangement. We will return to this surprising observation in Discussion. The colored arrows indicate the medians of the across-session distributions, which were  $-0.38$  for the artificial jump data (significantly different from 0;  $p = 6 \times 10^{-5}$ ),  $+0.01$  for the artificial ramp data (not significantly different from 0;  $p = 0.79$ ), and  $+0.08$  for the real LIP data (not significantly different from 0;  $p = 0.30$ ). The SIs for the real data were significantly to the right of the midpoints of the SIs for the two artificial datasets ( $p = 0.0001$ ) and therefore closer to the SIs for the artificial ramps. However, the median SI for the real data was not only significantly larger than that for the artificial jump data ( $p = 6 \times 10^{-5}$ ) but also than that for the artificial ramp data ( $p = 0.05$ ). Potential reasons for this significant rightward shift, in particular one that prompted us to repeat the analysis with a different bin width, will be addressed later in Discussion.

The results of an SI analysis using a bin width of 40 ms (instead of 25 ms) are shown in Figure 3*b*. Sixteen of the 18 sessions met the above-mentioned inclusion criteria. Twelve of the real LIP datasets had a mean SI that was not significantly different from 0. Four sessions had a negative SI that was significantly different from 0 (marked with blue stars). The median SI for the artificial jump data was  $-0.35$  (significantly different from 0;  $p = 3 \times 10^{-5}$ ),  $+0.01$  for the artificial ramp data (not significantly different from 0;  $p > 0.99$ ), and  $-0.08$  for the real LIP data (not significantly different from 0;  $p = 0.21$ ). The real LIP data were significantly different from both types of artificial datasets ( $p = 0.03$  for the artificial ramps and  $p = 0.0005$  for the artificial jumps) and significantly to the right of the midpoints between the two artificial datasets ( $p = 0.05$ ). This is a more reasonable result because any variability in the rate function across trials should place the SI of the real data to the left of the SI for the artificial ramps (no variability in the rate function across trials). However, the results indicate that the real LIP data, at the population level, were more similar to artificial ramps than to artificial jumps.

To be able to make a more direct comparison with the results reported by Horwitz and Newsome (2001), which were based on SU recordings in the SC, we also attempted an analysis at the single-neuron level, following the same procedure as for the combined spike trains. This turned out to be more challenging, and only 10 of the 33 neurons that had contributed to the SI results at the population level (when using a bin width of 25 ms) had properties that allowed a reliable classification of the artificial datasets with matched properties. The results are shown in Figure 3*c*; stars again mark mean SIs that were significantly different from 0 ( $p < 0.05$ ). Five of the 10 analyzed neurons had a mean SI that was not significantly different from 0, one had a negative mean SI that was significantly different from 0, and four had a positive mean SI that was significantly different from 0. The colored arrows again indicate the medians of the across-neuron distributions, which were  $-0.27$  for the artificial jump data (significantly different from 0;  $p = 0.002$ ),  $-0.02$  for the artificial ramp data (not significantly different from 0;  $p = 0.75$ ), and  $-0.05$  for the real LIP data (not significantly different from 0;  $p = 0.75$ ). The median SIs for the real data and the artificial ramp data were not significantly



**Figure 2.** SI analysis at the population level: example session. *a*, Firing rates as a function of time: real LIP data (blue), linear fit used for extracting initial and final firing rates (black dashed line), artificial ramp data (green), and artificial jump data (red). *b–d*, Raster plots of the real LIP data (*b*), artificial ramp data (*c*), and artificial jump data (*d*). *e–g*, Streak plots of the real LIP data (*e*), artificial ramp data (*f*), and artificial jump data (*g*): spikes have been counted in 16 successive 25 ms bins. Each line reflects a trial, and the colors indicate whether the spike count was above (black or red) or below (white or blue) the median spike count across trials. The combinations black/white and red/blue are used on alternating trials. *h–j*, SI distributions across trials for the real LIP data (*h*), artificial ramp data (*i*), and artificial jump data (*j*). The colored arrows indicate the means of the distributions, the dashed vertical line marks zero, and the red dotted line shows a standard normal distribution.

different from each other ( $p = 0.56$ ), whereas there was a significant difference between the real data and the artificial jump data ( $p = 0.002$ ). The midpoint test, however, did not quite reach significance ( $p = 0.08$ ). The results for an analysis bin width of

40 ms are shown in Figure 3*d*. Eight of 35 neurons met the inclusion criteria. Six neurons had a mean SI that was not significantly different from 0, one was negative, and one was positive. The median SIs across neurons were  $-0.29$  for the artificial steps ( $p =$

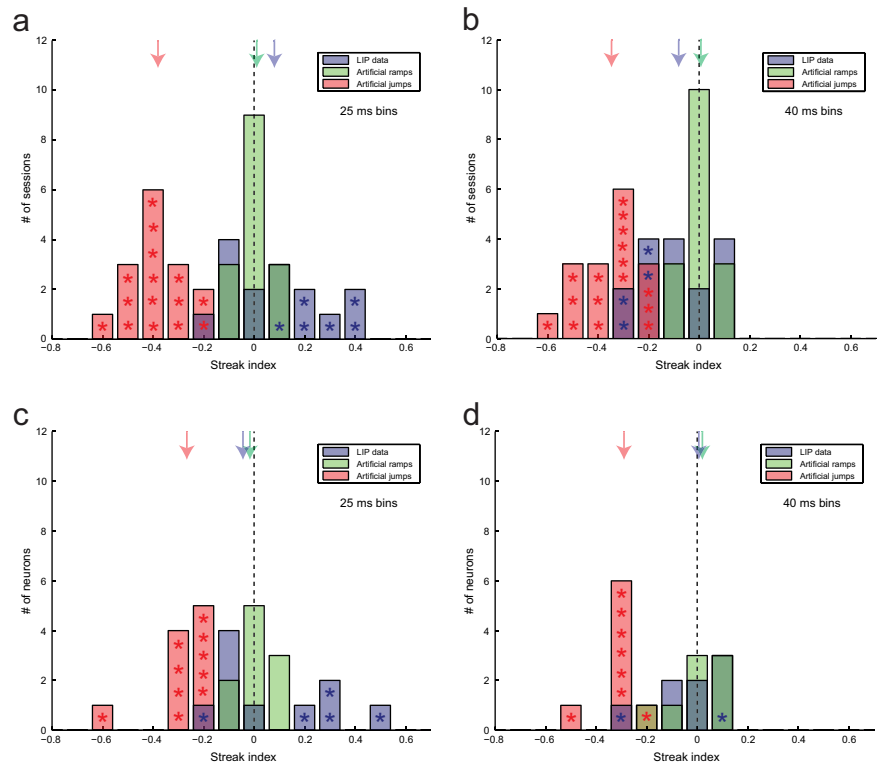


0.008), +0.02 for the artificial ramps ( $p = 0.45$ ), and +0.01 for the real LIP data ( $p > 0.99$ ). The median SI for the real data was not significantly different from the one for artificial ramps ( $p = 0.94$ ) but significantly different from the one for artificial jumps ( $p = 0.02$ ) and significantly to the right of the midpoint between the two artificial datasets ( $p = 0.05$ ). Thus, the results at the single-neuron level were similar to those found at the population level and suggest that both single LIP neurons with a trial-averaged ramping response as well as the population undergo gradual changes in firing rate during individual decisions.

### Global likelihood ratio between a linear model and a step model

Because the SI analysis still relies on a comparison of individual trial data with the data across trials and because it provides a rather indirect measure, we also aimed at a more direct test of competing models: a linear model, assuming that each spike train is an instantiation of the same non-homogeneous Poisson process with the rate changing linearly over time, and a step model, assuming that each spike train is an instantiation of a non-homogeneous Poisson process whose rate steps from a common initial rate to a common final rate at some random time (different across trials). Poisson models were chosen for computational tractability, because cortical spike trains are generally believed to be Poisson-like (Softky and Koch, 1993; Shadlen and Newsome, 1998). In contrast to Maimon and Assad (2009) who reported substantially increased regularity in LIP spiking, our data are mostly consistent with Poisson statistics. When fitting gamma distributions to the interspike interval distributions (without applying any rate correction), the median shape parameter (1.16) was indistinguishable ( $p = 0.55$ , Wilcoxon's test) from that predicted by our artificial ramp data based on Poisson spiking (1.17) when analyzing the single-trial combined LIP spike trains, which are used for all of our Poisson model comparisons. The value slightly exceeds 1 as a result of the non-stationarity. When analyzing our SU spike trains, the median shape parameter (1.35) was significantly ( $p = 0.004$ , Wilcoxon's test) but only slightly larger than that predicted by our artificial ramp data based on Poisson spiking (1.10). In contrast, Maimon and Assad reported a median (rate-corrected) shape parameter of 1.95 for their LIP neurons. Our data are much closer to what these authors reported for middle temporal area MT/medial superior temporal area MST (1.23). The use of Poisson-based models is therefore justified.

We set out to estimate the global likelihood ratio between the two non-homogeneous Poisson models. Ideally, the computation should be performed across all trials, but because the likelihoods have to be computed for all possible parameter combinations and because the possible combinations of step times explode as more and more trials are being added, we took a simplified approach to make the problem computationally tractable. First, we only used a 300 ms time window, ranging from

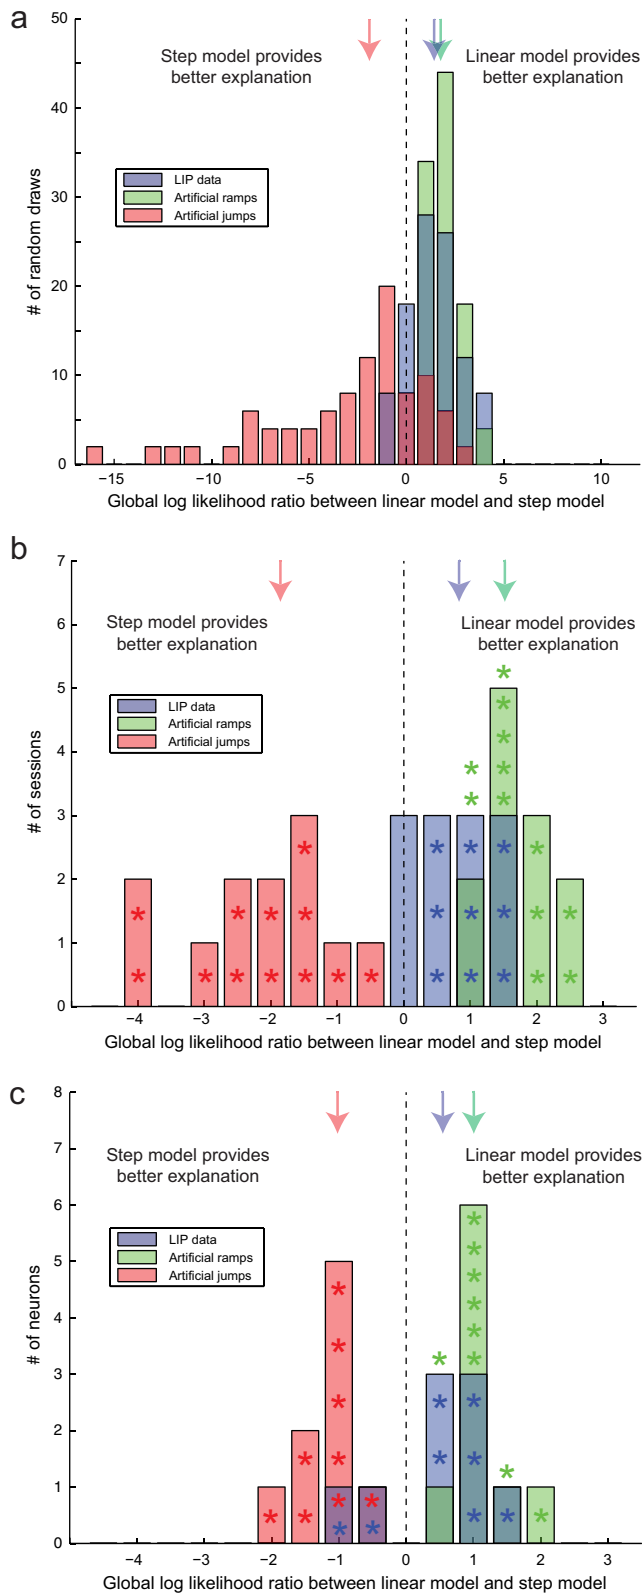


**Figure 3.** Mean streak indices across sessions/neurons. *a*, Distribution of mean SIs at the population level across experimental sessions using a bin width of 25 ms: real LIP data (blue), artificial ramp data (green), and artificial jump data (red). Stars mark mean SIs that were significantly different from zero ( $p < 0.05$ ). Colored arrows mark the medians of the distributions. *b*, Like *a* but using a bin width of 40 ms. *c*, Distribution of mean SIs at the single-neuron level across neurons (otherwise like *a*). *d*, Like *c* but using a bin width of 40 ms.

350 to 50 ms before saccade onset. Second, we only computed the global likelihood ratio for four trials at a time, but we repeated this calculation 100 times for each session. Thus, for each session, we made 100 random draws of four trials each from the real LIP data, from the artificial ramp data with matched initial and final firing rates, and from the artificial jump data with again matched initial and final firing rates and a uniform distribution of step times. Additional details can be found in Materials and Methods.

The results for one example session are shown in Figure 4*a*. The blue histogram shows the distribution of global LLRs across the 100 random draws for the real LIP data. The green histogram shows the same for the artificial ramp data and the red histogram for the artificial jump data. Negative values (to the left of the dashed vertical line) indicate that the step model provided the better explanation for the data. Positive values (to the right of the dashed vertical line) indicate that the linear model provided the better explanation for the data. Because the natural logarithm was used, a value of +2.3 (+4.6) would indicate that the data are 10 (100) times more likely according to the linear model compared with the step model. The colored arrows mark the medians of the distributions: +1.39 for the real LIP data (significantly different from 0;  $p = 4 \times 10^{-8}$ ), +1.77 for the artificial ramp data ( $p = 4 \times 10^{-23}$ ), and -1.88 for the artificial jump data ( $p = 4 \times 10^{-8}$ ). Thus, in this case, both the real LIP data and the artificial ramp data were better explained by the linear model, whereas the artificial jump data were better explained by the step model.

When analyzing the data across sessions, we only included sessions with a positive median global LLR that was significantly different from 0 ( $p < 0.05$ ) for the artificial ramp data and a negative median global LLR that was also significantly different

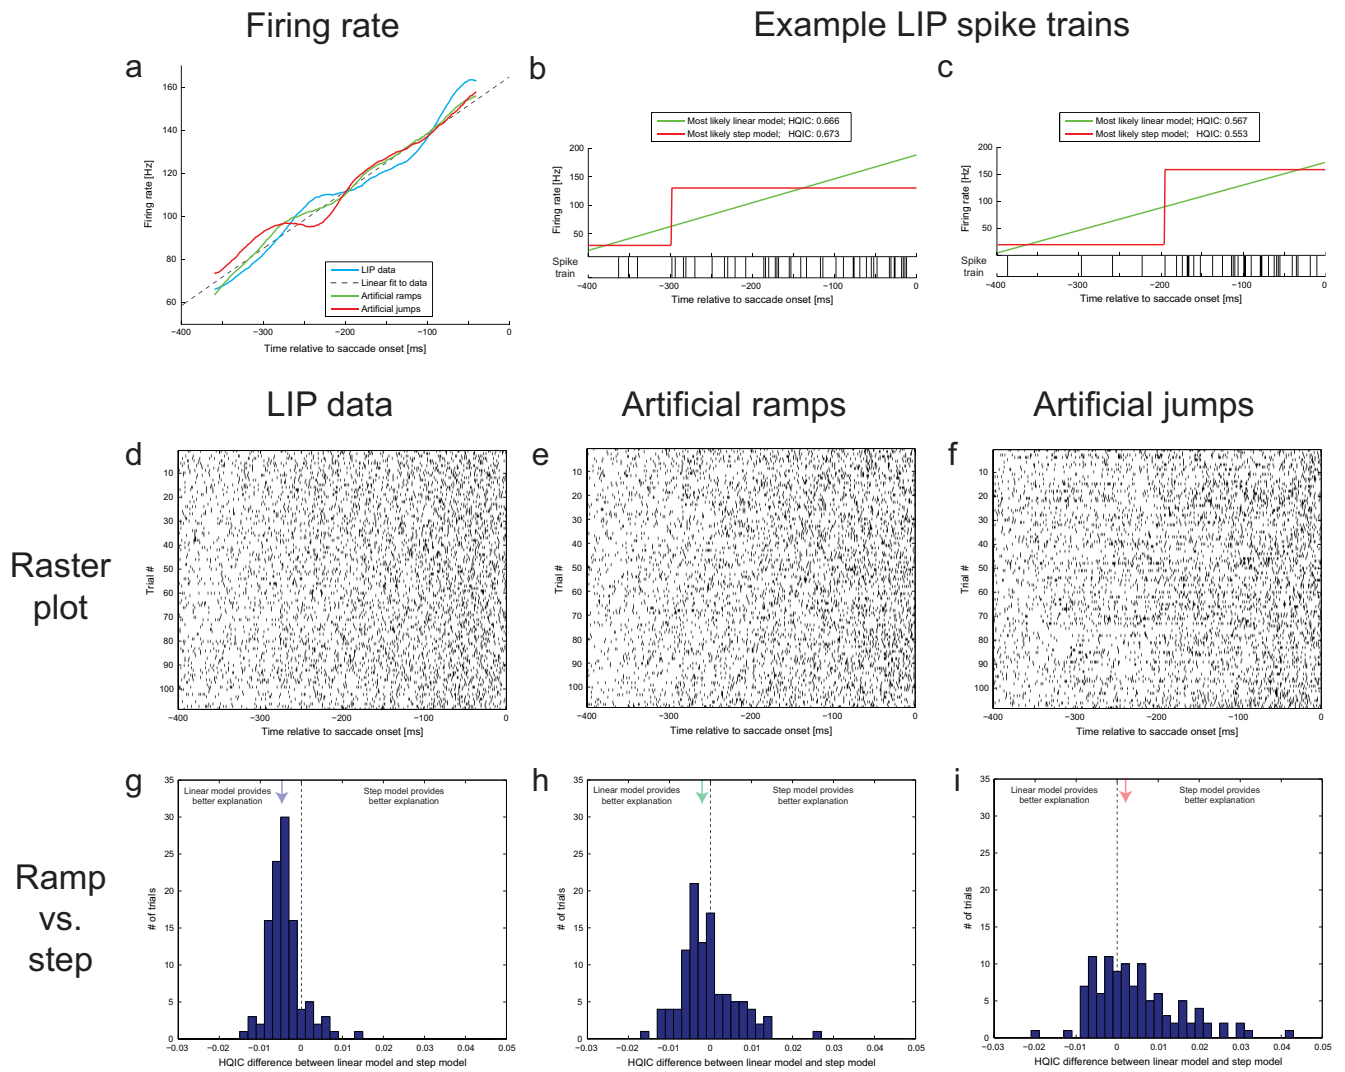


**Figure 4.** Global likelihood ratio between a linear model and a step model. **a**, Example session. One hundred random draws of four trials each were analyzed. The histograms show the distributions of global LLRs between the linear and the step model for the real LIP data (blue), the artificial ramp data (green), and the artificial jump data (red). Negative values (to the left of the dashed vertical line) indicate that the step model provided the better explanation. Positive values (to the right of the dashed vertical line) indicate that the linear model provided the better explanation. The colored arrows indicate the medians of the distributions. **b**, Median global LLRs at the population level across experimental sessions: real LIP data (blue), artificial

from 0 for the artificial jump data to ensure that our analysis method was sensitive enough to discriminate between both models given the initial and final firing rates of this particular session. Twelve of the 18 analyzed sessions met this criterion. Of those, nine had a positive median global LLR that was significantly different from 0 ( $p < 0.05$ ) for the real LIP data, indicating that the linear model provided the better explanation for the single-trial combined LIP spike trains. The LIP data from the remaining three sessions had a median global LLR that was not significantly different from 0, indicating that the data were inconclusive and neither better explained by the linear model nor by the step model. The results are shown in Figure 4*b*. The stars mark median global LLRs that were significantly different from 0. The colored arrows mark the medians of the global LLR distributions across sessions. Those were +0.83 for the real LIP data (significantly different from 0;  $p = 0.006$ ), +1.52 for the artificial ramp data ( $p = 0.0005$ ), and -1.86 for the artificial jump data ( $p = 0.0005$ ). The results for the real data were significantly to the right of the midpoints between the two artificial datasets ( $p = 0.005$ ). Thus, overall, the LIP spike trains were better explained by the linear model, which is consistent with the result of the SI analysis. However, in addition to the distribution for the real LIP data being different from the distribution for the artificial jump data ( $p = 0.0005$ ), the distribution for the real LIP data and the distribution for the artificial ramp data were also different from each other ( $p = 0.001$ ). The real data were therefore more ramp-like than jump-like but not indistinguishable from the artificial ramp data. This, however, is not too surprising. Even in the case of the LIP activity being the result of a perfect integration of sensory evidence, the resulting activity profiles would not be expected to be perfectly linear ramps attributable to fluctuations in the momentary evidence and potentially attributable to integration of some noise. Another possibility, and we will see some evidence for this later, is that some firing rate transitions are indeed jump-like. However, those would clearly have to be the minority.

We also repeated this analysis at the single-neuron level, which again proved more challenging: only 9 of the 26 neurons that had contributed to the results at the population level had properties that allowed a reliable classification of the artificial ramp and jump data. The results are shown in Figure 4*c*. Six LIP neurons were more consistent with the linear model, and two were more consistent with the step model. The median global LLR (+0.55) was not significantly different from 0 ( $p = 0.18$ ), not significantly different from that of the artificial ramp data (+1.00;  $p = 0.07$ ) but significantly different from that of the artificial jump data (-1.02;  $p = 0.004$ ). The midpoint test, however, did not reach significance ( $p = 0.13$ ), and the results therefore do not provide a clear answer whether the SU responses are closer to ramping or stepping responses. The absolute values of the median global LLRs associated with the artificial data are smaller than the ones that had been observed for the combined spike trains (+1.52 and -1.86), indicating that having access to at least two simultaneously recorded neurons provides increased statistical power for discriminating between models. The results at the SU level are therefore not as clear cut as at the population level.

←  
ramp data (green), and artificial jump data (red). Median global LLRs that were significantly different from 0 ( $p < 0.05$ ) are marked with stars. The medians of the distributions are indicated by the colored arrows. **c**, Median global LLRs at the single-neuron level across neurons: same format as **b**.



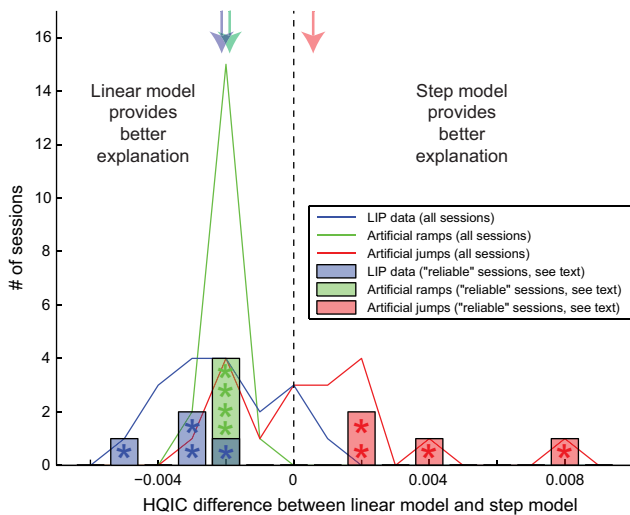
**Figure 5.** Maximum likelihood fits of non-homogeneous Poisson processes: example session. *a*, Firing rates as a function of time: real LIP data (blue), linear fit used for extracting initial and final firing rates (for generating the artificial datasets; black dashed line), artificial ramp data (green), and artificial jump data (red). *b, c*, Two real example LIP spike trains and the associated maximum likelihood estimates of a ramping non-homogeneous Poisson process (green) and a jumping non-homogeneous Poisson process (red). The example on the left (*b*) is better explained by the linear model (as indicated by the smaller HQIC value), and the example on the right (*c*) is better explained by the step model. *d–f*, Raster plots of the real LIP data (*d*), the artificial ramp data (*e*), and the artificial jump data (*f*). *g–i*, Distributions of HQIC differences across trials for the real LIP data (*g*), the artificial ramp data (*h*), and the artificial jump data (*i*). Negative differences (to the left of the dashed vertical line) indicate that the linear model provided the better explanation, and positive differences (to the right of the dashed vertical line) indicate that the step model provided the better explanation. The colored arrows indicate the medians of the distributions.

**Maximum likelihood fits of non-homogeneous Poisson processes**

The models used for the global likelihood ratio analysis still imposed a common firing rate constraint. Given that the underlying rates might vary across trials, we also aimed at a more independent classification of the shape of the firing rate function underlying individual spike trains. To this end, we subjected each single-trial combined spike train to maximum likelihood fits of two different non-homogeneous Poisson models: one assuming a linear change of the firing rate over time, the other one assuming a firing rate step. Initial and final firing rates as well as the step time were free parameters in each single fit. Thus, only the shape of the underlying rate function was constrained. It turned out that, even when working with artificial ramp data, direct model comparisons based on the maximum likelihood were heavily biased toward the step model because of its extra degree of freedom (three instead of two parameters). We therefore had to introduce a penalty term and used the HQIC (Hannan and Quinn, 1979),

which turned out to be less biased when classifying the artificial datasets than the more commonly used AIC or BIC.

The general approach was similar to the SI analysis. An example session with two simultaneously recorded neurons is shown in Figure 5. Figure 5*b* shows one example of a combined LIP spike train and the two best-fitting non-homogeneous Poisson models (green, linear; red, step). In this case, the linear model had the lower HQIC, indicating that it provided the better explanation of the data. An example spike train that was better explained by the step model is shown in Figure 5*c*. The blue line in Figure 5*a* shows the average firing rate of the LIP data (combined spike trains). A linear fit (black dashed line) was again used to estimate the initial (59 Hz) and the final (165 Hz) firing rate. Note that these rates have again not been divided by the number of contributing neurons. The green and red lines show the average firing rates of the two artificial datasets. The raster plot of the real data is shown in Figure 5*d*. Figure 5, *e* and *f*, shows the raster plots of the artificial ramp and jump data, respectively. The distribution of HQIC dif-



**Figure 6.** Maximum likelihood fits of non-homogeneous Poisson processes: median HQIC differences across sessions. Real LIP data (blue), artificial ramp data (green), and artificial jump data (red). The bars illustrate the results for the “reliable” sessions (see Results). Stars indicate median HQIC differences that were significantly different from 0 ( $p < 0.05$ ). Negative values (to the left of the dashed vertical line) indicate that the linear model provided the better explanation, and positive values (to the right of the dashed vertical line) indicate that the step model provided the better explanation. The lines reflect the results across all sessions. The colored arrows mark the medians of these latter distributions.

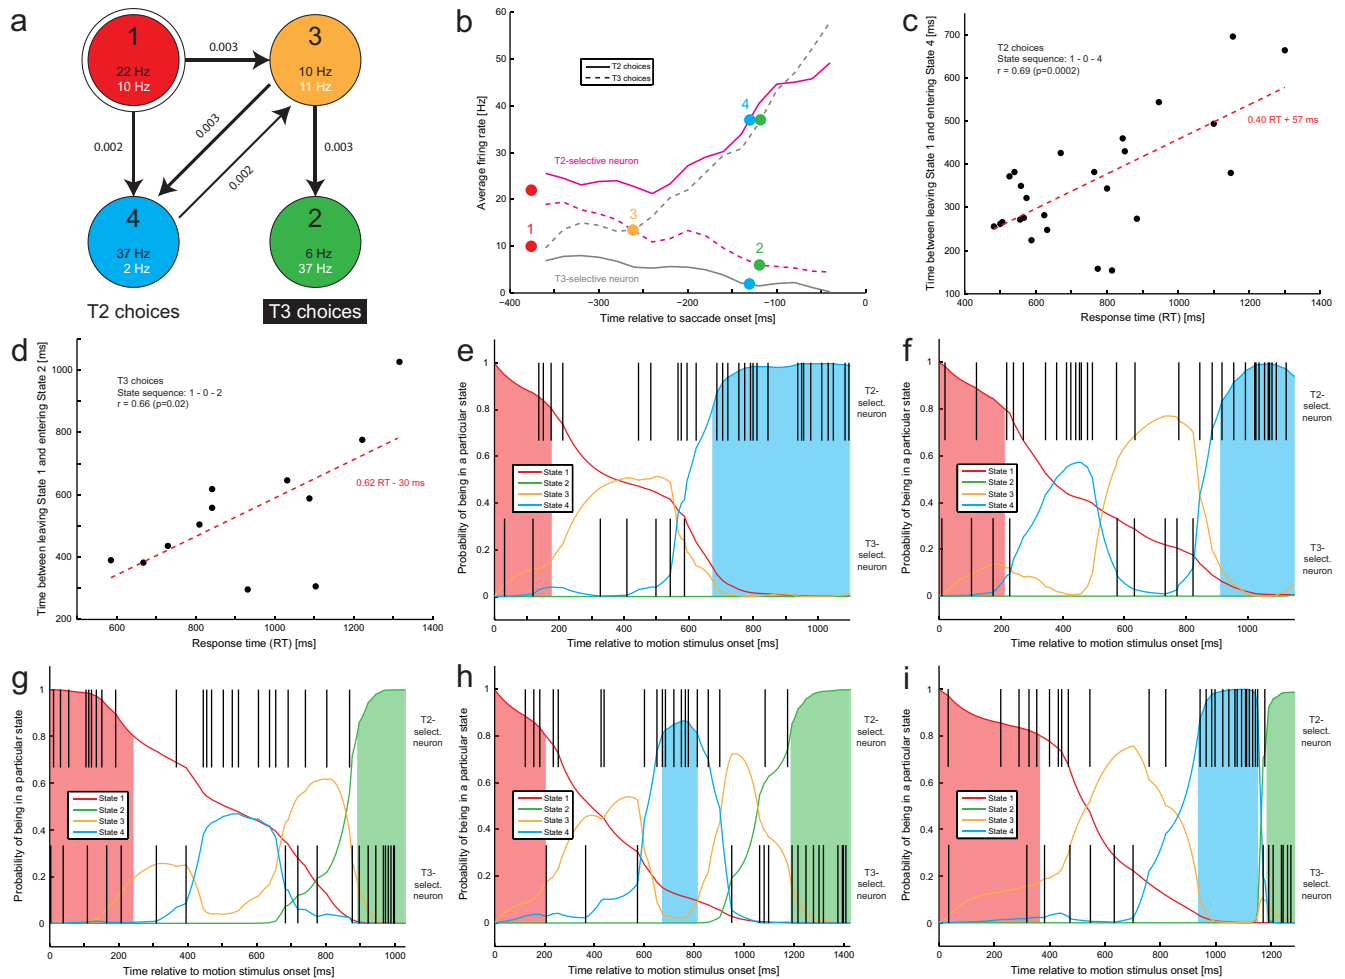
ferences between the linear and the step model across trials for the artificial ramp data is shown in Figure 5*h*. This dataset should be better explained by the linear model because it has been generated according to a linear model. The median HQIC difference was  $-0.0020$  and significantly different from 0 ( $p = 0.04$ ). A negative HQIC difference indicates that the linear model provides the better explanation. An HQIC difference of  $-0.0035$  ( $-0.012$ ) would indicate that the data are 2 (10) times more likely according to the linear model compared with the step model. Figure 5*i* shows the distribution of HQIC differences for the artificial jump data. In this case, the step model should be providing the better explanation. The median HQIC difference was  $+0.0020$  (significantly different from 0;  $p = 0.02$ ). The distribution of HQIC differences for the real LIP data is shown in Figure 5*g*. The median was  $-0.0046$  ( $p = 7 \times 10^{-15}$ ), indicating that, on average, the LIP spike trains were better explained by the linear model than by the step model.

The direct model comparison based on maximum likelihood fits turned out to be more challenging than the previous analyses. Again, we required that three instantiations of both artificial datasets were classified correctly, meaning that the median HQIC difference had to be negative and significantly different from 0 ( $p < 0.05$ ) for the artificial ramp data and positive and significantly different from 0 ( $p < 0.05$ ) for the artificial jump data. Our reasoning was that, if the analysis method cannot reliably classify the artificial datasets that were matched in number of trials and firing rates to the real data but were otherwise perfect implementations of linear ramps and steps, respectively, we cannot expect it to provide a reliable classification of the real data. Only four of our sessions had properties (number of trials and combination of initial and final firing rates) that allowed a reliable classification of the artificial datasets. The results for these sessions are shown by the bar graphs in Figure 6 (“reliable” sessions). Like in the case of the artificial ramp data, but unlike the artificial jump data, the median HQIC differences for the real LIP data were all negative and significantly different from 0 ( $p$  values ranging from  $7 \times$

$10^{-15}$  to 0.0005; indicated by stars). For the purpose of statistical comparisons, we also looked at the distributions across all sessions. These are shown by the line plots in Figure 6. The colored arrows mark the medians of these distributions. The median HQIC differences were  $-0.0021$  for the real LIP data (significantly different from 0;  $p = 0.0001$ ),  $-0.0019$  for the artificial ramps (significantly different from 0;  $p = 8 \times 10^{-6}$ ), and  $+0.0006$  for the artificial jumps (not significantly different from 0;  $p = 0.81$ ). This indicates that the HQIC was still slightly biased, which probably explains the small number of sessions that passed the inclusion criteria for reliable sessions and therefore shows the importance of using artificial control data for verifying the classification performance. The median HQIC difference for the real LIP data was not significantly different from that for the artificial ramps ( $p = 0.53$ ) but significantly different from that for the artificial jumps ( $p = 0.009$ ) and significantly to the left of the median midpoint between the two artificial datasets ( $p = 0.03$ ). This demonstrates that, overall, even without rate restrictions, the linear ramp model provides a better explanation of the LIP data than the step model. The result is therefore consistent with the outcome of the previous analyses. All of them suggest that, given a trial-averaged ramping response, the majority of single-trial LIP spike trains are also more consistent with a gradual than a sudden firing rate change.

**HMM analysis and changes of mind**

The analyses that have been discussed so far were limited to simultaneous recordings from neurons with the same choice selectivity. Furthermore, the recorded neurons had to be part of the winning pool, and only the last 400 ms of each trial were analyzed. To be able to also analyze simultaneous recordings from neurons with different choice selectivity (different targets in their RFs) and trials resulting in different choices and with different lengths, we made use of HMMs. The original idea behind HMM analysis is that a neural system can be described by a limited set of discrete states. Although these states cannot be directly observed, the spike pattern that is emitted by the neurons is modulated by the current state. In particular, the HMM used here assumes that each state is associated with a set of firing rates for each of the recorded neurons and that the neurons generate Poisson spike trains with these rates. Each state is further associated with a set of transition probabilities to other possible states. These transition probabilities are not influenced by the history of how the system arrived at the current state. To be able to apply HMM analysis, each single-trial set of spike trains has to be converted into an emission sequence. This process is illustrated in Figure 1*h*. The whole trial (from motion stimulus onset to saccade onset) was subdivided into 2 ms bins. If none of the recorded neurons had spiked in a given bin, the emitted symbol was set to 0. If exactly one of the neurons had spiked in a given bin the emitted symbol was set to the ordinal number of the neuron that had spiked. If multiple neurons had spiked in the same bin, one of the spiking neurons was picked randomly. The bin width of 2 ms was chosen to achieve a good compromise between length of the emission sequences (shorter is better because of faster model identification) and losing spikes because of “simultaneous” firing. The emission sequences are then used to find the HMM parameters (a set of firing rates for each state as well as a matrix of state transition probabilities) using the Baum–Welch algorithm. This algorithm assumes that the number of states in the model is known. Estimating the number of states in an HMM, the so-called “order estimation problem” is a difficult problem. In previous applications of HMMs to neural data, the model order was picked more



**Figure 7.** HMM analysis: first example session. **a**, Identified HMM with four states based on a simultaneous recording from two LIP neurons with different choice selectivity. Each state is associated with a state number (top), a color (red/green/orange/blue), and two firing rates (black, FR of T2-selective neuron; white, FR of T3-selective neuron). Arrows indicate state transitions with probabilities of at least 0.001 (when rounded to 3 decimals). Thicker arrows reflect higher transition probabilities. State 1 (top left, double circle) is the initial state on each trial. The two states at the bottom are associated with T2 and T3 choices, respectively. **b**, Approximate correspondence between states and trial-averaged firing rates: average firing rates of the T2-selective (magenta) and the T3-selective (gray) neuron for T2 (solid lines) and T3 (dashed lines) choices. The colored, filled circles with corresponding numbers show how the HMM states can be approximately mapped onto these average firing rate responses. **c**, Timing of trials with a 1–0–4 state sequence: significant correlation between the transition duration and RT. The dashed red line shows the linear regression. **d**, Timing of trials with a 1–0–2 state sequence: significant correlation between the transition duration and RT. The dashed red line again shows the linear regression. **e–i**, Examples of decoded individual trials. The top row shows the single-trial spike train of the T2-selective neuron, and the bottom row shows the spike train of the T3-selective neuron. The colored lines indicate the probabilities of being in any of the four identified states of the HMM. Probabilities exceeding 0.8 are marked with shaded areas. We refer to these time intervals as the system being in one of the discrete states. White time intervals (without any colored shading) are referred to as state transitions. **e, f**, Examples of T2 choices with a 1–0–4 state sequence. **g**, Example of a T3 choice with a 1–0–2 state sequence. **h, i**, Examples of T3 choices involving a change of mind: the system spent some time in State 4, the state associated with T2 choices, before switching to State 2. **h** shows an example of a slow change of mind, whereas **i** shows an example of a fast change of mind.

or less arbitrarily (Radons et al., 1994; Abeles et al., 1995; Seidemann et al., 1996; Gat et al., 1997; Jones et al., 2007; Miller and Katz, 2010). We used a more data-driven approach, which is explained in detail in Materials and Methods.

The model parameters were inferred from the full set of trials (all possible choices). Figure 7a shows the HMM for one example session, a simultaneous recording from two neurons with different choice selectivity (76 trials). This is the session that required the smallest number of states (four), and it is therefore particularly straightforward to interpret. The HMM always starts in State 1 (top left, double circle). Possible state transitions with probabilities of at least 0.001 (rounded to three decimals) are indicated by arrows. Thicker arrows indicate higher transition probabilities, the numbers next to the arrows. The two numbers at the bottom of each state are the associated firing rates of the two neurons. The black number shows the firing rate of the neuron that was selective for choosing

Target 2 (“T2 choice”), the white number the firing rate of the neuron that was selective for choosing Target 3 (“T3 choice”). As it turns out, the states can be approximately mapped onto trial-averaged firing rates. Figure 7b shows the average firing rate of both neurons (magenta, T2-selective neuron; gray, T3-selective neuron) for both T2 (solid lines) and T3 choices (dashed lines), aligned with the time of saccade onset. State 1 is the initial state and therefore reflects the firing rates of both neurons at the beginning of the trial before the rates split according to choice. State 2 approximately maps onto the average firing rates of both neurons ~120 ms before a saccade indicating a T3 choice. We therefore call it the state associated with T3 choices. State 3 approximately maps onto the average firing rates ~260 ms before a saccade to T3. State 4 approximately maps onto the average firing rates ~130 ms before a saccade indicating a T2 choice, and we therefore call it the state associated with T2 choices.

Once the HMM has been specified, the spike trains that have been observed in individual trials can be decoded to obtain the most likely underlying state sequence. For each time bin, a set of four posterior state probabilities can be obtained, indicating how likely the system was in a particular state at a particular time when emitting this sequence. Similar to previous work (Seidemann et al., 1996; Jones et al., 2007; Miller and Katz, 2010), we required one of these probabilities to exceed 0.8 to accept that the system was actually in one of the identified states. All of the probabilities being smaller than 0.8 was interpreted as the system currently undergoing a state transition. If the system were well described by the HMM, the state transitions should be very fast (see previous work, but see also Discussion). Conversely, if the system does not jump from discrete state to discrete state but its activity rather undergoes gradual changes, we would expect to see long transition periods when decoding the spike trains. We determined the state sequence for each trial according to this rule. More than 70% of the T2 choices were characterized by the sequence 1–0–4, with 0 indicating that the system was not in a defined state for some time after leaving State 1 and before entering State 4. Figure 7c shows the time between leaving State 1 and entering State 4 as a function of the RT of each trial. The time between states was positively correlated with RT ( $r = 0.69$ ;  $p = 0.0002$ ) and, according to a linear regression (see figure), typically accounted for >40% of the RT. Thus, the spike trains are typically not well described by sudden transitions between discrete states but rather very gradual transitions. The majority of T3 choices were characterized by the sequence 1–0–2. The time between leaving State 1 and entering State 2 is shown in Figure 7d. It was again positively correlated with RT ( $r = 0.66$ ;  $p = 0.02$ ) and, according to a linear regression (see figure), typically accounted for more than half of the RT, thus again an indication of a gradual change in activity rather than a sudden state transition. We repeated the same analysis for an artificial dataset with sudden firing rate jumps (see section on the first change-of-mind control analysis below for additional details on how these data were generated and matched to the real LIP data). We did not find any significant correlations between RT and the time between states for the artificial jump data ( $p \geq 0.17$ ), and the time between states typically accounted for <20% of the RT.

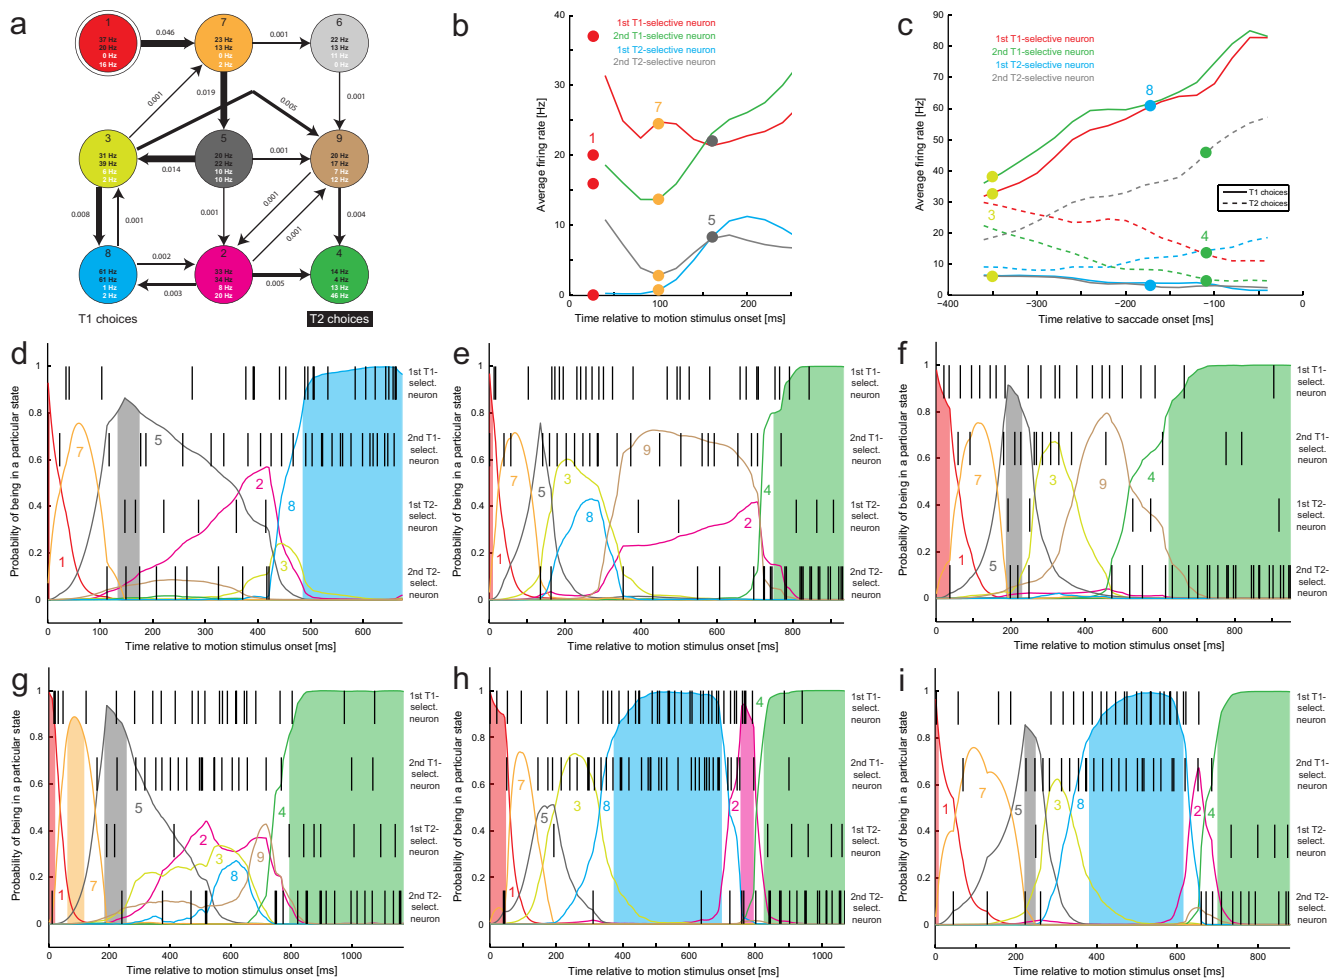
Several example trials are shown in Figure 7e–i. Figure 7e is a typical example of a T2 choice with a 1–0–4 sequence. Time ranges from motion stimulus onset (0) to the RT. The trial ends with saccade onset. The top spike train is from the T2-selective neuron, and the bottom spike train is from the T3-selective neuron. The four colored lines show the posterior probabilities of being in any of the HMM states. Time periods during which one of the probabilities exceeded 0.8 are marked with colored shading. The color corresponds to the currently active state (Fig. 7a). In this example, the system stayed in State 1 for ~180 ms and was in a transitional/undefined state for >400 ms (characterized by a continuous, gradual decline of the probability of being in the initial state), before entering State 4 ~400 ms before a saccade to T2. A similar example is illustrated in Figure 7f. The difference here is that, initially, the system was moving fast toward a T2 choice, indicated by the earlier increase in the probability of being in State 4, but returned to a more neutral state (higher probability of being in State 3), before finally entering State 4. Figure 7g illustrates a typical T3 choice with a state sequence of 1–0–2.

What makes this method attractive is that changes of mind can be identified on a single-trial basis. Figure 7h shows such an example. The system had entered State 4 ~700 ms into the trial, meaning that it was close to making a T2 choice, but then left this

state again ~800 ms into the trial, spent ~400 ms in an undecided state, before it finally entered State 2, initiating a saccade to T3. Even more interesting, Figure 7i shows an example of a rapid change of mind. In this case, the probability of being in State 4 had already approached 1, meaning that the system was extremely close to making a T2 choice, but it suddenly flipped into State 2 and a saccade to T3 was initiated. The transition period was only 30 ms, indicating that the system is apparently capable of fast state transitions.

The HMM analysis of a second recording session, this time a simultaneous recording from four decision-related LIP neurons, two of them selective for one of the choices and the other two for one of the other choices (312 trials), is shown in Figure 8. Figure 8a shows the structure of the HMM, which had nine states. The system always starts in State 1 (top left, double circle). The four numbers at the bottom of each state indicate the associated firing rates of the four recorded neurons, with the two black numbers reflecting the T1-selective neurons and the two white numbers reflecting the T2-selective neurons. Arrows indicate possible state transitions as in Figure 7a. As in the previous example, many of the states can be mapped onto trial-averaged firing rates. Figure 8b shows the average firing rates of the four neurons during the first 250 ms after motion stimulus onset. The state sequence 1–7–5 approximately reflects the stereotyped (stimulus-independent) response during the first ~200 ms after motion stimulus onset (Roitman and Shadlen, 2002; Churchland et al., 2008; Bollimunta and Ditterich, 2012). The average firing rates as a function of time relative to saccade onset and for both T1 (solid lines) and T2 choices (dashed lines) are shown in Figure 8c. Each line color (red, green, blue, gray) stands for one of the four recorded neurons. As can be seen, State 3 approximately maps onto the average activity ~350 ms before a saccade indicating a T1 choice. State 8 has an activity pattern that matches the average firing rates ~170 ms before a saccade to T1. We therefore call it the state associated with T1 choices. State 4 maps onto the average activity pattern ~110 ms before a saccade to T2. We therefore call it the state associated with T2 choices.

The model can again be used to decode individual trials. The majority of T1 choices were characterized by state sequences of 1–0–8, 1–0–7–0–5–0–8, 1–0–7–0–8, or 1–0–5–0–8. Figure 8d shows a typical example of a 1–0–5–0–8 sequence. The top two spike trains reflect the activity of the T1-selective neurons, and the bottom two spike trains reflect the activity of the T2-selective neurons. Again, typical decision trials were characterized by spending a substantial amount of time (several hundred milliseconds) in a transitional/undefined state. Typical T2 choices were characterized by state sequences of either 1–0–4, 1–0–9–0–4, 1–0–5–0–4, 1–0–7–0–9–0–4, or 1–0–7–0–4. Figure 8e shows an example of a 1–0–4 sequence with >700 ms between leaving State 1 and entering State 4. Figure 8f illustrates an example of a 1–0–5–0–4 sequence, and Figure 8g shows an example of a 1–0–7–0–5–0–4 sequence. All of these again have in common that the system spends a considerable amount of time (several hundred milliseconds) in a transitional/undefined state. Like in the case of the previous session, we can also look at changes of mind on a single-trial basis. Figure 8h shows an example of a trial in which the system had spent 300 ms in State 8, meaning that it was already close to making a T1 choice but ended up switching to State 4, and the trial ended with a T2 choice. Approximately 120 ms passed between leaving State 8 and entering State 4, and the system passed through State 2, a state characterized by all neurons having intermediate firing rates. Another change-of-mind example is illustrated in Figure 8i. Again, the system had



**Figure 8.** HMM analysis: second example session. **a**, Identified HMM with nine states based on a simultaneous recording from four LIP neurons. Two neurons were T1 selective (black FRs), and two were T2-selective (white FRs). State 1 (top left) is again the initial state, and States 8 and 4 at the bottom are associated with T1 and T2 choices, respectively. Otherwise, the conventions are the same as in Figure 7a. **b**, Approximate mapping between the average firing rates (colored lines) during the first 200 ms after motion stimulus onset and States 1, 7, and 5 (colored, filled circles with associated numbers). **c**, Approximate mapping between trial-averaged firing rates of all four neurons (red/green/blue/gray lines) for T1 (solid lines) and T2 (dashed lines) choices and States 3, 4, and 8 (colored, filled circles with associated numbers). **d–i**, Examples of decoded individual trials (same conventions as in Fig. 7e–j). **d**, Example of a T1 choice. **e–g**, Examples of T2 choices with different state sequences. **h, i**, Examples of T2 choices with a change of mind: the system spent some time in State 8, the state associated with T1 choices, before switching to State 4.

spent already 200 ms in State 8 when it switched to State 4. In this case, ~90 ms passed between leaving State 8 and entering State 4.

To obtain a better picture how often and under what circumstances changes of mind occurred, we analyzed a total of six sessions with two to four simultaneously recorded neurons that were selective for two different choices. These recordings are extremely difficult to obtain, but they provided us with a total of 2062 decoded individual decisions. A total of 63.7% of the state sequences ended in the state that was the most frequent final state for this particular choice, which we call the state that is associated with this choice. A total of 35.6% ended in an undefined state, meaning that none of the posterior probabilities exceeded 0.8, but in 83% of these cases, the associated state was the most likely one. Thus, 93.2% of all decoded trials ended with the associated state being the most likely one. A total of 0.5% of the trials ended in a state that was not associated with any of the two choices represented by the recorded neurons, and only three trials (0.1%) ended in the state that was associated with the other choice. Thus, it was extremely rare that the final state suggested a choice that was different from the one that the monkey had actually made. Fifty-one (2.5%) of the 2062 analyzed trials were change-of-mind trials, meaning that both states associated with choosing the tar-

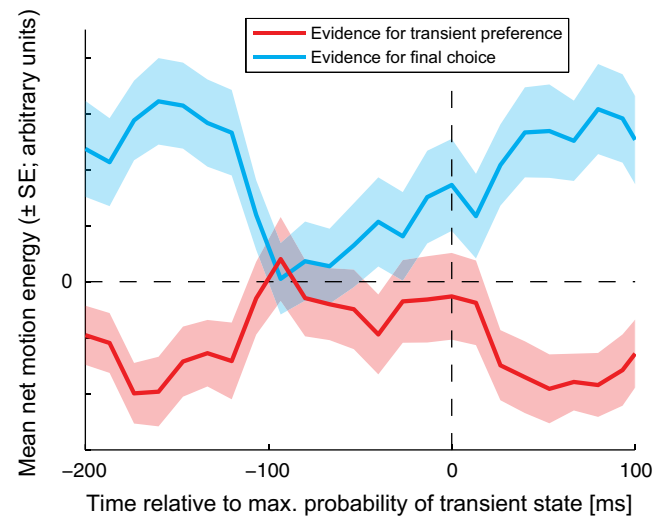
gets in the RFs of the recorded neurons were members of the decoded state sequence. The average net motion strength on these trials was 9.4%, which is significantly smaller than the average net motion strength on all trials (14.7%;  $p = 0.001$ , Wilcoxon’s test), indicating that changes of mind were more likely to occur during difficult discriminations than during easy discriminations. Twenty-seven of these trials (53%) led to an improved outcome because the monkey switched from a wrong choice to the correct choice, only five (10%) led to a worse outcome because of a switch from the correct choice to a wrong choice, and 19 switches (37%) had no impact on the outcome of the trial because it was either a switch from one wrong choice to the other wrong choice (eight trials) or because the net motion strength was 0 (11 trials), meaning that all targets were equally likely to be rewarded. An improvement in the outcome of 27 of the 40 change-of-mind trials with non-zero net motion strength is significantly more than what would be expected by chance ( $p = 2 \times 10^{-5}$ ,  $\chi^2$  test). If one were to construct random state sequences with two states that are associated with two different choices in them, the expectation would be that in one-third of them a state associated with the correct choice should be followed by a state associated with a wrong choice, in one-third of them both states

should be associated with wrong choices, and in one-third of them a state associated with a wrong choice should be followed by a state associated with the correct choice, leading to an improved outcome. We also analyzed the RT of change-of-mind trials and found that 90% of them had an RT that was larger than the median RT of trials with this particular combination of choice and net coherence. Thus, changes of mind tended to be associated with significantly longer RTs than normal ( $p = 2 \times 10^{-9}$ , binomial test).

Because the changes of mind were still low-frequency events, we sought additional confirmation that our analysis was not driven by random fluctuations but indeed identified a temporary preference for a choice that was different from the monkey's final choice. To this end, we performed two additional analyses. The first one was based on creating two artificial datasets for each real dataset, which were matched for number of trials, number of neurons, and firing rate profiles but did not contain any changes of mind, and repeating the HMM analysis. In the first dataset, each neuron emitted a Poisson spike train that followed the average firing rate profile of real neuron during the first 200 ms of each trial and then spent the remaining time of the trial ramping linearly to the average final firing rate of the neuron for the monkey's choice on this particular trial. In the second dataset, the first 200 ms were the same, but, instead of ramping to the final firing rate, all neurons would initially stay at the average firing rates 200 ms into the trial and, at some random time during the remaining time of the trial, jump to their final average firing rates for this particular choice and stay there until the end of the trial. For neither of these two different types of artificial datasets did the HMM analysis report more than two change-of-mind trials, indicating that false positives seem to be extremely rare.

The second analysis was based on the idea that, if the monkey had indeed developed a transient preference for a different choice, it might have, at least partly, been driven by fluctuations in the sensory stimulus. Although the stimuli were generated on the fly during the experiment, we had stored all necessary information to be able to recreate each of them exactly. We therefore recreated all stimuli that had been shown during change-of-mind trials and sent them through a set of motion energy filters, similar to those used by Kiani et al. (2008), to extract the time courses of motion energy in all three relevant directions. We then averaged the time courses of the motion energy component supporting the monkey's final choice, as well as the motion energy component supporting the monkey's transient preference, time-locked to the maximum posterior probability of the state associated with the temporary preference, as reported by the HMM analysis. The result of this analysis is shown in Figure 9. If the identified changes of mind were simply false positives or if they were only driven by factors other than the sensory stimulus, the average motion energy would not be expected to show any significant modulation around the time of the transient preference. However, as can be seen in Figure 9, the motion energy component supporting the monkey's final choice tended to drop just before the temporary preference for a different choice, whereas the motion energy component supporting the transiently preferred choice tended to increase just before such an event. This does not only indicate that the identified changes of mind were not just false positives but also that temporal fluctuations in the sensory stimulus made a significant contribution.

To summarize, the HMM analysis demonstrates that being in a firing rate state that resembles the typical activity pattern shortly before a particular choice does not necessarily mean that the trial



**Figure 9.** Fluctuations in motion energy and changes of mind. The plots show mean net motion energy in the random-dot stimulus (difference between motion energy in one particular direction and average of the motion energies in the other two directions) across all identified change-of-mind trials, time-locked to the probability peak of the transient state reflecting the temporary preference for a different choice option (vertical dashed line). The red trace shows the net motion energy driving the sensory evidence for the transiently preferred option, and the blue trace shows the net motion energy driving the sensory evidence for the option that was chosen last by the monkey. The confidence bands reflect  $\pm 1$  SE. The temporary preference for a different choice option during identified change-of-mind trials tended to be preceded by a drop in the net motion energy supporting the final choice and an increase in the net motion energy supporting the temporarily preferred option.

will be terminated by exactly this choice. We have seen quite a number of examples of changes of mind, meaning that the system transitioned into a state associated with a different choice. At times, these transitions were surprisingly fast (down to 30 ms). Thus, the neural system that is reflected by the activity of decision-related neurons in parietal cortex seems capable of fast state transitions, but the majority of trials, as supported by the other types of analyses, are characterized by gradual changes in neural activity.

## Discussion

In contrast to previous SU recordings of decision-related activity in parietal cortex during perceptual decision tasks, which did not provide sufficient data for assessing the neural dynamics at the single-trial level, we have made use of simultaneous recordings from up to four decision-related neurons to analyze how the neural activity changes over time during individual decisions. The results consistently indicated that, overall, the LIP population response during individual decisions is better explained by a gradual change of neural activity, as expected from an integrator-based mechanism, than by a sudden change of activity, as expected from a mechanism based on stochastic transitions between discrete states. The HMM analysis further allowed us to identify and analyze changes of mind during individual decisions at the neural level.

## SI analysis

In contrast to Horwitz and Newsome (2001), who reported that the majority of SUs recorded from the SC during a similar task had negative SIs, only a minority of our single-trial estimates of the LIP population activity (based on two to three neurons) had significantly negative SIs. Whereas Horwitz and Newsome interpreted their result as potentially indicating that SC neurons make



transitions between discrete firing rate states, our results indicate that the average single-trial LIP population response is more consistent with a gradual firing rate change. The discrepancy between the findings could be attributable to either having recorded from different areas (SC vs parietal cortex) or our results being based on at least two simultaneously recorded neurons, whereas the SC result was based on SUs. The results at the population level and at the SU level do not necessarily have to agree. For example, a single-trial ramping response at the population level could be implemented by different neurons in the population stepping at different times. When repeating our analysis at the SU level, again only a minority of the neurons had significantly negative SIs. Thus, the difference in results is most likely attributable to SC and LIP having different neural dynamics during individual decisions.

A number of the analyzed LIP sessions had significantly positive mean SIs when, like Horwitz and Newsome, using an analysis bin width of 25 ms. One possible scenario that could lead to a positive SI, as mentioned by Horwitz and Newsome, would be an extremely long refractory period. However, an analysis of the interspike interval distributions argues against this being the origin of our positive SIs. Alternatively, we found that robust periodic firing rate modulations with a frequency between 15 and 20 Hz could also reliably induce positive SIs. Because this effect critically depends on the choice of the bin width used for the SI analysis, we tested whether the SIs would change if the data were analyzed using a bin width of 40 ms instead of 25 ms. Almost all significantly positive SIs disappeared at both the population and the single-neuron level. Thus, it is quite likely that the observed positive SIs originated from periodic firing rate modulations. This is consistent with our finding of a task-relevant modulation in the 10–20 Hz band of the local field potential (Bollimunta and Ditterich, 2012).

### HMM analysis and changes of mind

The initial assumption behind the HMM analysis is that the activity pattern of a population of neurons can be described as sequences of discrete states, which are all characterized by a particular set of firing rates across the analyzed neurons. The technique had previously been used to reveal sequences of task-relevant states with sharp (on the order of 50 ms) transitions (Seidemann et al., 1996; Jones et al., 2007). Although our HMM analysis resulted in states that were associated with particular choices, the state transitions were mostly slow. Their average duration was 342 ms, and again, on average, only 25% of the trial duration was spent in one of the identified states. This seems to argue against a neural system that makes fast, stochastic transitions between discrete states. However, an analysis of the state transitions that were observed for our artificial datasets with either ramping or jumping responses yielded very similar results. These parameters are therefore apparently not particularly sensitive to how gradual or how sudden the firing rates change in our case. One reason is probably that, even in the case of artificial jumps in the later part of the trial, the firing rates were always undergoing gradual changes during the first 200 ms. Another reason is probably that the firing rate changes were sometimes too small to allow a reliable identification of fast transitions by the HMM analysis. The observation of longer transition durations than what has been reported previously in the literature should therefore not automatically be interpreted as evidence against fast changes in firing rate. However, the SI analysis and the comparisons between ramping and stepping non-homogeneous Poisson models, which were more sensitive in this

respect, suggest that firing rate changes in the LIP population during individual decisions are mostly gradual.

Because the HMM analysis estimates the most likely state sequence for each individual decision, one can identify trials whose state sequence contains states associated with two different choice alternatives. These are decisions during which the pattern of neural activity was already very similar to one that is typically obtained just before making a particular choice, but then the pattern changed and the trial ended with a different choice. We have referred to these trials as changes of mind. These neural observations share similarities with the behavioral observations made by Resulaj et al. (2009) in a version of the random-dot motion discrimination task that required human subjects to move a manipulandum from a starting position to one of two possible target positions. On a fraction of trials, these authors observed movements that started toward one of the targets, but ended at the other. Resulaj et al. referred to these trials also as changes of mind, which were more likely to occur during difficult discriminations (low coherence) than during easy discriminations (high coherence). Furthermore, they were more likely to change an incorrect choice into a correct one than vice versa. Our neural changes of mind had similar properties: they occurred mostly during difficult discriminations, and the majority of switches were from an incorrect choice to the correct choice. We further demonstrated that our neural changes of mind tended to be associated with longer-than-average RTs and that it was possible to identify features in the motion energy time course of the sensory stimuli that contributed to the monkeys' transient choice preference. Our results therefore indicate that simultaneous recordings from multiple LIP neurons can be used to decode the time course of individual perceptual decisions.

### Neural dynamics of perceptual decisions

As has been pointed out in the literature (Gold and Shadlen, 2001; Bogacz et al., 2006; McMillen and Holmes, 2006; Ditterich, 2010), nearly optimal perceptual decisions based on continuously inflowing sensory information should be based on (Multi-hypothesis) Sequential Probability Ratio Test-like algorithms, which require an operation close to perfect mathematical integration. Neural integrators can be implemented with attractor networks (Seung, 1996), with perfect integration being characterized by a flat energy landscape. If the integrator is leaky, the energy landscape has a slope that tends to move the system toward background activity levels. Other decision-making networks are based on the choice alternatives being represented by point attractors, which means that the energy landscape is sloped such that the system tends to move toward one of these attractors once it enters its attractive basin (Wong and Wang, 2006; Deco et al., 2009). The starting point can be either stable or unstable. In the first case, the system needs to be pushed out of the attractive basin of the initial state by strong enough sensory evidence, and, in the second case, it is free to move toward any of the choice attractors. If the attractor basins are deep enough and the slopes therefore steep, the system has the tendency to make stochastic transitions between states associated with the different point attractors and therefore exhibits jump-like behavior.

Our analysis of the single-trial activity in parietal cortex of monkeys performing perceptual decisions indicates that jump-like transitions are not the typical operating mode of the decision circuitry. The analyses that have been described in this paper point toward a gradual change of neural activity during most individual decisions, which suggests that the energy landscape, at least around the starting point, is relatively flat. Furthermore,

change-of-mind trials indicate that being close to one of the choice “attractors” does not mean that the outcome of the decision is irreversible. The basins of attraction around the choice alternatives therefore cannot be so steep that an escape is impossible. Interestingly, we have also seen examples of fast state transitions (from a state associated with one of the choice alternatives to a state associated with another alternative in  $<50$  ms), indicating that the system also seems capable of jump-like behavior. Thus, it will be interesting to see in future theoretical work how a decision network can exhibit both gradual changes in activity and also less frequent jump-like state transitions.

## References

- Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, Tishby N, Vaadia E (1995) Cortical activity flips among quasi-stationary states. *Proc Natl Acad Sci U S A* 92:8616–8620.
- Baras JS, Finesso L (1992) Consistent estimation of the order of hidden Markov chains. *Lect Notes Contr Inf* 184:26–39.
- Bennur S, Gold JI (2011) Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. *J Neurosci* 31:913–921.
- 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.
- Bollimunta A, Ditterich J (2012) Local computation of decision-relevant net sensory evidence in parietal cortex. *Cereb Cortex* 22:903–917.
- Bollimunta A, Knuth KH, Ding M (2007) Trial-by-trial estimation of amplitude and latency variability in neuronal spike trains. *J Neurosci Methods* 160:163–170.
- Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11:693–702.
- Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN (2011) Variance as a signature of neural computations during decision making. *Neuron* 69:818–831.
- Deco G, Rolls ET, Romo R (2009) Stochastic dynamics as a principle of brain function. *Prog Neurobiol* 88:1–16.
- Ditterich J (2006) Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw* 19:981–1012.
- Ditterich J (2010) A comparison between mechanisms of multi-alternative perceptual decision making: ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Front Neurosci* 4:184.
- Durstewitz D, Deco G (2008) Computational significance of transient dynamics in cortical networks. *Eur J Neurosci* 27:217–227.
- Gat I, Tishby N, Abeles M (1997) Hidden Markov modelling of simultaneously recorded cells in the associative cortex of behaving monkeys. *Network* 8:297–322.
- Gold JL, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5:10–16.
- Gregory PC (2005) Bayesian logical data analysis for the physical sciences: a comparative approach with Mathematica support. Cambridge, UK: Cambridge UP.
- Hannan EJ, Quinn BG (1979) Determination of the order of an autoregression. *J R Stat Soc Series B Stat Methodol* 41:190–195.
- Horowitz GD, Newsome WT (2001) Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. *J Neurophysiol* 86:2543–2558.
- 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.
- Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB (2007) Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A* 104:18772–18777.
- 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.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat Neurosci* 11:505–513.
- Maimon G, Assad JA (2009) Beyond Poisson: increased spike-time regularity across primate parietal cortex. *Neuron* 62:426–440.
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cereb Cortex* 13:1257–1269.
- McCane B, Caelli T (2004) Diagnostic tools for evaluating and updating hidden Markov models. *Pattern Recogn* 37:1325–1337.
- McMillen T, Holmes P (2006) The dynamics of choice among multiple alternatives. *J Math Psychol* 50:30–57.
- Miller P, Katz DB (2010) Stochastic transitions between neural states in taste processing and decision-making. *J Neurosci* 30:2559–2570.
- Radons G, Becker JD, Dülfer B, Krüger J (1994) Analysis, classification, and coding of multielectrode spike trains with hidden Markov models. *Biol Cybern* 71:359–373.
- Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput* 20:873–922.
- Resulaj A, Kiani R, Wolpert DM, Shadlen MN (2009) Changes of mind in decision-making. *Nature* 461:263–266.
- 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.
- Rorie AE, Gao J, McClelland JL, Newsome WT (2010) Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS One* 5:e9308.
- Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E (1996) Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci* 16:752–768.
- Seung HS (1996) How the brain keeps the eyes still. *Proc Natl Acad Sci U S A* 93:13339–13344.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- 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.
- Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350.
- Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev* 108:550–592.
- Wong KF, Wang XJ (2006) A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 26:1314–1328.