Discrete stepping and nonlinear ramping dynamics underlie spiking responses of LIP neurons during decision-making

David M. Zoltowski[1,⋆], Kenneth W. Latimer[2], Jacob L. Yates[3], Alexander C. Huk[4,5,6], and Jonathan W. Pillow[1,7]

[1]Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08540, USA
[2]Department of Neurobiology, University of Chicago, Chicago, IL, USA
[3]Center for Visual Science, University of Rochester, Rochester, NY, USA
[4]Department of Neuroscience, The University of Texas at Austin, Austin, TX 78712, USA
[5]Department of Psychology, The University of Texas at Austin, Austin, TX 78712, USA
[6]Center for Perceptual Systems, The University of Texas at Austin, Austin, TX 78712, USA
[7]Department of Psychology, Princeton University, Princeton, NJ 08540, USA
⋆Correspondence: zoltowski@princeton.edu

Abstract

Neurons in macaque area LIP exhibit gradual ramping in their trial-averaged spike responses during sensory decision-making. However, recent work has sparked debate over whether single-trial LIP spike trains are better described by discrete “stepping" or continuous drift-diffusion (“ramping”) dynamics. Here we address this controversy using powerful model-based analyses of LIP spike responses. We extended latent dynamical models of LIP spike trains to incorporate non-Poisson spiking, baseline firing rates, and various nonlinear relationships between the latent variable and firing rate. Moreover, we used advanced model-comparison methods, including cross-validation and a fully Bayesian information criterion, to evaluate and compare models. These analyses revealed that when non-Poisson spiking was incorporated into existing stepping and ramping models, a majority of neurons remained better described by stepping dynamics, even when conditioning on evidence level or choice. However, an extended ramping model with a non-zero baseline and a compressive output nonlinearity accounted for roughly as many neurons as the stepping model. The latent dynamics inferred under this model exhibited high diffusion variance for many neurons, making them qualitatively different than slowly-evolving continuous dynamics. These findings generalized to alternative tasks, suggesting that a robust fraction of LIP neurons are better described by each model class.

1 Introduction

Perceptual decision-making provides an opportunity to probe the role of different brain regions in cognitive tasks (Gold and Shadlen, 2007; Hanks and Summerfield, 2017). In direction discrimination tasks with choices conveyed by a saccadic eye movement (Newsome and Pare, 1988; Britten et al., 1992, 1996), macaque lateral intraparietal area (LIP) responses exhibit positive correlation with choice (Shadlen and Newsome, 1996, 2001). An important series of papers provided support for the idea that the firing rates of LIP neurons reflect the accumulation of sensory evidence in favor of a "preferred" choice target; this hypothesis unified neural responses and behavior under a single theoretical framework known as the drift-diffusion or accumulation-to-bound model (Roitman and Shadlen, 2002; Mazurek et al., 2003; Gold and Shadlen, 2007; Shadlen and Kiani, 2013). An extensive literature has examined this hypothesis in a variety of experimental and theoretical settings (Huk and Shadlen, 2005; Palmer et al., 2005; Ditterich, 2006a,b; Hanks et al., 2006; Kiani et al., 2008; Churchland et al., 2008; Kiani and Shadlen, 2009; de Lafuente et al., 2015)

Although the trial-averaged responses in LIP typically resemble ramps, the average responses do not directly reveal a neuron's single-trial dynamics. This shortcoming has motivated recent work to determine the single-trial dynamics of LIP responses in direction discrimination tasks (Churchland et al., 2011; Bollimunta et al., 2012; Latimer et al., 2015). In particular, Latimer et al. (2015) compared a discrete switching process or "stepping" model and a accumulation-to-bound or "ramping" model of LIP dynamics, both of which can give rise to ramping trial-averaged activity. They found that the majority of LIP cells were better explained by the stepping model. However, subsequent literature has sparked debate over the interpretation of these results (Shadlen et al., 2016; Zylberberg and Shadlen, 2016; Chandrasekaran et al., 2016; Latimer et al., 2017; Zhao and Kording, 2018).

In this paper, we aim to settle the debate about single-trial LIP dynamics using improved models and model comparison methods. We have extended the classic ramping and stepping models of LIP dynamics in several important ways. First, we incorporated spike-history dependencies into both models to account for the temporal autocorrelations of spike trains in LIP. Second, we investigated nonlinear ramping models with a non-zero baseline firing rate and several possible nonlinear relationships between the latent variable and firing rate. We compared these models using a principled, fully Bayesian information criterion and Bayesian leave-one-out cross-validation.

In our analyses, the stepping model outperformed the ramping model for a majority of neurons when both models were extended to incorporate spike-history. This result was robust to partitioning of the data by choice or sensory evidence level, showing that (in contrast to recent analyses in Zylberberg and Shadlen (2016)) model selection was not driven by anti-preferred choice or evidence trials. On the other hand, an extended ramping model with a non-zero baseline firing rate and a decelerating nonlinearity outperformed the stepping model for slightly more than half the neurons in our population. The latent firing rates inferred under this model, however, were often governed by high diffusion variability, making the distinction between continuous and discrete models less concrete. These analyses revealed that spike responses in LIP are more complex than simple ramping or stepping models, while confirming that discrete dynamics provide the best account for a substantial fraction of neurons in LIP.

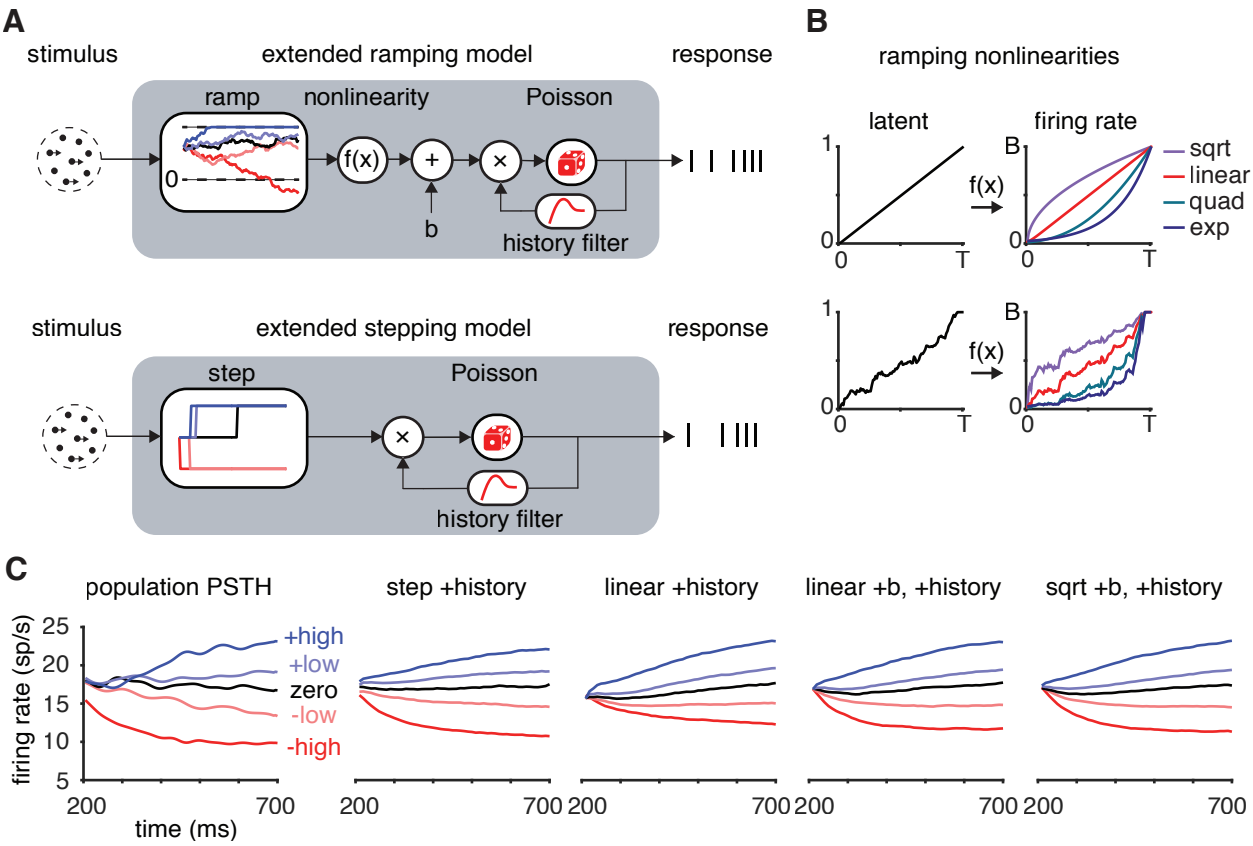

Figure 1. Latent variable models for LIP spike responses during decision-making. **A.** Schematics of extended ramping (above) and stepping (below) models. For the ramping model, the stimulus coherence sets the drift rate of a latent diffusion-to-bound process; this process is transformed by a rectifying nonlinearity $f(\cdot)$ and added to a baseline firing rate b. For the stepping model, stimulus coherence determines the distribution over time and direction of a discrete step to one of two possible latent firing rates. In both models, the latent firing rate is multiplied by the exponentiated output of a spike history filter, allowing it to capture non-Poisson firing statistics. **B.** Output nonlinearities considered for the extended ramping model. The nonlinearity transforms example latent diffusion paths (left) into different latent firing rate trajectories (right). **C.** Average firing rates for different motion coherence levels, aligned to motion onset and averaged across neurons (left), along with averaged responses simulated from models considered in this paper, where "+b" indicates inclusion of a non-zero baseline and "linear" and "sqrt" indicate the choice of nonlinearity in the extended ramping model.

2 Results

We formulated explicit statistical models of latent dynamics underlying single-trial spike trains in area LIP during perceptual decision-making and used two different statistical methods to compare them. Our analysis builds on Latimer et al. (2015), which formulated the ramping and stepping latent variable models of LIP spike trains. The basic ramping model, often referred to as the drift-diffusion or accumulation-to-bound model, consists of a continuous latent diffusion process, which is passed through a soft-rectifying nonlinearity to obtain a Poisson firing rate. The basic stepping model, on the other hand, consists of a discrete switching process that jumps from an initial firing rate to one of two levels with a probability that depends on the stimulus. Here we extended these two models in order to incorporate non-Poisson spike-history effects and to allow additional forms of nonlinearity in the ramp-

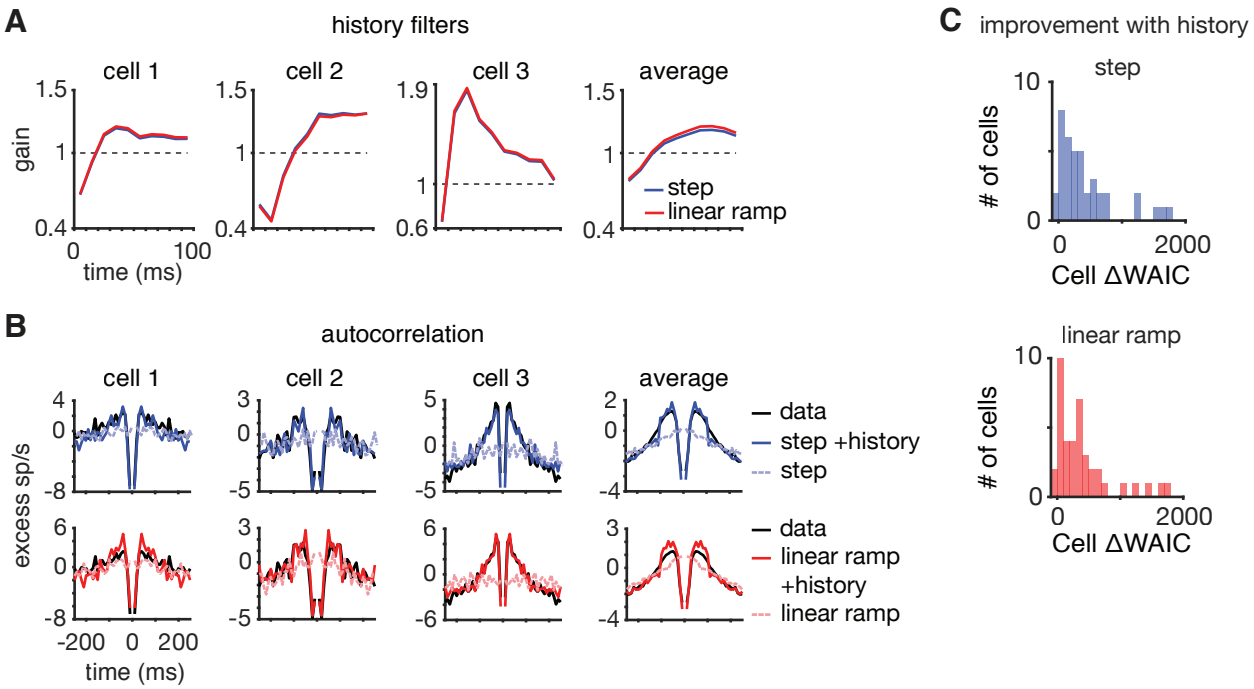

Figure 2. Models with spike-history dependence outperformed classic stepping and ramping models. **A.** Spike-history filters for three example neurons, and the average spike-history filter across neurons, for extended stepping (blue) and extended ramping (red) models. Inferred filters were remarkably similar between models, implying that spike-history effects did not vary with the choice of latent dynamics model. **B.** Spike train autocorrelations of LIP neurons and models with and without spike-history filters, revealing that classic stepping and ramping models could not account for the temporal statistics of real spike trains. **C.** Models with spike-history filters performed better than classic stepping (above) and ramping models (below) for the majority of LIP neurons, as quantified by WAIC. Positive WAIC differences favor the model with spike-history.

ing model (Figure 1).

To compare these models, we used two different methods: the Watanabe-Akaike information criterion (WAIC, Section 4; Watanabe, 2010; Gelman et al., 2014) and Bayesian leave-one-out cross-validation (Vehtari et al., 2017). The WAIC has multiple features that make it robust for model comparison. First, it uses the full posterior over the parameters for model evaluation, and therefore does not rely on a point estimate of the parameters (which is the case for other model-selection criteria, e.g., AIC, BIC, or DIC). Also, the penalty term in the WAIC is stable and guaranteed to be non-negative, in contrast with the DIC (Gelman et al., 2014; Vehtari et al., 2017). Finally, the WAIC has solid theoretical grounding as it is asymptotically equivalent to Bayesian leave-one-out cross-validation (Gelman et al., 2014; Vehtari et al., 2017). We find these benefits are realized empirically, as the WAIC outperforms the DIC at identifying the true model in simulations (Figure 7B).

2.1 Incorporating spike-history dependencies

The basic ramping and stepping models from Latimer et al. (2015) described spiking as Poisson conditioned on the latent ramping or stepping process, which ignores spike-history effects present in real spike trains (e.g. refractoriness, bursting, spike-rate adaptation). We therefore extended both mod-

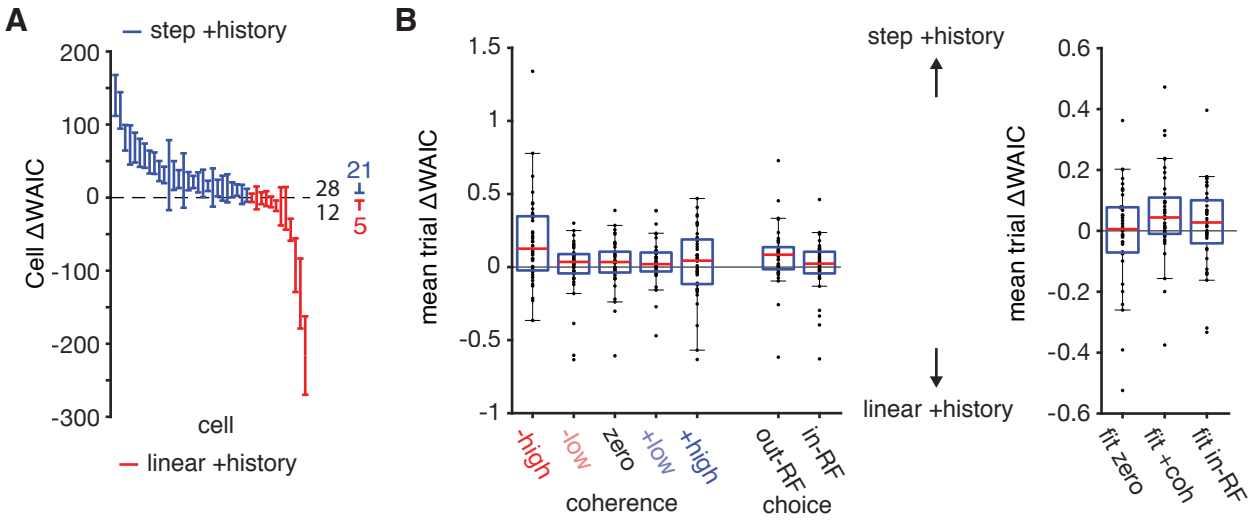

Figure 3. Stepping model with spike-history outperforms linear ramping model with spike-history for a majority of cells across conditions. **A**. Sorted WAIC differences between the stepping and linear ramping models with spike-history for all neurons (error bars indicate ±1SEM). Blue (red) lines indicate cells for which the stepping (ramping) model had a better WAIC value. Colored numbers indicate the number of cells where the WAIC difference favored stepping or ramping by more than the standard error. **B**. *Left*: WAIC differences computed for subsets of trials conditioned on motion coherence and choice. The differences were normalized by the number of trials in each condition. Positive differences favor the stepping model. *Right*: WAIC differences for models fit only to data from zero coherence, positive coherence, or in-RF trials (respectively), normalized by the number of trials in each condition. Similar results were obtained with cross-validation (Figure 8).

els to include autoregressive spike-history filters, like those in the generalized linear modeling (GLM) framework (Figure 1, Section 4). These filters capture non-Poisson spike-history dependencies (Truccolo et al., 2005; Weber and Pillow, 2017) and allow for a dissociation of latent dynamics from spiking activity that can be explained more parsimoniously by past spikes.

We fit stepping and ramping models both with and without spike-history to the responses of 40 LIP cells during a variable-duration random dot motion task (Section 4; Meister et al., 2013). Simulated data from the fitted models captured the shape of true population PSTHs for different coherence levels (Figure 1C). The inferred spike history filters typically exhibited short-timescale inhibition (or refractoriness) and longer-timescale self-excitation, although there was substantial variability across neurons (Figure 2A). These filters conferred a dramatic improvement in the ability to capture temporal autocorrelations in spiking activity under both models (Figure 2B). Intriguingly, the history filters were nearly identical between stepping and ramping models for the majority of neurons. This suggests that the filters captured similar structure across models, and were not strongly influenced by model-specific assumptions about latent dynamics. Spike-history filters substantially improved prediction accuracy under both models for the vast majority of neurons, giving better WAIC for 38 out of 40 cells (Figure 2C).

2.2 Stepping model robustly outperforms linear ramping model

Model comparison using the WAIC and cross-validation revealed that the stepping model with spike-history outperformed the linear ramping model with spike-history for 28 out of 40 cells (Figures 3A and

8A). We quantified uncertainty in the model comparison using the standard error of the WAIC difference across trials (Section 4). For the 26 cells whose WAIC differences were more than the standard error from zero, 21 favored the stepping model. This result was not driven by the penalty term in the WAIC, as the stepping model had higher predictive accuracy for 33 of the 40 cells (higher lppd, Section 4, Equation 68).

To examine how different trial conditions contributed to the model comparison across cells, we then compared the models across subsets of trials partitioned by the motion coherence and choice. For each coherence and choice, the stepping model with spike-history outperformed the linear ramping model with spike-history for the majority of cells (Figures 3B and 8B). The negative high coherence trials had the largest median difference between cells, which reflects concerns about the inability of the ramping model to handle strong negative drifts in firing rate (Shadlen et al., 2016; Zylberberg and Shadlen, 2016). However, the median difference favored stepping for each coherence level, including those for weak or strong motion into the RF. Also, with the negative high coherence or all negative coherence trials excluded, the stepping model still outperformed the ramping model for a majority of cells (26 out of 40). This suggests that the overall comparison did not depend on the negative coherence trials and that the dynamics were consistent across stimulus conditions and choice.

We next evaluated the possibility that inference of the model parameters was influenced by subsets of trials, which would also affect the model comparison. We re-fit the models to three different subsets of trials (Figures 3C and 8B). The first contained all zero-coherence trials, which putatively have long integration times (and therefore might be expected to have the slowest or most gradual "ramp-like" dynamics). The other two re-fits used data from all positive coherence trials and data from all in-RF choice trials, in which the animal made a saccade to the "preferred" target. These latter two fits restricted analysis to trials with putatively positive values of accumulated sensory evidence. For all three analyses, model comparison favored the stepping model with spike-history for more than half the neurons: 21/40 for zero-coherence trials, 25/40 for positive coherence trials, and 27/40 for in-RF choice trials.

Recent work has argued that trials with strong initial negative diffusion and termination at non-zero rates might bias model comparison in favor of the stepping model (Zylberberg and Shadlen, 2016). The competing accumulators model assumes that these trials occur most prevalently during negative coherence trials or out-RF choice trials (Mazurek et al., 2003). These are trials where either the presented or perceived evidence is to the target outside the RF of the LIP neuron, and therefore are trials where the LIP neurons are more likely to decrease their firing rates throughout a trial. Given that assumption and the analyses performed above, we conclude that the comparison between the ramping and stepping models with history fit to all trials was not driven by these trials, even though a mechanism such as a baseline firing rate to stop strongly going negative rates was not included.

2.3 Nonlinearities and non-zero baselines improve the ramping model

Although the ramping model formulated in (Latimer et al., 2015, 2017) was motivated to capture the hypothesized linear relationship between a biased diffusion-to-bound process and single-neuron firing rates, neurons might exhibit nonlinear relationships between a putative latent diffusion process and their firing rates. To investigate this possibility, we fit nonlinear ramping models with a variety of different

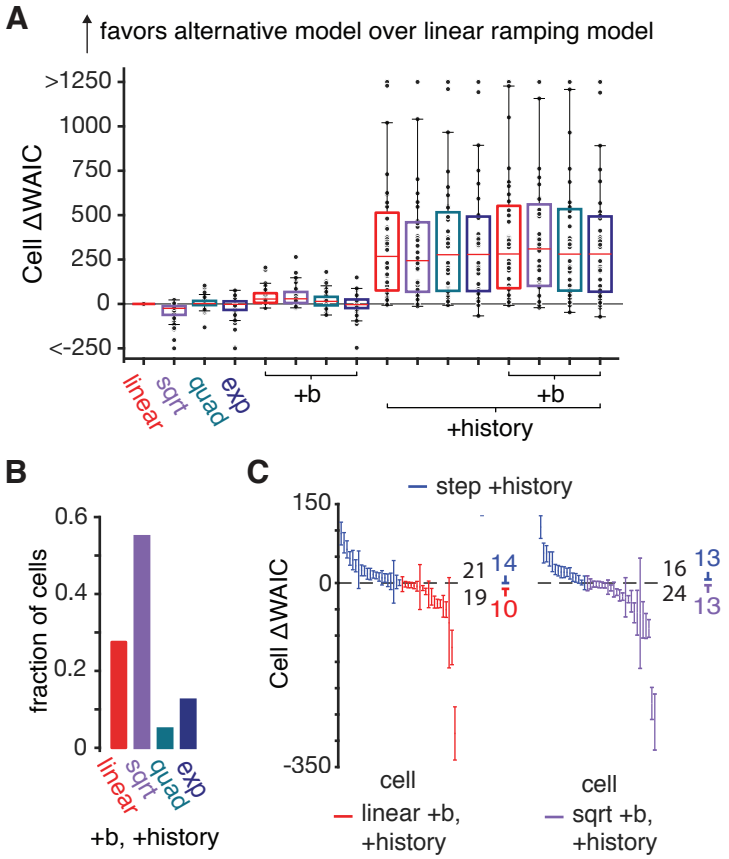

Figure 4. Comparison of extended ramping models. **A**. Comparison of nonlinear ramping models with non-zero baseline firing rates and spike-history against the classic (linear) ramping model. Positive values indicate improvements relative to the classic linear ramping model in terms of WAIC difference. The box color indicates the choice of nonlinearity and "+b" refers to models with a non-zero baseline. **B**. The fraction of cells for which each model (all with non-zero baseline and spike-history) achieved the best WAIC, showing that the square root nonlinearity performed best for more than half the population. **C**. WAIC differences between the stepping model with spike-history and the linear (*left*) and square root (*right*) ramping models with baseline and spike-history. See Fig. 8 for comparable analysis using cross-validation instead of WAIC.

nonlinearities: a soft-rectified square root function ("sqrt"), a soft-rectified quadratic ("quad"), or an exponential ("exp") (Section 4). These models can capture varying degrees of nonlinear response saturation or acceleration as a function of the latent variable (Figure 1B). We also included a non-zero baseline firing rate parameter b that acts as a (non-absorbing) lower bound on the firing rate (Section 4). This prevents firing rate from going to zero when the drift term is strongly negative (Shadlen et al., 2016; Zylberberg and Shadlen, 2016; Latimer et al., 2017).

For each nonlinearity, we fit the ramping model with and without an additive baseline and spike-history filters. We compared each model to the original linear ramping model without spike-history from Latimer et al. (2015) (Figure 4A). Including the non-zero baseline improved all models, and models with the largest median improvements over the linear ramping model included both the non-zero baseline and spike-history filters. Across all nonlinear ramping models with baseline, the model with square root nonlinearity performed best, achieving the best WAIC for more than 50% of neurons (Figure 4B).

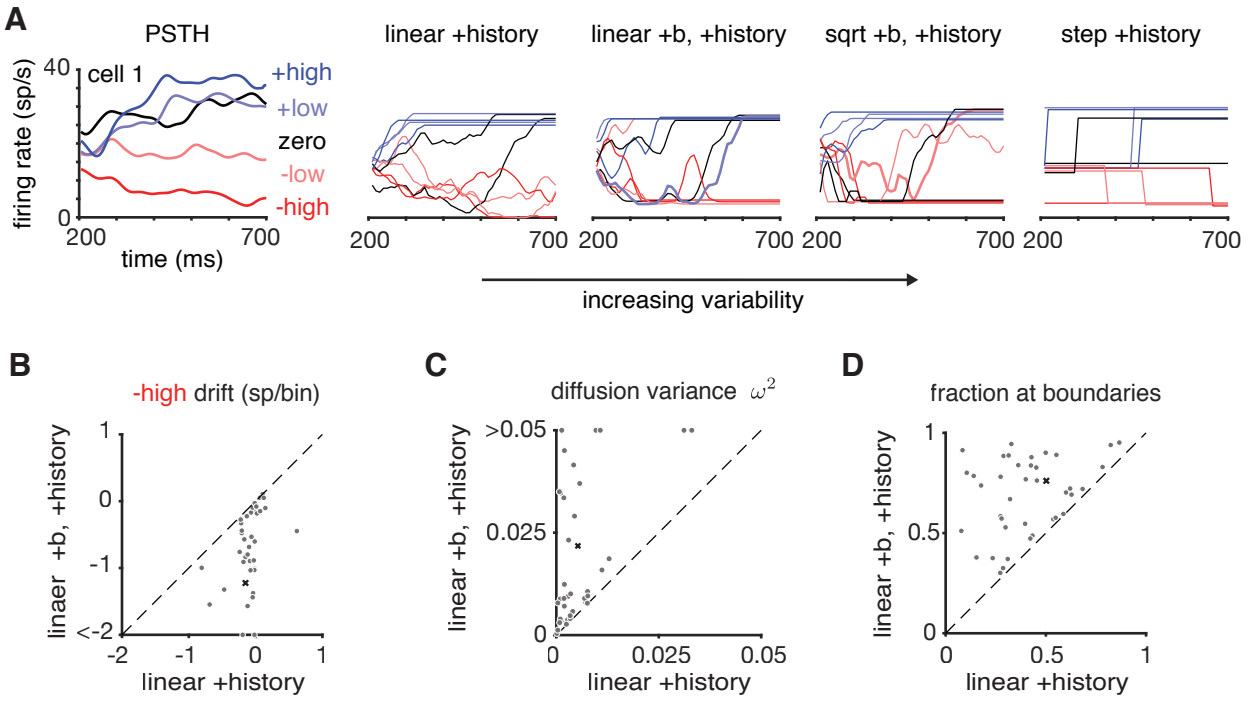

Figure 5. Analysis of latent firing rates under different latent dynamical models. **A**. Coherence-conditional PSTHs for example cell #1 (left) and simulated latent firing rate paths from fitted nonlinear ramping models and stepping model (right). The data were simulated using samples from the posterior. The linear and square root ramping models with non-zero baseline and spike-history can produce latent trajectories with high variability. The bolded trajectories are simulated trajectories that quickly evolved from the baseline to the upper boundary. **B-D**. Scatter plots comparing latent firing rate dynamics under the ramping model with and without a non-zero baseline. Each point corresponds to a pair of model fits for a single neuron. **B**. Drift rates for the highest negative-coherence stimulus are more negative under a ramping model with non-zero baseline than without. This indicates that adding a non-zero baseline allows the firing rate to ramp downwards more rapidly for negative-coherence motion (Zylberberg and Shadlen, 2016). **C**. Diffusion variance increased with the addition of the non-zero baseline. **D.** Fraction of time the simulated latent firing rates were equal to the baseline rate or upper absorbing boundary. See Figure 9 for comparable analyses with the square root ramping model with non-zero baseline.

As the linear and square root ramping models with non-zero baseline and spike-history were the best-performing among all models with continuous latents, we performed a direct comparison with the stepping model with spike-history. These extended ramping models closely matched the performance of the stepping model with spike-history (Figures 4C and 8C). The square root ramping model with non-zero baseline and spike-history achieved better WAIC than the stepping model with spike-history for more than half of the cells, although they each had an equal number of cells with WAIC differences more than the standard error from zero.

2.4 Inferred single-trial trajectories of nonlinear ramping models are more discrete

In the previous section, we showed that the linear and square root ramping models with non-zero baseline outperformed the linear ramping model for the majority of cells. We therefore investigated how these extensions affected different aspects of the model behavior. First, we observed that including a

non-zero baseline helped to capture the steep initial decrease in the average firing rate for negative-coherence trials via steeper negative drift rates, which was missing for the simpler ramping model (Figure 5A-B and Figure 1C).

Second, we observed that single-trial trajectories sampled from the models with non-zero baseline looked more discrete, exhibiting more rapid jumps to maximal or minimal firing rate, than trajectories from the linear ramping model without baseline (Figure 5A). Adding a non-zero baseline led to an increase in diffusion variability for most of the cells (Figure 5C), which allowed for larger changes in firing rate between time bins. Simulated trajectories also spent larger fractions of the time at the lower and upper boundaries compared to the models with a zero baseline (Figure 5D). Because the lower bound is non-absorbing, the firing rate can hit the lower bound and still evolve to the upper bound during the course of the trial; examples of these trials are shown in bold in Figure 5A. The effects of increased variability and increased time at minimal or maximal firing rates were also observed in the model with square-root nonlinearity and non-zero baseline when compared to the linear model (Figure 9).

Overall, these findings suggest that including a non-zero baseline rate can improve the ramping model in multiple ways. It does help the ramping model capture strong negative going rates (Zylberberg and Shadlen, 2016). However, with these modifications, the ramping model produced highly variable latent trajectories, which transitioned rapidly to minimal or maximal rates, making them qualitatively less similar to the gradually drifting rates expected from a perfect accumulator. Put simply, inclusion of a non-zero baseline moved both linear and nonlinear ramping models closer to discrete dynamics.

2.5 Generalization to two additional decision-making tasks

To see if the results for this dataset generalized to recordings from LIP in other direction discrimination tasks, we repeated a subset of our analyses on recordings during a discrete-pulse accumulation task (Yates et al., 2017) and a reaction-time (RT) task (Roitman and Shadlen, 2002; Section 4). The trial-averaged responses of neurons in both tasks resembled gradual ramps (Figure 6). However, statistical comparison of stepping and ramping models yielded results consistent with our findings above (Figures 6 and 10). For both additional datasets, spike-history filters improved the fits of both ramping and stepping models. The linear ramping model performed better with inclusion of a non-zero baseline, and the square-root ramping with non-zero baseline model outperformed the linear ramping with non-zero baseline model (both with history) for the majority of units in both datasets: 65/115 in the discrete-pulse task and 9/16 in the RT task. The stepping model with spike-history outperformed the linear ramping and square root ramping models with non-zero baseline and spike-history for a large majority of units in the discrete-pulse task. In the RT task, the models were evenly split, with models in the ramping class favored for cells with WAIC differences greater than the standard error from zero.

It is worth summarizing the comparison between the stepping model and the square root ramping model with non-zero baseline (both with spike-history), as neither model outperformed the other across all datasets. Between the two models, the stepping model with spike-history was favored for slightly less than half of cells in the variable duration task (16/40), a majority of cells in the discrete-pulse task (73/115), and exactly half of cells in the RT task (8/16). However, for cells with WAIC differences greater than the standard error from zero, the stepping model with spike-history was favored for exactly half of cells in the variable duration task (13/26), a majority in the discrete-pulse task (46/72), and a minority in

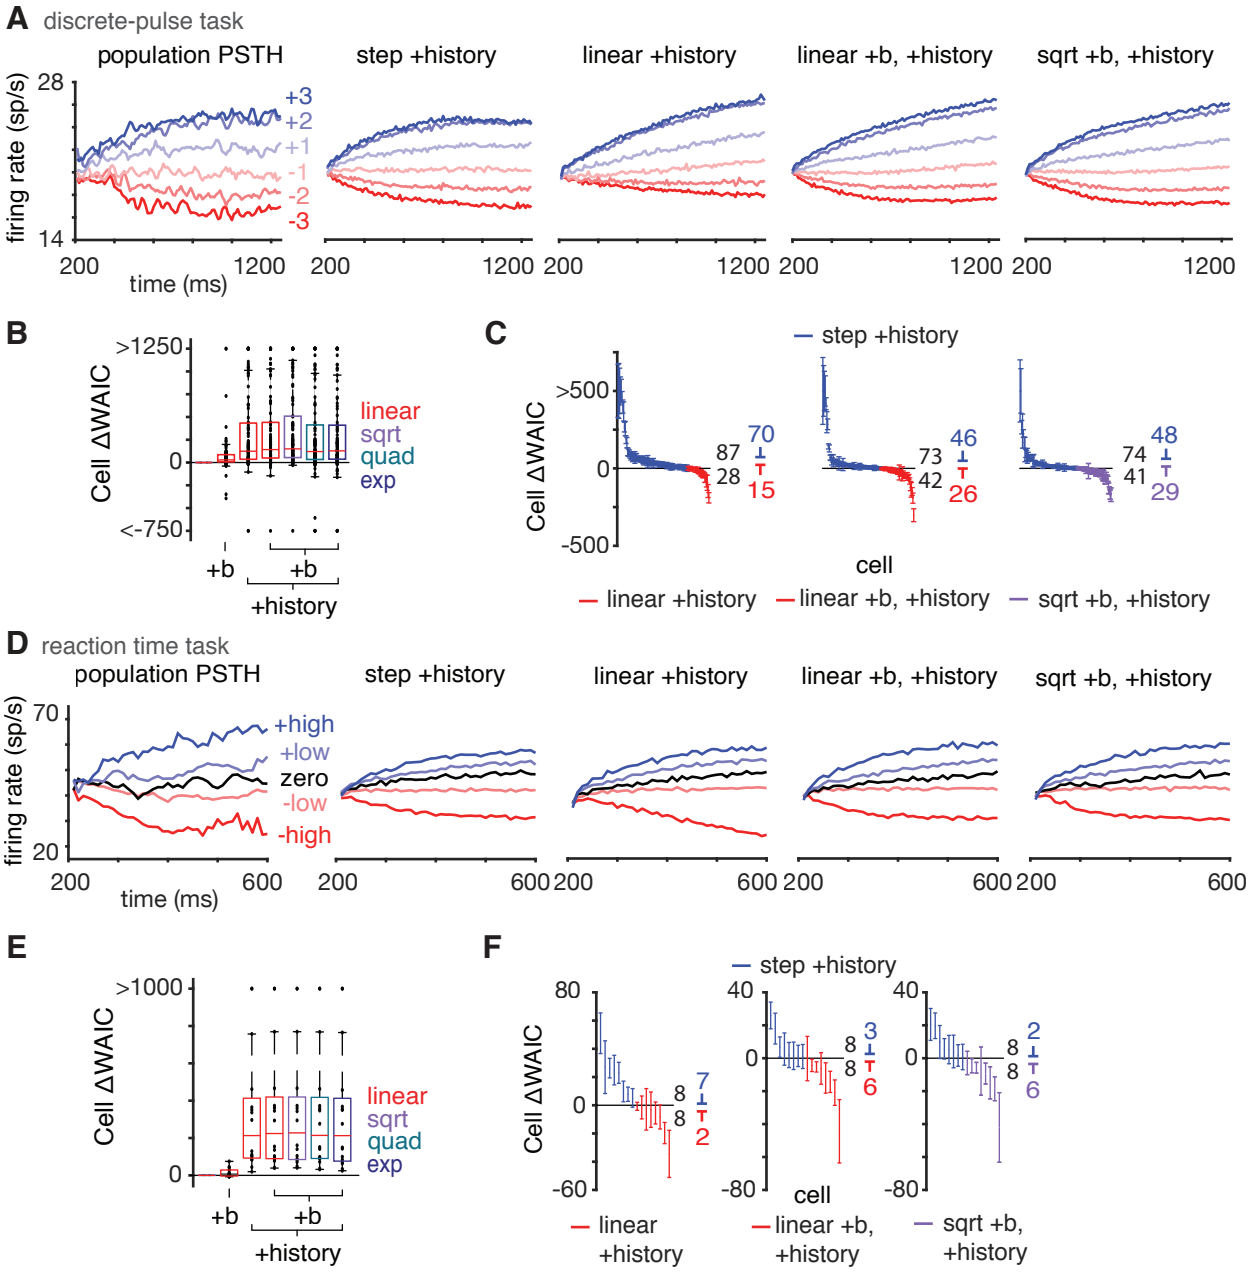

Figure 6. Analyses of neural data from two additional decision-making tasks. **A**. Population PSTHs from LIP neurons and simulated data from fitted models in a discrete-pulse task (Yates et al., 2017). **B.** Model comparison of extended nonlinear ramping models and classic linear ramping model using WAIC. **C**. The stepping model with spike-history achieved higher WAIC differences for the majority of cells. **D**. The population PSTH and simulated PSTHs from the fitted models, aligned to motion onset, in a reaction-time (RT) task (Roitman and Shadlen, 2002). **E.** Spike-history and non-zero baseline firing rates improved ramping model performance, but different nonlinearities performed equally well. **F.** Comparison of the stepping model and different ramping models, all with spike-history. See Figure 10 for comparable analyses using cross-validation.

the RT task (2/8). In this view, across datasets these two models achieved relatively equal performance.

3 Discussion

Our study strengthens the evidence for discrete-state models of the single-trial dynamics of many LIP cells during decision-making. Importantly, we found that LIP dynamics are heterogeneous, with discrete stepping and continuous diffusion-to-bound models both accounting for a substantial fraction of neurons (Meister et al., 2013; Park et al., 2014; Latimer et al., 2015). Our findings are supported by dynamical models that account for spiking autocorrelation and that allow for nonlinear mappings from diffusion-to-bound to firing rate in the ramping model. We obtained the same results when using a fully Bayesian information criterion and leave-one-out cross validation.

Although we have significantly extended the models considered in Latimer et al. (2015), which allowed for more accurate descriptions of spike trains in LIP, there are a variety of other proposed extensions that we have not yet explored. For example, the ramping model could use a random start time to the diffusion process on each trial (Churchland and Kiani, 2016). We speculate that the ramping model has more flexibility to handle small deviations in the onset of a diffusion-to-bound process because it has continuous latent states that place mass over all possible trajectories in the latent space, although we have not performed this comparison explicitly. Next, one could formulate a ramping model in which negative drifting rates are stopped by a competing accumulator on each trial, rather than a non-zero baseline that is constant across trials (Mazurek et al., 2003; Zylberberg and Shadlen, 2016; Latimer et al., 2017). However, this would require a two-dimensional latent diffusion process, which would be computationally more demanding than the one-dimensional models we have considered here, and might prove more difficult to identify with single-neuron data. Fitting such a model using multi-neuron recordings therefore presents one promising direction for future work. Finally, one could extend the stepping model to a general hidden Markov model, allowing for more than three discrete firing rates, with more flexible transition dynamics that allow for more than one transition per trial (Bollimunta et al., 2012). Such a model would have more flexibility than the stepping model we considered, which might allow better generalization to alternate tasks (Janssen and Shadlen, 2005; Yang and Shadlen, 2007; Kira et al., 2015; Morcos and Harvey, 2016).

Our findings appear to contradict a recent study from Zhao and Kording (2018), which reported that the best model of LIP responses, according to a cross-validation analysis, was a model with a constant firing rate on every trial. Although the specific models differed from those we have considered here in multiple ways, we believe the discrepancy is likely due to the fact that Zhao and Kording (2018) treated latent firing rates as parameters to be estimated, instead of marginalized or integrated over. This resulted in models with one fitted parameter per trial (a step time, ramp slope, or constant firing rate), making for hundreds of parameters per neuron, which is far more than the models we have considered here. We suspect that this approach therefore suffered from overfitting, leading to the dubious conclusion that firing rates are constant over time (a result that is inconsistent with the basic ramping apparent in trial-averaged activity). We have shown here that WAIC and cross-validation give virtually identical results when integrating over the unobserved latent firing rates. Nonetheless, the suggestion from Zhao and Kording (2018) that baseline firing rates may vary stochastically across trials is interesting, and consistent with recent findings about spike count variability (Goris et al., 2014; Charles et al., 2018). Incorporating slow changes in gain or excitability over trials therefore represents an additional promising direction for future work.

Some authors have raised the concern that animals performing the task are adopting a strategy that involves integration over a shorter time period than the entire trial, which could produce discrete-looking neural dynamics even in neurons that are accumulating evidence (Shadlen et al., 2016). Further, alternative strategies without accumulation can also match some behavioral features of evidence accumulation (Ditterich, 2006b; Stine et al., 2018). We recognize the ambiguity in determining behavioral strategy, and in future work we expect that including the time-varying evidence stream in the models could help identify the behavioral strategies used by the animals.

The GLM framework we have used to incorporate spike-history effects could naturally be expanded to include regressors for experimental variables related to the stimulus or behavior of the animal. A worthwhile future direction would be to fit latent variable models with GLM regressors in tasks with structured stimuli, to better disentangle latent dynamics from sensorimotor variables that affect neural activity on single trials (Brunton et al., 2013; Hanks et al., 2015; Morcos and Harvey, 2016; Katz et al., 2016; Scott et al., 2017; Yates et al., 2017; Huk et al., 2017). As behavioral paradigms become richer, and as the numbers of recorded neurons increase, we expect that population latent variable models with regressors and GLM outputs will provide a powerful framework for studying the neural computations underlying sensory decision-making in a wide variety of tasks and brain areas.

4 Methods

4.1 Data

We analyzed the responses of LIP cells during three motion-discrimination tasks. Our primary analyses were performed on the responses of 40 LIP cells (single-units) recorded from two rhesus monkeys during a variable-duration random dot motion task, originally described in Meister et al. (2013). In the task, a random dot motion stimulus was presented for durations uniformly drawn in the range 500 to 1000 ms after a variable delay. The dot motion coherence, or the expected percentage of dots moving in the true direction at each time point, on each trial was taken from the set of values: 0.0, 3.2, 6.4, 12.8, 25.6, or 51.2%. The monkey reported its estimate of the dot motion direction via a saccade to one of two targets. One saccade target was placed in the response field of the neuron ("in-RF") while the other was placed outside of the response field ("out-RF"). The animal had to wait for 500 ms after the stimulus was extinguished before it could indicate its choice. The original study recorded from 80 LIP neurons and the 40 LIP cells used in this study were the 40 most choice-selective responses during the period 200-700 ms after motion onset, determined by the d' criterion (Latimer et al., 2015). For analysis, the trials were grouped into five coherence levels: zero included 0% trials, positive/negative low included 3.2%, 6.4%, and 12.8% trials, and positive/negative high included 25.6% and 51.2% trials, where positive motion is towards the target in the response field. The coherence-dependent PSTHs for each unit in the variable duration task were smoothed using a Gaussian filter with 20 ms standard deviation.

We analyzed two additional datasets. The first consisted of 16 LIP cells (single-units) recorded from two rhesus monkeys during a reaction-time (RT) version of the random dot motion task (Roitman and Shadlen, 2002), where the monkey chooses when to respond. The details for the selection of the 16 LIP cells are in (Latimer et al., 2015). We also divided the trials in this dataset into the five levels described above. We included spikes starting at 200 ms after stimulus onset and up to 50 ms before the saccade for analysis. The final spike bin contained the time point 50 ms before the saccade and we included all spikes that fell into this bin. We only included trials in which we had 100 ms of data in this period (Latimer et al., 2015).

The second additional dataset consisted of 115 LIP units (single- and multi-units) from two rhesus monkeys performing a discrete-pulse accumulation task (Yates et al., 2017). In this task, the animal viewed a set of Gabor patches that either flickered or drifted during seven discrete portions of the trial (pulses). In each pulse, all of the drifting Gabor patches moved in the same direction. The task of the animal was to report the net motion direction across the seven pulses with a saccade to one of two targets. In this task, the net motion levels did not map directly to discrete-coherence levels, as each trial could have a different amount of net pulses in either direction. Therefore, we partitioned the data for each experimental session into six levels by sectioning the net pulses in each direction into thirds. We selected neurons with d' statistic magnitudes larger than 0.2 for analysis (Yates et al., 2017).

4.2 Models

4.2.1 Ramping model with history

In the ramping model, the firing rate is linked to a latent diffusion process (Latimer et al., 2015). In the following, we describe the generative model. Individual trajectories in the latent space are initiated with a sample from a Gaussian with mean x_0 and variance ω^2. The trajectory then evolves according to drift-diffusion dynamics with a diffusion variance ω^2 and drift β_c. The drift β_c depends on the coherence c of the current trial. The latent trajectory is scaled by a factor of γ and passed through a nonlinearity $f(x)$ to map it to a positive firing rate space. In the linear ramping model, the nonlinearity is the softplus function $f(x) = \log(1 + \exp(x))$. The output of the nonlinearity is multiplied by history dependence g_t such that the firing rate in spikes per second is $\lambda_t = f(x_t\gamma)\, g_t$. If the trajectory crosses an absorbing upper boundary at 1 in the latent space then the firing rate is held fixed for the remainder of the trial at the boundary rate. The generative model for a trial of length T is

$$x_1 \sim \mathcal{N}(x_0, \omega^2) \tag{1}$$

$$x_{t+1} \sim \mathcal{N}(x_t + \beta_c, \omega^2), \quad t > 1 \tag{2}$$

$$y_t \sim \begin{cases} \text{Poisson}(f(x_t\,\gamma)\, g_t\, \Delta), & t < \tau \\ \text{Poisson}(f(\gamma)\, g_t\, \Delta), & t \geq \tau \end{cases} \tag{3}$$

where τ is the first time bin that $x_t \geq 1$ (otherwise $\tau = \infty$) and $\Delta = 0.01s$ is the bin size. The bin size is equal to one frame of the stimulus in Meister et al. (2013).

The history dependence modulates the firing rate through a multiplicative interaction. At time t, the history dependence g_t is the exponential of the weighted sum of the previous H bins of spiking activity of the neuron

$$g_t = \exp\left(\sum_{h=1}^{H} w_h y_{t-h}\right) \tag{4}$$

with $\mathbf{w} = [w_1, \cdots, w_H]^{\mathrm{T}}$ a vector of weights. In the models without history $g_t = 1$ for all t. We used $H = 10$ bins for 100 ms of history dependence.

The ramping model parameters are $\Theta = \{\beta_{1:C}, x_0, \omega^2, \gamma, \mathbf{w}\}$ where C is the number of coherence levels. The latent variables in the ramping model $\mathbf{x}$ are the latent diffusion trajectories for each trial.

4.2.2 Nonlinear ramping models

The nonlinear ramping models are linked to the latent diffusion process of the ramping model through alternative nonlinearities. We used three alternative nonlinearities: a soft square root $f(x) = \log(1 + \exp(x))^{1/2}$ (sqrt), a soft quadratic $f(x) = \log(1 + \exp(x))^2$ (quad), and an exponential $f(x) = \exp(x)$ (exp). The parameters and latent variables of the nonlinear ramping models are unchanged from the ramping model.

4.2.3 Non-zero baseline firing rate

In the linear and nonlinear ramping models with non-zero baseline rates, the output of the nonlinearity is shifted by a positive baseline firing rate parameter b before multiplication with the history term (Zylberberg and Shadlen, 2016; Latimer et al., 2017) such that the firing rate is $\lambda_t = (f(x_t\gamma) + b)\, g_t$. The generative model is

$$x_1 \sim \mathcal{N}(x_0, \omega^2) \tag{5}$$

$$x_{t+1} \sim \mathcal{N}(x_t + \beta_c, \omega^2), \quad t > 1 \tag{6}$$

$$y_t \sim \begin{cases} \text{Poisson}((f(x_t\,\gamma) + b)\, g_t\, \Delta), & t < \tau \\ \text{Poisson}((f(\gamma) + b)\, g_t\, \Delta), & t \geq \tau. \end{cases} \tag{7}$$

The parameters of the ramping and nonlinear ramping models with non-zero baseline are $\Theta = \{\beta_{1:C}, x_0, \omega^2, \gamma, b, \mathbf{w}\}$, where C is the number of coherence levels. The latent variables are $\mathbf{x}$.

4.2.4 Stepping model with history

In the stepping model, the initial firing rate starts at a state α_0. During the trial, the state can either remain constant or it can switch to one of two other states, a down state α_1 or an up state α_2 (Latimer et al., 2015). The step direction d is sampled from a Bernoulli distribution such that the probability of a step to α_2 is ϕ_c and the probability of a step to α_1 is $1 - \phi_c$. The step time z is drawn from a negative binomial (NB) distribution with a shape parameter r and coherence-dependent mean step time m_c. Both the step direction and step time vary from trial to trial. The stepping model firing rate is the product of the state and the spike-history dependent gain g_t. The generative model for a trial of coherence c is

$$z \sim \text{NB}(m_c, r) \tag{8}$$

$$d \sim \text{Bernoulli}(\phi_c) \tag{9}$$

$$y_t \sim \begin{cases} \text{Poisson}(\alpha_0\, g_t\, \Delta), & t \leq z \\ \text{Poisson}(\alpha_1\, g_t\, \Delta), & t > z, d = 1 \\ \text{Poisson}(\alpha_2\, g_t\, \Delta), & t > z, d = 2. \end{cases} \tag{10}$$

The stepping model parameters are $\Theta = \{\alpha_0, \alpha_1, \alpha_2, m_{1:C}, \phi_{1:C}, r, \mathbf{w}\}$. The bin size is $\Delta = 0.01s$. The history-dependence has the same parameterization as the ramping models. The latent variables in the stepping model are the step times and step directions on each trial. This stepping model is a reparameterization of the stepping model in (Latimer et al., 2015), which used scale parameters p_c instead of m_c, where $p_c = \frac{m_c}{m_c + r}$. We used the mean step time parameterization because the parameters r and m_c are less correlated than r and p_c, which improved mixing in the MCMC algorithm described below.

4.3 Model Inference

4.3.1 Prior Distributions

We used the following priors on the parameters of the ramping and nonlinear ramping models

$$p(x_0) = \mathcal{N}(x_0; \mu_x, \sigma_x^2) \tag{11}$$
$$p(\beta_c) = \mathcal{N}(\beta_c; \mu_\beta, \sigma_\beta^2) \tag{12}$$
$$p(\omega^2) = \text{Inv-Gamma}(\omega^2; \alpha_\omega, \beta_\omega) \tag{13}$$
$$p(\gamma) = \text{Gamma}(\gamma; \alpha_\gamma, \beta_\gamma) \tag{14}$$
$$p(b) = \text{Gamma}(b; \alpha_b, \beta_b) \tag{15}$$
$$p(w_i) = \mathcal{N}(w_i; \mu_h, \sigma_h^2). \tag{16}$$

For all models, the priors on the diffusion drifts and variance were $\mu_\beta = 0$, $\sigma_\beta = 0.1$, $\alpha_\omega = 1.1$, and $\beta_\omega = 1\text{e}{-3}$. In models with a zero baseline, we set $\mu_x = 0$ and $\sigma_x = 10$. In models with a non-zero baseline, we set $\mu_x = 0.5$ and $\sigma_x = 0.5$ and used $\alpha_b = 1$ and $\beta_b = 0.01$ for the prior on the baseline parameter. The prior on the bound height varied for each nonlinearity. We used $\alpha_\gamma = 2$ and $\beta_\gamma = 0.05$ for the softplus, $\alpha_\gamma = 1$ and $\beta_\gamma = 1\text{e}{-4}$ for the soft square root, $\alpha_\gamma = 3$ and $\beta_\gamma = 0.5$ for the soft quadratic, and $\alpha_\gamma = 3$ and $\beta_\gamma = 3$ for the exponential. The parameters for the prior on the history weights were $\mu_h = 0$ and $\sigma_h^2 = 10$.

The priors on the stepping model with history were

$$p(\alpha_0) = \text{Gamma}(\alpha_0; \alpha_\alpha, \beta_\alpha) \tag{17}$$
$$p(\alpha_1, \alpha_2) \propto 1(\alpha_2 > \alpha_1)\,\text{Gamma}(\alpha_1; \alpha_\alpha, \beta_\alpha)\,\text{Gamma}(\alpha_2; \alpha_\alpha, \beta_\alpha) \tag{18}$$
$$p(m_c) = \text{Gamma}(m_c; \alpha_m, \beta_m) \tag{19}$$
$$p(r) = \text{Gamma}(r; \alpha_r, \beta_r) \tag{20}$$
$$p(\phi_c) = \text{Beta}(\phi_c; \alpha_\phi, \beta_\phi) \tag{21}$$
$$p(w_i) = \mathcal{N}(w_i; \mu_h, \sigma_h^2) \tag{22}$$

where $1(\cdot)$ is the indicator function. The joint prior on the rates $p(\alpha_0, \alpha_1, \alpha_2) = p(\alpha_0)p(\alpha_1, \alpha_2)$ enforces identifiability and we set $\alpha_\alpha = 1$ and $\beta_\alpha = 0.01$. The prior over the mean step times was $\alpha_m = 2$ and $\beta_m = 0.02$. This prior has a peak at 50, a mean of 100, and it places significant mass over a broad range of m.The prior over the step direction probabilities was uniform over $[0, 1]$ with $\alpha_\phi = 1$ and $\beta_\phi = 1$. The prior over r used $\alpha_r = 2$ and $\beta_r = 1$. The history weights had the same prior as in the ramping model with $\mu_h = 0$ and $\sigma_h^2 = 10$. For the stepping model without history, the priors were those specified in (Latimer et al., 2015).

4.3.2 MCMC Overview

We used Markov chain Monte Carlo (MCMC) methods to obtain approximate samples from the posterior of the model parameters Θ given the data $\mathbf{y}$, $p(\Theta|\mathbf{y})$. Specifically, we used MCMC to approximately sample from the joint posterior of the parameters and the latents $p(\Theta, \mathbf{x}|\mathbf{y})$ and ignored the samples of the latents to obtain samples from $p(\Theta|\mathbf{y})$.

The following is an overview of the MCMC methods we used to obtain samples from $p(\Theta|\mathbf{y})$ (Latimer et al., 2015). First, we sampled a value for each latent variable given the parameters and the observed spike counts from the distribution $p(\mathbf{x}|\Theta, \mathbf{y})$. Then, conditioned on the new value of the latent variables and the data, we sampled new parameter values from $p(\Theta|\mathbf{x}, \mathbf{y})$. By repeating this procedure many times we obtained samples from the distribution $p(\Theta, \mathbf{x}|\mathbf{y})$. We marginalized over $\mathbf{x}$ to obtain the posterior distribution over the parameters $p(\Theta|\mathbf{y})$ by simply discarding the values of $\mathbf{x}$.

For each MCMC simulation, we simulated a chain of 60000 samples from $p(\Theta|\mathbf{y})$. We discarded the first 10000 samples from the chain as a burn in period. We kept every fifth sample afterwards, leaving $S = 10000$ samples from the posterior distribution over the parameters. For each posterior sample $\{\Theta^s\}_{s=1:S}$, we computed the likelihood of the data given the posterior sample $p(\mathbf{y}|\Theta^s)$ by marginalizing over the latent variables. In the ramping models, we used 5000 Monte Carlo samples of the latent trajectories given Θ^s to compute the likelihood. In the stepping model, we performed the marginalization by integrating the step times and step directions on a grid (Latimer et al., 2015).

For the variable duration dataset, we simulated two MCMC chains for the ramping, linear ramping with non-zero baseline, square root ramping with non-zero baseline, and stepping models (all with spike-history) to examine convergence in the MCMC chains before comparing these models. We assessed convergence using the potential scale reduction factor (PSRF) convergence diagnostic (Gelman et al., 2013) on the trial likelihoods from the two chains. We chose to monitor the convergence of the likelihoods because our model comparison is based solely on the likelihoods. If the diagnostic indicated that the two chains had not converged to the same likelihood distribution (PSRF > 1.1), we simulated additional chains until we obtained two chains that passed the diagnostic. This required increasing the number of burn in samples for a few cells.

4.3.3 MCMC for Ramping and Nonlinear Ramping Models

The MCMC sampling procedure for the ramping and nonlinear ramping models proceeded as follows. We first initialized the parameters to $\Theta^{(1)}$. We set $\beta_{1:C}$ by sampling C values from the distribution $\mathcal{N}(\beta_c; 0, 0.001)$ and sorting the values in the order of the coherence levels. We set the initial bound height γ to be a sample from a Gaussian distribution with mean equal to the average spike rate in the final time bin of in-RF choice trials and with unit variance. We sampled the initial x_0 from a Gaussian with mean equal to the average spike rate in the first time bin divided by the initial γ and with standard deviation 0.01. We constrained the initial x_0 to be in $[0.1, 0.9]$. We sampled the initial variance ω^2 uniformly in the range $[5\text{e}{-4}, 5\text{e}{-3}]$. We sampled the initial history weights $w_h^{(1)} \sim \mathcal{N}(w_h; 0, 0.1)$. With a non-zero baseline, we sampled the initial baseline parameter $b^{(1)} \sim \mathcal{N}(b; 0.5, 0.01)$ and also subtracted the baseline from the mean of the distribution for sampling γ.

After initializing the parameters, we alternated between sampling the latent diffusion paths conditioned on the current parameters and sampling new values of the parameters conditioned on the previous latent path. Formally, we obtained the s^{th} sample, for $s > 1$, with

$$\mathbf{x}^{(s)} \sim p(\mathbf{x}|\Theta^{(s-1)}, \mathbf{y}) \tag{23}$$

$$\beta_{1:C}^{(s)} \sim p(\beta_{1:C}|\mathbf{x}^{(s)}, x_0^{(s-1)}, \omega^{2,(s-1)}) \tag{24}$$

$$x_0^{(s)} \sim p(x_0|\mathbf{x}^{(s)}, \beta_{1:C}^{(s)}, \omega^{2,(s-1)}) \tag{25}$$

$$\omega^{2,(s)} \sim p(\omega^2|\mathbf{x}^{(s)}, \beta_{1:C}^{(s)}, x_0^{(s)}) \tag{26}$$

$$\gamma^{(s)}, b^{(s)}, \mathbf{w}^{(s)} \sim p(\gamma, b, \mathbf{w}|\mathbf{x}^{(s)}, \mathbf{y}). \tag{27}$$

For step (23), we used a particle filter to estimate the distribution of latent paths below the boundary and the distribution of boundary crossing times (Latimer et al., 2015). Given those two distributions, we used a backwards sampling scheme to sample the latent paths $\mathbf{x}^{(s)}$. We modified the firing rate observation likelihood in this step for each model to include the appropriate nonlinearity, baseline, and history dependence. We exploited conjugacy in steps (24), (25), and (26) for Gibbs steps, which were identical to those presented in Latimer et al. (2015).

For the final step (27), we used a manifold Metropolis-adjusted Langevin (MMALA) step to jointly sample the parameters $\boldsymbol{\theta} = [\gamma, b, \mathbf{w}]^{\mathrm{T}}$ (Girolami and Calderhead, 2011; Latimer et al., 2015). The vector of parameters $\boldsymbol{\theta}$ has dimension $J = 2 + H$, where H is the number of history weights. In the following derivation, for models with a subset of the parameters $\boldsymbol{\theta}$, the terms unrelated to the subset of parameters are disregarded. Each Metropolis step consisted of sampling a new value of the parameters $\boldsymbol{\theta}^\star$ from a proposal distribution $q(\boldsymbol{\theta}^\star|\boldsymbol{\theta}^{(s-1)}, \mathbf{y}, \mathbf{x}^{(s)})$ and accepting the newly sampled values (that is, set $\boldsymbol{\theta}^{(s)} = \boldsymbol{\theta}^\star$) with probability

$$p_{\text{accept}} = \min\left(1, \frac{p(\boldsymbol{\theta}^\star|\mathbf{x}^{(s)}, \mathbf{y})}{p(\boldsymbol{\theta}^{(s-1)}|\mathbf{x}^{(s)}, \mathbf{y})} \frac{q(\boldsymbol{\theta}^{(s-1)}|\boldsymbol{\theta}^\star, \mathbf{y}, \mathbf{x}^{(s)})}{q(\boldsymbol{\theta}^\star|\boldsymbol{\theta}^{(s-1)}, \mathbf{y}, \mathbf{x}^{(s)})}\right). \tag{28}$$

If the proposed values were not accepted then we set $\boldsymbol{\theta}^{(s)} = \boldsymbol{\theta}^{(s-1)}$. The proposal distribution used the gradient of the log likelihood plus log prior $\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta})$ and the Fisher information matrix plus Hessian of the log prior $\mathbf{G}(\boldsymbol{\theta})$

$$q(\boldsymbol{\theta}^\star|\boldsymbol{\theta}^{(s-1)}, \mathbf{y}, \mathbf{x}^{(s)}) = \mathcal{N}\left(\boldsymbol{\theta}^\star; \boldsymbol{\theta}^{(s-1)} + \epsilon_r^2 \frac{1}{2}\mathbf{G}^{-1}(\boldsymbol{\theta}^{(s-1)})\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta}^{(s-1)}),\ \epsilon_r^2\mathbf{G}^{-1}(\boldsymbol{\theta}^{(s-1)})\right) \tag{29}$$

where $\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta}) = [\frac{\partial}{\partial\theta_1}\mathcal{L}(\boldsymbol{\theta}), \frac{\partial}{\partial\theta_2}\mathcal{L}(\boldsymbol{\theta}), \cdots, \frac{\partial}{\partial\theta_J}\mathcal{L}(\boldsymbol{\theta})]^{\mathrm{T}}$. The step size ϵ_r was initialized to 0.05 and gradually increased to 1 during the burn in period.

We define the firing rate function

$$\lambda(x_{i,t}^{(s)}, \boldsymbol{\theta}) = (f(x_{i,t}^{(s)}\,\gamma) + b)g_{i,t} \tag{30}$$

that is in general a function of γ, the baseline parameter b, the sampled latent path, and the history for trial i and time t. The log likelihood plus log prior is a sum over N trials and T_i time points on each trial

$$\mathcal{L}(\boldsymbol{\theta}) = \log p(\mathbf{y}|\mathbf{x}^{(s)}, \boldsymbol{\theta}) + \log p(\boldsymbol{\theta}) \tag{31}$$

$$= \left(\sum_{i=1}^{N}\sum_{t=1}^{T_i} y_{i,t}\log\lambda(x_{i,t}^{(s)}, \boldsymbol{\theta})\Delta - \lambda(x_{i,t}^{(s)}, \boldsymbol{\theta})\Delta - \log\Gamma(y_{i,t}+1)\right) + \log p(\boldsymbol{\theta}). \tag{32}$$

The gradient of $\mathcal{L}(\boldsymbol{\theta})$ is defined by the derivatives with respect to each parameter θ_j

$$\frac{\partial}{\partial\theta_j}\mathcal{L}(\boldsymbol{\theta}) = \left(\sum_{i=1}^{N}\sum_{t=1}^{T_i}\lambda'_{\theta_j}(x_{i,t}^{(s)}, \boldsymbol{\theta})\left(\frac{y_{i,t}}{\lambda(x_{i,t}^{(s)}, \boldsymbol{\theta})} - \Delta\right)\right) + \frac{\partial}{\partial\theta_j}\log p(\boldsymbol{\theta}) \tag{33}$$

where $\lambda'_{\theta_j}(x_{i,t}^{(s)}, \boldsymbol{\theta})$ is the derivative of the rate function with respect to θ_j. The $J \times J$ matrix $\mathbf{G}(\boldsymbol{\theta})$ is

$$\mathbf{G}_{j,k}(\boldsymbol{\theta}) = -\mathbb{E}_{\mathbf{y}|\mathbf{x}^{(s)},\boldsymbol{\theta}}\left[\frac{\partial^2}{\partial\theta_j\,\partial\theta_k}\mathcal{L}(\boldsymbol{\theta})\right] = -\frac{\partial^2}{\partial\theta_j\,\partial\theta_k}\log p(\boldsymbol{\theta}) + \sum_{i=1}^{N}\sum_{t=1}^{T_i}\Delta\frac{\lambda'_{\theta_j}(x_{i,t}^{(s)},\boldsymbol{\theta})\lambda'_{\theta_k}(x_{i,t}^{(s)},\boldsymbol{\theta})}{\lambda(x_{i,t}^{(s)},\boldsymbol{\theta})}. \quad (34)$$

The derivatives $\frac{\partial}{\partial\theta_j}\lambda(x_{i,t}^{(s)},\boldsymbol{\theta})$ for each parameter are

$$\frac{\partial}{\partial\gamma}\lambda(x_{i,t}^{(s)},\boldsymbol{\theta}) = f'_\gamma(x_{i,t}^{(s)},\gamma)\,g_{i,t} \quad (35)$$

$$\frac{\partial}{\partial b}\lambda(x_{i,t}^{(s)},\boldsymbol{\theta}) = g_{i,t} \quad (36)$$

$$\frac{\partial}{\partial w_h}\lambda(x_{i,t}^{(s)},\boldsymbol{\theta}) = (f(x_{i,t}^{(s)},\gamma)+b)\,g_{i,t}\,y_{i,t-h}$$

where $f'_\gamma(x_{i,t}^{(s)},\gamma)$ is the derivative of each nonlinearity with input $x_{i,t}^{(s)},\gamma$ with respect to γ

$$f'_\gamma(x_{i,t}^{(s)},\gamma) = \begin{cases} \frac{x_{i,t}^{(s)}}{1+\exp(-x_{i,t}^{(s)}\gamma)}, & \text{softplus (linear)} \\ \frac{1}{2}(\log(1+\exp(x_{i,t}^{(s)}\gamma)))^{-1/2}\frac{x_{i,t}^{(s)}}{1+\exp(-x_{i,t}^{(s)}\gamma)}, & \text{soft sqrt} \\ 2\log(1+\exp(x_{i,t}^{(s)}\gamma))\frac{x_{i,t}^{(s)}}{1+\exp(-x_{i,t}^{(s)}\gamma)}, & \text{soft quad} \\ x_{i,t}^{(s)}\exp(x_{i,t}^{(s)}\gamma), & \text{exponential.} \end{cases}$$

The derivative and second derivative of the log prior on γ are

$$\frac{\partial}{\partial\gamma}\log p(\gamma) = \frac{\partial}{\partial\gamma}\left(\alpha_\gamma\log(\beta_\gamma) - \log\Gamma(\alpha_\gamma) + (\alpha_\gamma-1)\log\gamma - \beta_\gamma\gamma\right) = \frac{\alpha_\gamma-1}{\gamma} - \beta_\gamma \quad (37)$$

$$\frac{\partial^2}{\partial\gamma^2}\log p(\gamma) = \frac{\partial}{\partial\gamma}\left(\frac{\alpha_\gamma-1}{\gamma} - \beta_\gamma\right) = -\frac{\alpha_\gamma-1}{\gamma^2}. \quad (38)$$

Similarly, for the baseline b these quantities are

$$\frac{\partial}{\partial b}\log p(b) = \frac{\alpha_b-1}{b} - \beta_b \quad (39)$$

$$\frac{\partial^2}{\partial b^2}\log p(b) = -\frac{\alpha_b-1}{b^2}. \quad (40)$$

The first two derivatives of the log likelihood of the history weights w_h are

$$\frac{\partial}{\partial w_h}\log p(w_h) = \frac{\partial}{\partial w_h}\left(-\frac{1}{2}\log 2\pi - \frac{1}{2}\log\sigma_h^2 - \frac{1}{2}\frac{(w_h-\mu_h)^2}{\sigma_h^2}\right) = -\frac{w_h-\mu_h}{\sigma_h^2} \quad (41)$$

$$\frac{\partial^2}{\partial w_h^2}\log p(w_h) = \frac{\partial}{\partial w_h}\left(-\frac{w_h-\mu_h}{\sigma_h^2}\right) = -\frac{1}{\sigma_h^2}. \quad (42)$$

Each of the priors are independent and therefore the prior terms contributing to the off-diagonal elements of $\mathbf{G}_{j,k}(\boldsymbol{\theta})$ are zero.

4.3.4 MCMC for Stepping Models

In the stepping model we alternated between sampling the latent step times $\mathbf{z}$ and directions $\mathbf{d}$ and sampling the parameters of the model Θ. We first initialized the parameters to $\Theta^{(1)}$. We set $\alpha_0^{(1)}$ to the firing rate in the first time bin, $\alpha_1^{(1)}$ to the firing rate in the final time bin of out-RF choice trials, and $\alpha_2^{(1)}$ to the firing rate in the final time bin of in-RF choice trials. We then added Gaussian noise to each $\alpha_{0:2}^{(1)}$. We sampled $r^{(1)} \sim \mathcal{N}(r; 1, 0.0025)$, $m_c^{(1)} \sim \mathcal{N}(m_c; 30, 25)$, $\phi_c^{(1)} \sim \mathcal{N}(\phi_c; 0.5, 0.05)$, and $w_h^{(1)} \sim \mathcal{N}(w_h; 0, 0.1)$.

After initialization, we performed the following sequence of steps to obtain the $s^{\text{th}} > 1$ sample

$$\mathbf{z}^{(s)}, \mathbf{d}^{(s)} \sim p(\mathbf{z}, \mathbf{d} | \Theta^{(s-1)}, \mathbf{y}) \tag{43}$$

$$\alpha_{0:2}^{(s)}, \mathbf{w}^{(s)} \sim p(\alpha_{0:2}, \mathbf{w} | \mathbf{z}^{(s)}, \mathbf{d}^{(s)}, \mathbf{y}) \tag{44}$$

$$\phi_{1:C}^{(s)} \sim p(\phi_{1:C} | \mathbf{d}^{(s)}) \tag{45}$$

$$m_{1:C}^{(s)} \sim p(m_{1:C} | \mathbf{z}^{(s)}, r^{(s-1)}) \tag{46}$$

$$r^{(s)} \sim p(r | \mathbf{z}^{(s)}, m_{1:C}^{(s)}). \tag{47}$$

We sampled the step directions and step times (step 43) by sampling from the distribution computed on a grid, truncated at 1500 time bins (Latimer et al., 2015), with the history dependence included in the observation likelihood. We employed Beta-Bernoulli conjugacy to directly sample the step direction probabilities (step 45) using a Gibbs' step.

Sampling the rates and spike history filters

We used an MMALA step in the stepping model to sample the rates $\alpha_{0:2}$ and history weights $\mathbf{w}$ with proposal distribution

$$q(\boldsymbol{\theta}^\star | \boldsymbol{\theta}^{(s-1)}, \mathbf{y}, \mathbf{z}^{(s)}, \mathbf{d}^{(s)}) = \mathcal{N}\left(\boldsymbol{\theta}^\star; \boldsymbol{\theta}^{(s-1)} + \epsilon_s^2 \frac{1}{2} \mathbf{G}^{-1}(\boldsymbol{\theta}^{(s-1)}) \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^{(s-1)}),\ \epsilon_s^2 \mathbf{G}^{-1}(\boldsymbol{\theta}^{(s-1)})\right) \tag{48}$$

where $\boldsymbol{\theta} = [\alpha_{0:2}, \mathbf{w}]^{\mathrm{T}}$. The step size ϵ_s was gradually increased from 0.05 to 1 during the burn in period. The firing rate for trial i and time t is

$$\lambda(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t) = \alpha(z_i^{(s)}, d_i^{(s)}, t) g_{i,t} \tag{49}$$

where

$$\alpha(z_i^{(s)}, d_i^{(s)}, t) = \begin{cases} \alpha_0, & z_i^{(s)} \geq t \\ \alpha_1, & z_i^{(s)} < t, d_i^{(s)} = 1 \\ \alpha_2, & z_i^{(s)} < t, d_i^{(s)} = 2. \end{cases} \tag{50}$$

The MMALA step uses $\mathcal{L}(\boldsymbol{\theta})$, $\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta})$, and $\mathbf{G}(\boldsymbol{\theta})$. The log likelihood plus log prior is

$$\mathcal{L}(\boldsymbol{\theta}) = \log p(\mathbf{y} | \mathbf{z}^{(s)}, \mathbf{d}^{(s)}, \boldsymbol{\theta}) + \log p(\boldsymbol{\theta}) \tag{51}$$

$$= \left(\sum_{i=1}^{N} \sum_{t=1}^{T_i} y_{i,t} \log \lambda(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t) \Delta - \lambda(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t) \Delta - \log \Gamma(y_{i,t} + 1) \right) + \log p(\boldsymbol{\theta}). \tag{52}$$

The gradient of the log likelihood plus log prior is

$$\frac{\partial}{\partial \theta_j}\mathcal{L}(\boldsymbol{\theta}) = \left(\sum_{i=1}^{N}\sum_{t=1}^{T_i}\lambda'_{\theta_j}(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t)\left(\frac{y_{i,t}}{\lambda(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t)} - \Delta\right)\right) + \frac{\partial}{\partial \theta_j}\log p(\boldsymbol{\theta}). \tag{53}$$

The elements of the Fisher information matrix plus the Hessian of the log prior are

$$\mathbf{G}_{j,k}(\boldsymbol{\theta}) = -\frac{\partial^2}{\partial\theta_j\,\partial\theta_k}\log p(\boldsymbol{\theta}) + \sum_{i=1}^{N}\sum_{t=1}^{T_i}\Delta\frac{\lambda'_{\theta_j}(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t)\lambda'_{\theta_k}(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t)}{\lambda(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t)}. \tag{54}$$

The derivative of the rate function given $z_i^{(s)}$ and $d_i^{(s)}$ is

$$\frac{\partial}{\partial \alpha_j}\lambda(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t) = \begin{cases} g_{i,t}, & \alpha(z_i^{(s)}, d_i^{(s)}, t) = \alpha_j \\ 0, & \text{otherwise.} \end{cases} \tag{55}$$

The derivative with respect to the history weights is

$$\frac{\partial}{\partial w_h}\lambda(z_i^{(s)}, d_i^{(s)}, \boldsymbol{\theta}, t) = \alpha(z_i^{(s)}, d_i^{(s)}, t)\, g_{i,t}\, y_{i,t-h}. \tag{56}$$

When evaluating the first and second derivatives of the log prior for the proposal distribution in (53) and (54), we used independent priors on α_1 and α_2

$$q(\alpha_1) = \Gamma(\alpha_1; \alpha_\alpha, \beta_\alpha) \tag{57}$$

$$q(\alpha_2) = \Gamma(\alpha_2; \alpha_\alpha, \beta_\alpha). \tag{58}$$

This simplifies computation of the gradient and Hessian for the proposal distribution. The derivatives of the log prior for $\alpha_{0:2}$ and $\mathbf{w}$ have the same form as γ and $\mathbf{w}$ in the ramping model.

We note that if across all trials the cell was never in state α_j then $\frac{\partial}{\partial \alpha_j}\mathcal{L}(\boldsymbol{\theta})$ is zero and the row and column of $\mathbf{G}(\boldsymbol{\theta})$ corresponding to α_j is zero. This matrix must be nonsingular such that we can use its inverse in the proposal distribution. Therefore, if this occurred, although rare, we set the diagonal element of each zero row and column to one.

Sampling step time means and shape

We sampled the mean $m_{1:C}$ and shape r parameters of the negative binomial distribution over the step times using Metropolis steps. The probability of a step time z_i on trial i with coherence c_i in terms of m_{c_i} and r is

$$p(z_i|r, m_{c_i}) = \frac{\Gamma(z_i + r)}{\Gamma(z_i + 1)\Gamma(r)}\left(\frac{m_{c_i}}{m_{c_i} + r}\right)^{z_i}\left(\frac{r}{m_{c_i} + r}\right)^{r}. \tag{59}$$

We alternated between sampling each m_c conditioned on r and sampling r conditioned on each m_c.

The proposal distribution for each m_c is

$$Q(m_c^\star|m_c^{(s-1)}, r^{(s-1)}, \mathbf{z}^{(s)}) \sim \mathcal{N}\left(m_c^\star;\, m_c^{(s-1)} + \frac{1}{2}\epsilon_m^2\mathbf{G}^{-1}(m_c^{(s-1)})\frac{\partial}{\partial m_c}\mathcal{L}(m_c^{(s-1)}),\, \epsilon_m^2\mathbf{G}^{-1}(m_c^{(s-1)})\right). \tag{60}$$

We gradually increased ϵ_m from 0.05 to 1 during the burn in period. The log likelihood of m_c plus log prior is the sum of the likelihoods of the step time z_i for each trial with coherence c

$$\mathcal{L}(m_c) = \sum_{i \in c_i = c} \log p(z_i | r, m_{c_i}) + \log p(m_c) \tag{61}$$

$$= \sum_{i \in c_i = c} \left[z_i \log \left(\frac{m_{c_i}}{m_{c_i} + r} \right) + r \log \left(\frac{r}{m_{c_i} + r} \right) \right] + \log p(m_c) + const. \tag{62}$$

The derivative of the log likelihood of m_c plus log prior with respect to m_c is

$$\frac{\partial}{\partial m_c} \mathcal{L}(m_c) = \sum_{i \in c_i = c} \left[z_i \frac{r}{m_{c_i}(m_{c_i} + r)} - \frac{r}{m_{c_i} + r} \right] + \frac{\partial}{\partial m_c} \log p(m_c). \tag{63}$$

The Fisher information plus the Hessian of the log prior is

$$\mathbf{G}(m_c) = -\mathbb{E}_{\mathbf{z}|m_c, r} \left[\frac{\partial^2}{\partial m_c^2} \mathcal{L}(m_c) \right] = -\frac{\partial^2}{\partial m_c^2} \log p(m_c) + \sum_{i \in c_i = c} \left[\frac{r}{m_{c_i}(m_{c_i} + r)} \right] \tag{64}$$

where we have used $\mathbb{E}_{z_i | m_{c_i}, r}[z_i] = m_{c_i}$.

The proposal distribution for r is

$$Q(r^\star | r^{(s-1)}, m_{1:C}^{(s)}, \mathbf{z}^{(s)}) \sim \mathcal{N}\left(r^\star;\, r^{(s-1)} + \frac{1}{2}\epsilon^2 \frac{\partial}{\partial r} \mathcal{L}(r^{(s-1)}),\, \epsilon^2 \right) \tag{65}$$

and the log likelihood of r plus log prior is a sum over all trials

$$\mathcal{L}(r) = \sum_{i=1}^{N} \left[\log \Gamma(z_i + r) - \log \Gamma(r) + z_i \log \left(\frac{m_{c_i}}{m_{c_i} + r} \right) + r \log \left(\frac{r}{m_{c_i} + r} \right) \right] + \log p(r) + const. \tag{66}$$

The derivative of the log likelihood plus log prior with respect to r is

$$\frac{\partial}{\partial r} \mathcal{L}(r) = \frac{\partial}{\partial r} \log p(r) + \sum_{i=1}^{N} \left[\psi(z_i + r) - \psi(r) - \frac{z_i}{m_{c_i} + r} + \log \left(\frac{r}{m_{c_i} + r} \right) + \frac{m_{c_i}}{m_{c_i} + r} \right] \tag{67}$$

where ψ is the digamma function $\psi(r) = \Gamma'(r)/\Gamma(r)$. The step size ϵ was initialized to 0.075 and was adjusted throughout the burn in period, after which it was fixed.

4.4 Model Comparison: WAIC

We used the WAIC to compare the models (Watanabe, 2010; Gelman et al., 2014; Vehtari et al., 2017; Piironen and Vehtari, 2017). The WAIC estimates the expected generalization of a fit model to new data from the true data generating distribution. This corresponds to an estimate of how well the model would predict spike trains recorded on new trials. As generally the experimenter neither has access to unlimited data nor the true data generating distribution, the WAIC and other information criterion methods estimate the expected generalization using how well the model describes the in-sample data with a correction factor.

The WAIC is a function of the probability of the data given each posterior sample, $\{p(\mathbf{y}_i|\Theta^s)\}_{s=1:S}$, for each trial i. We used the formula in (Gelman et al., 2014) to compute the WAIC across N trials as

$$\text{WAIC} = -2\Bigg(\underbrace{\sum_{i=1}^{N}\log\left(\frac{1}{S}\sum_{s=1}^{S}p(\mathbf{y}_i|\Theta^s)\right)}_{\text{lppd}} - \underbrace{\sum_{i=1}^{N}\text{Var}[\log p(\mathbf{y}_i|\Theta)]}_{p_{\text{WAIC}}}\Bigg). \tag{68}$$

The first term is the log pointwise predictive density (lppd) and it describes how well the model predicts the data to which it was fit. A strength of the WAIC is the lppd averages over the posterior rather than conditioning on a point estimate of the parameters. The second term p_{WAIC} is a penalty that corrects the bias induced by estimating the expected generalization to new data from the lppd. The penalty term p_{WAIC} is computed for each trial as the variance of the log likelihoods of a trial across the posterior samples $\{\Theta^s\}_{s=1:S}$, and therefore is guaranteed to be non-negative because it is a sum of variances, another strength of the WAIC (Gelman et al., 2014). Additional advantages of the WAIC are theoretical results showing its asymptotic equivalence to Bayesian leave-one-out cross validation, its applicability to singular statistical models, and its computational efficiency when compared to leave-one-out cross validation (Watanabe, 2010; Gelman et al., 2014; Piironen and Vehtari, 2017). For all of these reasons we used it to compare the relative fits of the models.

In model comparison, we computed the WAIC difference between two models

$$\Delta\text{WAIC} = \text{WAIC}_{\text{model 1}} - \text{WAIC}_{\text{model 2}}. \tag{69}$$

Since lower WAIC values are better, a positive difference favors model two while a negative difference favors model one. In some cases, we normalized the WAIC by the number of trials to put comparisons with differing numbers of trials on the same scale. We also considered the WAIC difference on subsets of trials by only summing across trials of certain conditions. We set the stepping model with spike-history as model 2 in model comparison with other models. Therefore, positive WAIC differences favor the stepping with spike-history model over the alternative model in these comparisons.

We quantified uncertainty in the model comparison using standard errors of the WAIC differences across trials

$$\text{se}(\Delta\text{WAIC}) = \sqrt{N\,\text{Var}(\Delta\text{WAIC}_i)} \tag{70}$$

where ΔWAIC_i is the WAIC difference computed for trial i.

4.5 Simulated Data

We computed simulated latent trajectories and PSTHs from a fit model using the following procedure. For 40 random samples from the posterior over the parameters, we simulated a spike train for each trial conditioned on the pre-trial spiking activity and parameters of the model. For a few cells, simulated spike trains from the models with spike-history generated unrealistically large numbers of spikes due to self-excitation. We enforced realistic spike trains in these cases by setting the multiplicative history gain to unitary whenever the generated spike history effect was larger than the largest inferred history gain in the data. We averaged the simulated spike trains corresponding to each coherence to compute the coherence-dependent simulated PSTHs.

We computed the autocorrelation of the observed and simulated data for each neuron with the normalized autocorrelation function

$$R(\tau) = \frac{1}{m}\left(\frac{1}{N_\tau}\sum_t y_t y_{t-\tau}\right) - m \tag{71}$$

where y_t is the spike count at time t and m is the mean spike count. The sum was computed over all valid time bins and N_τ is the number of valid time bins.

For computing the fraction at the boundaries in (Figure 5), we used the following criterion. A time point t was classified as at a boundary if $x_t \leq 0$ or $x_t \geq 1$.

Acknowledgements

This work was supported by the McKnight Foundation (J.W.P.), NSF CAREER Award IIS-1150186 (J.W.P.), and grants from the NIH (NEI grant EYE017366, to A.C.H., and NIMH grant MH099611, to A.C.H. & J.W.P.). David Zoltowski is supported by NIH grant T32MH065214. Jacob Yates is an Open Philanthropy Fellow of the Life Sciences Research Foundation. Kenneth Latimer is a Chicago Fellow.

Author Contributions

Conceptualization, D.M.Z., K.W.L., A.C.H., and J.W.P.; Methodology, D.M.Z., K.W.L., A.C.H., and J.W.P.; Software, D.M.Z.; Formal Analysis, D.M.Z.; Investigation, J.L.Y.; Data Curation, J.L.Y.; Writing - Original Draft, D.M.Z. and J.W.P.; Writing - Review & Editing, D.M.Z., K.W.L., J.L.Y., A.C.H., and J.W.P.; Visualization, D.M.Z.; Supervision, J.W.P.; Funding Acquisition, A.C.H. and J.W.P..

Declaration of Interests

The authors declare no competing interests.

References

Bollimunta, A., Totten, D., and Ditterich, J. (2012). Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex. *Journal of Neuroscience*, 32(37):12684–12701.

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., and Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque mt. *Visual neuroscience*, 13(1):87–100.

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12(12):4745–4765.

Brunton, B. W., Botvinick, M. M., and Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, 340(6128):95–98.

Chandrasekaran, C., Soldado-Magraner, J., Peixoto, D., Sahani, M., and Shenoy, K. V. (2016). Do decision-related firing rates of dorsal premotor cortex neurons “ramp” or “step” on single trials? In *Society for Neuroscience Annual Meeting*.

Charles, A. S., Park, M., Weller, J. P., Horwitz, G. D., and Pillow, J. W. (2018). Dethroning the fano factor: A flexible, model-based approach to partitioning neural variability. *Neural Computation*, 30(4):1012–1045.

Churchland, A. K. and Kiani, R. (2016). Three challenges for connecting model to mechanism in decision-making. *Current opinion in behavioral sciences*, 11:74–80.

Churchland, A. K., Kiani, R., Chaudhuri, R., Wang, X.-J., Pouget, A., and Shadlen, M. N. (2011). Variance as a signature of neural computations during decision making. *Neuron*, 69(4):818–831.

Churchland, A. K., Kiani, R., and Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature neuroscience*, 11(6):693.

de Lafuente, V., Jazayeri, M., and Shadlen, M. N. (2015). Representation of accumulating evidence for a decision in two parietal areas. *Journal of Neuroscience*, 35(10):4306–4318.

Ditterich, J. (2006a). Evidence for time-variant decision making. *European Journal of Neuroscience*, 24(12):3628–3641.

Ditterich, J. (2006b). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Networks*, 19(8):981–1012.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). Bayesian data analysis, 3rd edition.

Gelman, A., Hwang, J., and Vehtari, A. (2014). Understanding predictive information criteria for bayesian models. *Statistics and Computing*, 24(6):997–1016.

Girolami, M. and Calderhead, B. (2011). Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214.

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.*, 30:535–574.

Goris, R. L., Movshon, J. A., and Simoncelli, E. P. (2014). Partitioning neuronal variability. *Nature neuroscience*, 17(6):858.

Hanks, T. D., Ditterich, J., and Shadlen, M. N. (2006). Microstimulation of macaque area lip affects decision-making in a motion discrimination task. *Nature neuroscience*, 9(5):682.

Hanks, T. D., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C., and Brody, C. D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*, 520(7546):220–223.

Hanks, T. D. and Summerfield, C. (2017). Perceptual decision making in rodents, monkeys, and humans. *Neuron*, 93(1):15–31.

Huk, A. C., Katz, L. N., and Yates, J. L. (2017). The role of the lateral intraparietal area in (the study of) decision making. *Annual review of neuroscience*, 40:349–372.

Huk, A. C. and Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience*, 25(45):10420–10436.

Janssen, P. and Shadlen, M. N. (2005). A representation of the hazard rate of elapsed time in macaque area lip. *Nature neuroscience*, 8(2):234.

Katz, L. N., Yates, J. L., Pillow, J. W., and Huk, A. C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature*, 535(7611):285–288.

Kiani, R., Hanks, T. D., and Shadlen, M. N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *Journal of Neuroscience*, 28(12):3017–3029.

Kiani, R. and Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *science*, 324(5928):759–764.

Kira, S., Yang, T., and Shadlen, M. N. (2015). A neural implementation of wald?s sequential probability ratio test. *Neuron*, 85(4):861–873.

Latimer, K. W., Huk, A. C., and Pillow, J. W. (2017). No cause for pause: new analyses of ramping and stepping dynamics in lip (rebuttal to response to reply to comment on latimer et al. 2015). *bioRxiv*.

Latimer, K. W., Yates, J. L., Meister, M. L., Huk, A. C., and Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, 349(6244):184–187.

Mazurek, M. E., Roitman, J. D., Ditterich, J., and Shadlen, M. N. (2003). A role for neural integrators in perceptual decision making. *Cerebral cortex*, 13(11):1257–1269.

Meister, M. L., Hennig, J. A., and Huk, A. C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *Journal of Neuroscience*, 33(6):2254–2267.

Morcos, A. S. and Harvey, C. D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature neuroscience*, 19(12):1672.

Newsome, W. and Pare, E. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (mt). *Journal of Neuroscience*, 8(6):2201–2211.

Palmer, J., Huk, A. C., and Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of vision*, 5(5):1–1.

Park, I. M., Meister, M. L., Huk, A. C., and Pillow, J. W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature neuroscience*, 17(10):1395–1403.

Piironen, J. and Vehtari, A. (2017). Comparison of bayesian predictive methods for model selection. *Statistics and Computing*, 27(3):711–735.

Roitman, J. D. and Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, 22(21):9475–9489.

Scott, B. B., Constantinople, C. M., Akrami, A., Hanks, T. D., Brody, C. D., and Tank, D. W. (2017). Fronto-parietal cortical circuits encode accumulated evidence with a diversity of timescales. *Neuron*, 95(2):385–398.

Shadlen, M. N. and Kiani, R. (2013). Decision making as a window on cognition. *Neuron*, 80(3):791–806.

Shadlen, M. N., Kiani, R., Newsome, W. T., Gold, J. I., Wolpert, D. M., Zylberberg, A., Ditterich, J., de Lafuente, V., Yang, T., and Roitman, J. (2016). Comment on "single-trial spike trains in parietal cortex reveal discrete steps during decision-making". *Science*, 351(6280):1406–1406.

Shadlen, M. N. and Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the national academy of sciences*, 93(2):628–633.

Shadlen, M. N. and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of Neurophysiology*, 86(4):1916–1936. PMID: 11600651.

Stine, G. M., Zylberberg, A., and Shadlen, M. N. (2018). Disentangling evidence integration from memoryless strategies in perceptual decision making. *Cosyne*.

Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., and Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, 93(2):1074–1089.

Vehtari, A., Gelman, A., and Gabry, J. (2017). Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and Computing*, 27(5):1413–1432.

Watanabe, S. (2010). Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11(Dec):3571–3594.

Weber, A. I. and Pillow, J. W. (2017). Capturing the dynamical repertoire of single neurons with generalized linear models. *Neural Computation*, 29(12):3260–3289.

Yang, T. and Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, 447(7148):1075.

Yates, J. L., Park, I. M., Katz, L. N., Pillow, J. W., and Huk, A. C. (2017). Functional dissection of signal and noise in mt and lip during decision-making. *Nature neuroscience*, 20(9):1285.

Zhao, X. and Kording, K. (2018). Rate fluctuations not steps dominate lip activity during decision-making. *bioRxiv*, page 249672.

Zylberberg, A. and Shadlen, M. N. (2016). Cause for pause before leaping to conclusions about stepping. *bioRxiv*, page 085886.

A Supplementary Figures

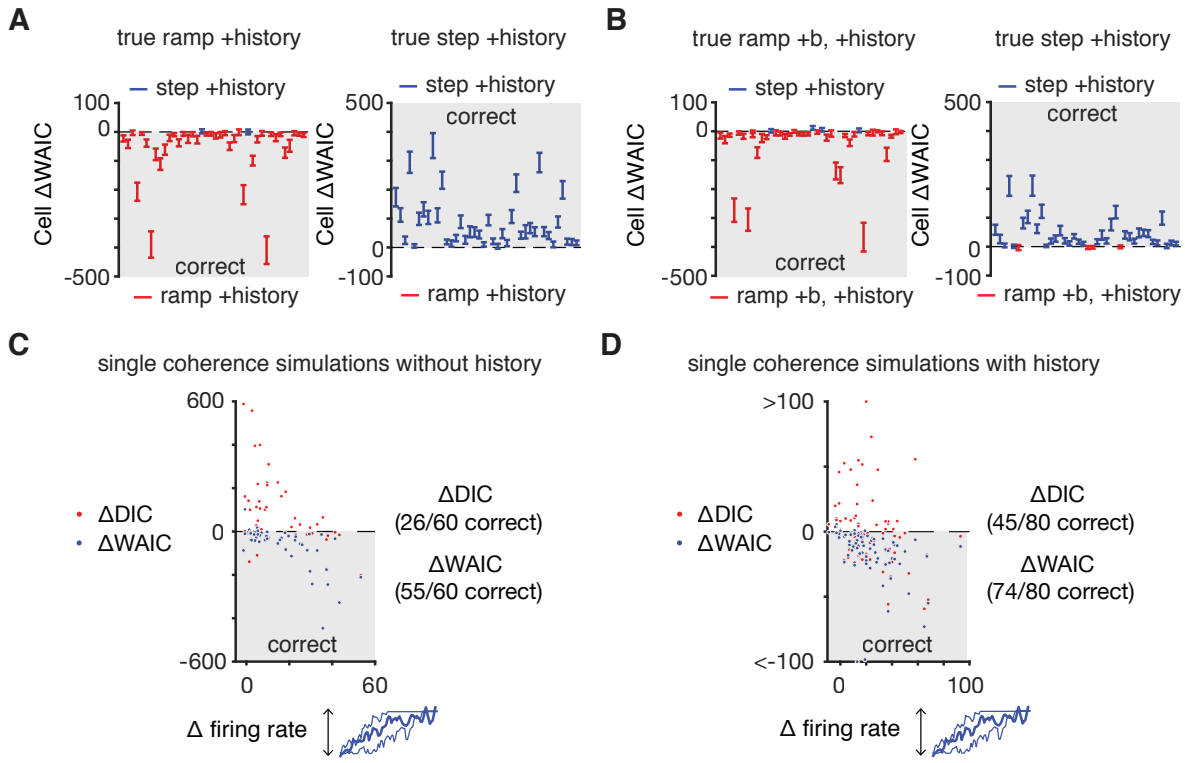

Figure 7. The WAIC correctly identifies simulated data from different models. Related to Figures 3 and 4. **A.** The WAIC reliably distinguished data from the ramping and stepping models with spike-history simulated from the mean parameters of the ramping (left) and stepping (right) models with spike-history for each cell. The number of simulated trials per motion coherence was matched to the true number for each cell. **B.** Same as **A**, except for simulated data from the ramping with non-zero baseline and history model. **C.-D.** Model comparison results from simulated single-coherence ramping data as a function of the change in firing rate from the beginning to the end of the trial using the DIC and WAIC. The WAIC correctly identifies data simulated from single-coherence ramping models (Chandrasekaran et al., 2016).

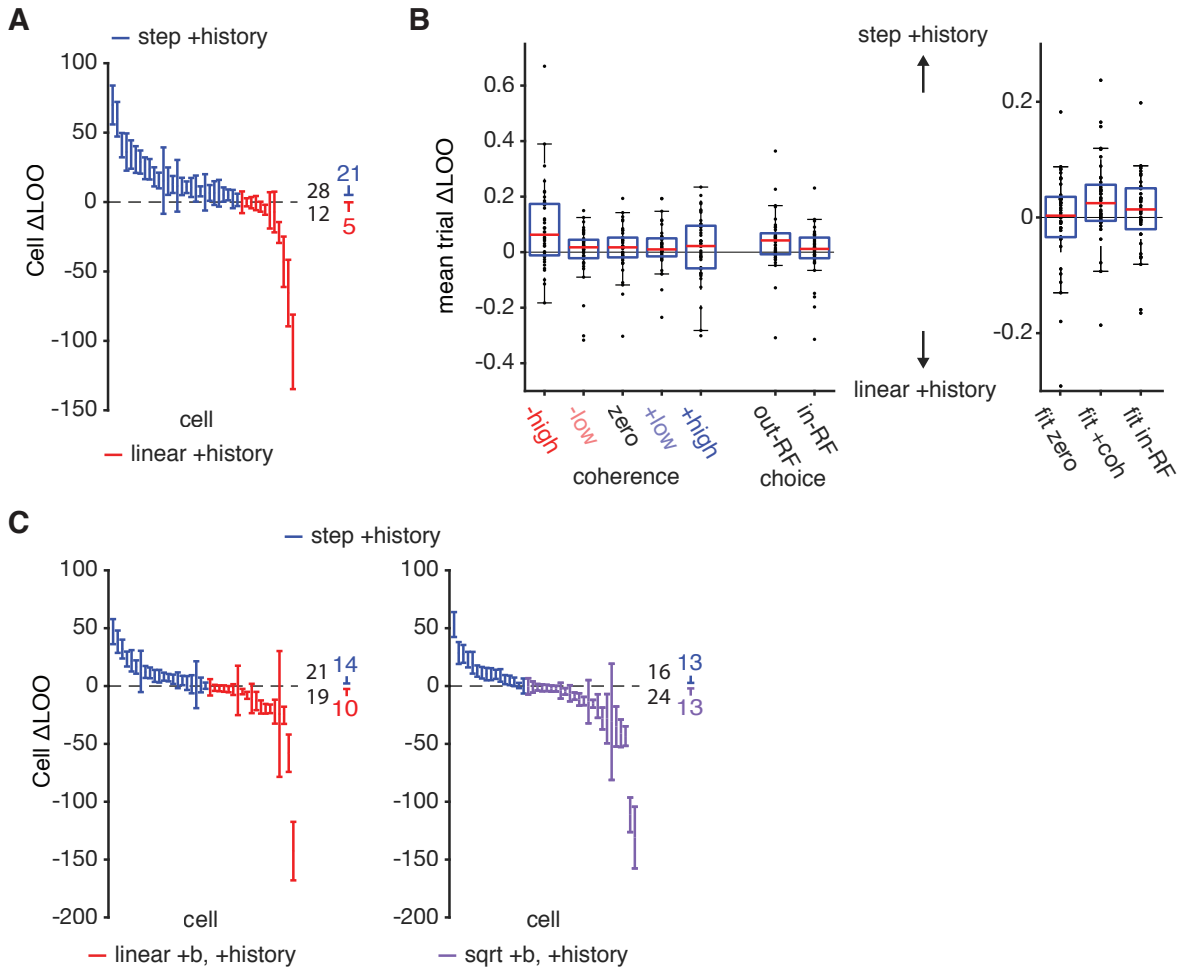

Figure 8. Model comparison using leave-one-out cross validation. Related to Figures 3 and 4. We used PSIS-LOO to estimate the Bayesian leave-one-out cross validation performance (LOO) on held out trials (Vehtari et al., 2017). In Bayesian LOO, predictive performance on a held-out trial is evaluated by integrating over the posterior distribution of the parameters conditioned on the rest of the trials. PSIS-LOO uses importance sampling to estimate this intergral. The proposal distribution is the posterior distribution conditioned on all of the trials (which we sample from in our MCMC procedure) and the importance weights are stabilized through regularization. Model comparison using LOO provided nearly identical results to the model comparison using WAIC.

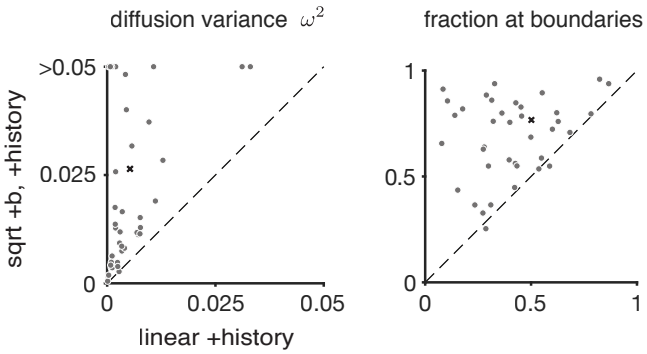

Figure 9. Related to Figure 5. Quantitative comparison of simulated trajectories from the square root ramping model with non-zero baseline and history against the linear ramping model with history and zero baseline. *Left:* Diffusion variance increased in the square root ramping model with non-zero baseline. *Right:* Fraction of time the simulated latent firing rates were equal to the baseline rate or upper absorbing boundary. Cell 1 is marked by an (x) in these plots.

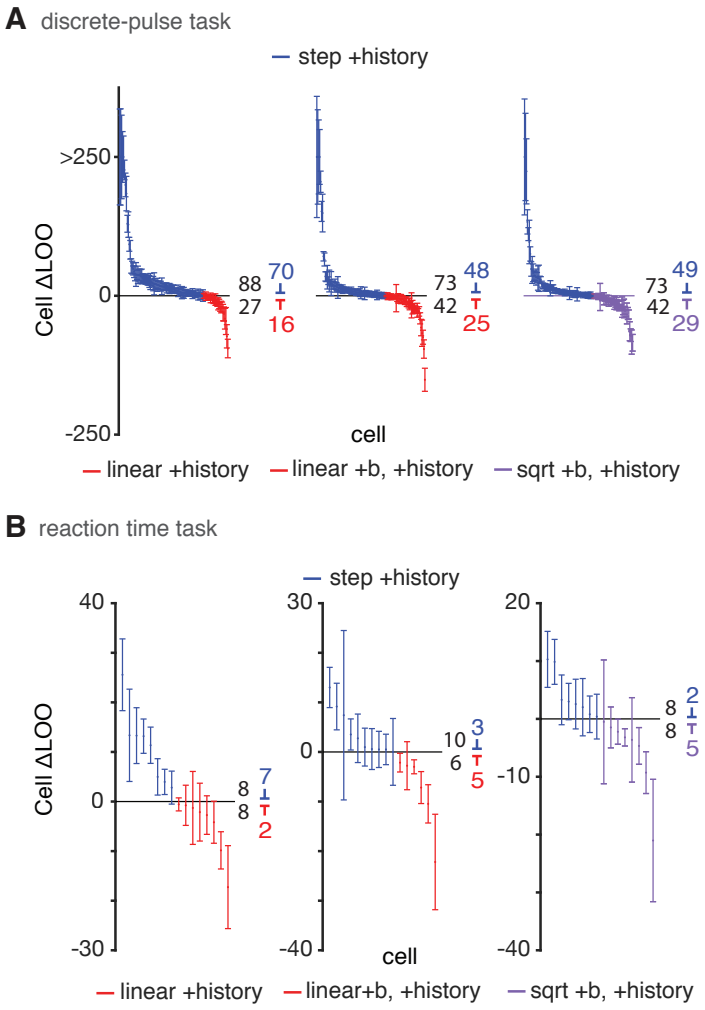

Figure 10. Model comparison using leave-one-out cross validation (LOO) for the discrete-pulse and reaction time data. Related to Figure 6. LOO was estimated using PSIS-LOO (Vehtari et al., 2017, see Figure 8). Model comparison using LOO for these datasets provided very similar results to model comparison using WAIC.