

Mice alternate between discrete strategies during perceptual decision-making

Zoe C. Ashwood,^{*1, 2}, Nicholas A. Roy², Iris R. Stone², The International Brain Laboratory³, Anne K. Churchland⁴, Alexandre Pouget⁵, and Jonathan W. Pillow,^{†2, 6}

¹Dept. of Computer Science, Princeton University

²Princeton Neuroscience Institute

³University College London

⁴David Geffen School of Medicine, The University of California, Los Angeles

⁵Faculty of Medicine & Dept. of Basic Neurosciences, University of Geneva

⁶Dept. of Psychology, Princeton University

May 10, 2021

Abstract

Classical models of perceptual decision-making assume that animals use a single, consistent strategy to form decisions, or that decision-making strategies evolve slowly over time. Here we present new analyses suggesting that this common view is incorrect. We analyzed data from two mouse decision-making experiments and found that choice behavior relies on an interplay between multiple interleaved strategies. These strategies, characterized by states in a hidden Markov model, persist for tens to hundreds of trials before switching, and may alternate multiple times within a session. The identified strategies were highly consistent across animals, consisting of a single “engaged” state, in which decisions relied heavily on the sensory stimulus, and several biased or disengaged states in which errors frequently occurred. These results provide a powerful alternate explanation for “lapses” often observed in psychophysical experiments, and suggest that standard measures of performance mask the presence of dramatic changes in strategy across trials.

1 Introduction

Understanding the computations performed in the brain will require a comprehensive characterization of behavior [6, 29, 39]. This realization has fueled a recent surge in methods devoted to the measurement, quantification, and modeling of natural behaviors [2, 12, 37, 59, 67]. Historically, studies of perceptual decision-making behavior have tended to rely on models derived from signal detection theory (SDT) [30, 38] or evidence accumulation [8, 27, 52]. More recently, approaches based on reinforcement learning have also been used to model the effects of context, reward, and trial-history on perceptual decision-making behavior [34, 40, 50, 63]. In all cases, however, these approaches describe decision-making in terms of a single strategy that does not change abruptly across trials or sessions.

*zashwood@princeton.edu

†pillow@princeton.edu

One puzzling aspect of sensory decision-making behavior is the presence of so-called “lapses”, in which an observer makes an error despite the availability of strong sensory evidence. The term itself suggests an error that arises from a momentary lapse in attention or memory, as opposed to an inability to perceive the sensory stimulus. Lapses arise in all species, but are surprisingly frequent in rodent experiments, where lapses can comprise up to 10-20% of all trials [35, 48, 49].

The standard approach for modeling lapses involves augmenting the classic psychometric curve with a “lapse parameter”, which characterizes the probability that the observer simply ignores the stimulus on any given trial [13, 51, 66]. This model can be conceived as a mixture model [20, 43, 46] in which, on every trial, the animal flips a biased coin to determine whether or not to pay attention to the stimulus when making its choice. Previous literature has offered a variety of explanations for lapses, including inattention, motor error, and incomplete knowledge of the task [13, 41, 66], and recent work has argued that they reflect an active process of uncertainty-guided exploration [50]. However, a common thread to these explanations is that lapses arise independently across trials, in a way that does not depend on the time course of other lapses.

Here we show that lapses do not arise independently, but depend heavily on latent states that underlie decision-making behavior. We use model-based analyses to show that mice rely on discrete decision-making strategies that persist for tens to hundreds of trials. One of these states corresponds to an “engaged” strategy, in which the animal’s choices are strongly influenced by the sensory stimulus, while other states correspond to biased or weakly stimulus-dependent strategies. These analyses show that lapses arise primarily during long sequences of trials when the animal is in a biased or disengaged state. Conversely, we find that animals with high apparent lapse rates may nevertheless be capable of high-accuracy performance for extended blocks of trials.

Our modeling framework consists of a hidden Markov Model (HMM) with states corresponding to different decision-making strategies. Within each state, the animal’s strategy is parameterized by a Bernoulli generalized linear model (GLM). The resulting “GLM-HMM” framework [5, 12, 22] includes the classic lapse model as a special case, where the probability of entering the stimulus-independent “lapse state” is the same on every trial. However, the framework allows for a variety of more complex behaviors in which multiple decision-making strategies trade off in a state-dependent manner over longer timescales.

We used the GLM-HMM to analyze choice data from two large cohorts of mice performing different visual detection tasks [35, 48]. We showed that lapses in these experiments were the result of mice switching between relatively engaged and disengaged decision-making strategies *within* a session. While lapse events are typically expected to last for a single trial, these strategies persisted for tens to hundreds of trials. We studied the statistics of data simulated from the best-fitting GLM-HMM, and found that the long choice run-lengths (a run refers to a sequence of consecutive trials in which a mouse makes the same repeated choice) that are observed in the real data only exist in data simulated from a GLM-HMM and not in data simulated from the classic lapse model, indicating that the GLM-HMM better captures the true behavior of real mice. Finally, we validated our interpretation of the GLM-HMM’s latent states when we compared the response times and violation rates (quantities not used to fit our model) associated with engaged and disengaged states. We found that the most extreme response times were typically associated with the disengaged states, which is consistent with previous findings linking accuracy and response times [32, 53, 61]. Taken together, these results shed substantial new light on the factors governing sensory decision-making in rodents, and provide a powerful set of tools for identifying hidden states in behavioral data.

2 Results

2.1 The classic lapse model for sensory decision-making

A common approach for analyzing data from two-choice perceptual decision-making experiments involves the psychometric curve, which describes the probability that the animal chooses one option (e.g., “rightward”) as a function of the stimulus value [13, 51, 66]. The psychometric curve is commonly parameterized as a sigmoidal function that depends on a linear function of the stimulus plus an offset or bias. This sigmoid rises from a minimum value of γ_r to a maximal value of $1 - \gamma_l$, where γ_r and γ_l denote “lapse” parameters, which describe the probability of making a rightward or leftward choice independent of the stimulus value. Thus, the probability of a rightward choice is always at least γ_r , and it cannot be greater than $1 - \gamma_l$. In what follows, we will refer to this as the “classic lapse model of choice behavior”, which can be written:

$$p(y_t = 1 \mid \mathbf{x}_t) = (1 - (\gamma_r + \gamma_l)) \frac{1}{1 + e^{-\mathbf{x}_t \cdot \mathbf{w}}} + \gamma_r, \quad (1)$$

where $y_t \in \{0, 1\}$ represents the choice (left or right) that an animal makes at trial t , $\mathbf{x}_t \in \mathbb{R}^M$ is a vector of covariates, and $\mathbf{w} \in \mathbb{R}^M$ is a vector of weights that describes how much each covariate influences the animal’s choice. Note that \mathbf{x}_t includes both the stimulus and a constant ‘1’ element to capture the bias or offset, but it may also include other covariates that empirically influence choice, such as previous choices, stimuli, and rewards [11, 25, 44].

Although the classic lapse model can be viewed as defining a particular sigmoid-shaped curve relating the stimulus strength to behavior (Fig. 1c), it can equally be viewed as a mixture model [20, 43, 46]. In this interpretation, we regard the animal as having an internal state z_t that takes on one of two different values on each trial, namely “engaged” or “lapse”. If the animal is engaged, it makes its choice according to the classic sigmoid curve (which saturates at 0 and 1). If lapsing, it ignores the stimulus and makes its choice based only on the relative probabilities of a left and right lapse. Mathematically, this can be written:

$$p(y_t = 1 \mid \mathbf{x}_t) = \begin{cases} \frac{1}{1 + e^{-\mathbf{w} \cdot \mathbf{x}_t}}, & z_t = \text{“engaged”} \\ \frac{\gamma_r}{\gamma_r + \gamma_l}, & z_t = \text{“lapse”}, \end{cases} \quad (2)$$

where $p(z_t = \text{“lapse”}) = (\gamma_r + \gamma_l)$ and $p(z_t = \text{“engaged”}) = 1 - (\gamma_r + \gamma_l)$. In this interpretation, the animal flips a biased coin, with fixed probability $(\gamma_r + \gamma_l)$, on each trial and then adopts one of two strategies based on the outcome—a strategy that depends on the stimulus vs. one that ignores it. Note that the animal can make a correct choice in the lapse state through pure luck; we use “lapse” here simply to indicate that the animal is not relying on the stimulus when making its decision.

Viewing the classic lapse model as a mixture model highlights some of its limitations. First, it assumes that animals switch between only two decision-making strategies. Second, it assumes that lapses occur independently in time, according to an independent Bernoulli random variable on each trial. Finally, the model assumes that choices in the “lapse” state are fully independent of the stimulus, neglecting the possibility that they are still weakly stimulus dependent [13], or are influenced by other covariates such as reward or choice history [11, 25]. These limitations motivate us to consider a more general family of models, which includes the classic lapse model as a special case.

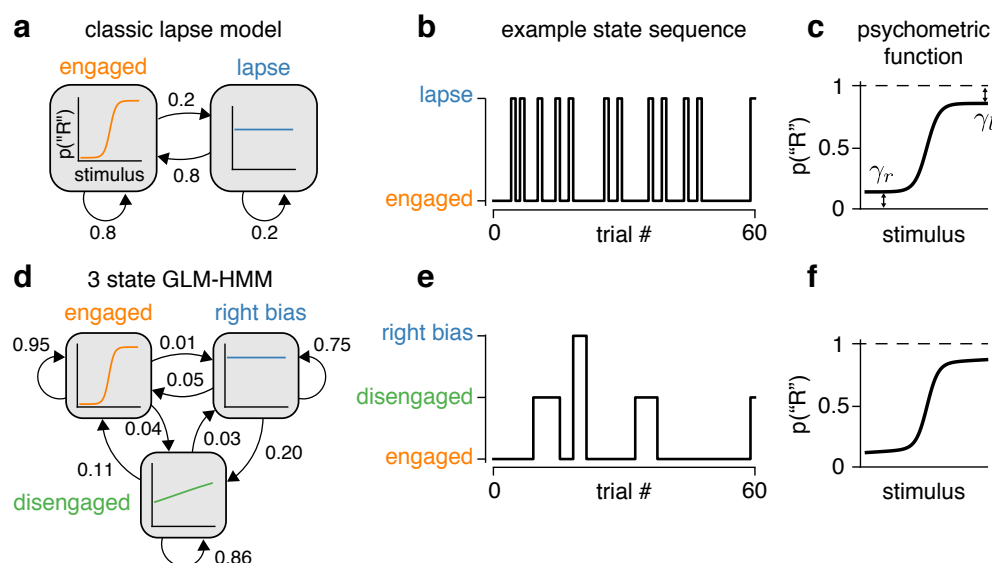


Figure 1: The GLM-HMM generalizes the classic lapse model. (a) The classic lapse model formulated as a 2-state GLM-HMM. Each box represents a generalized linear model (GLM) describing how the probability of a binary choice depends on the stimulus in the corresponding state: “engaged” (left) or “lapse” (right). Arrows between boxes indicate the transition probabilities between states. Note that the probability of switching to the engaged state at the next trial is always 0.8 (it is independent of state at the current trial) and similarly there is always a 0.2 probability of entering the lapse state on each trial. (b) An example (simulated) sequence for the animal’s internal state when the transitions between states are governed by the probabilities in (a). Notice that the lapse state tends to last for only a single trial at a time. (c) Psychometric function arising from the model shown in (a), depicting the probability of a rightward choice as a function of the stimulus. The parameters γ_r and γ_l denote the probability of a rightward and leftward lapse, respectively. As specified by the transition probabilities in (a), the total lapse probability for this model is $\gamma_r + \gamma_l = 0.2$. (d) Example 3-state GLM-HMM, with three different GLMs corresponding to different decision-making strategies (labeled “engaged”, “disengaged” and “right biased”). The high self-transition probabilities of 0.95, 0.86 and 0.75 ensure that these states typically persist for many trials in a row. (e) An example sequence for the animal’s internal state sampled from the GLM-HMM shown in (d). (f) The psychometric curve arising from the model shown in (d), which corresponds to a weighted average of the psychometric functions associated with each state. Note that although the decision-making models shown in (a) and (d) are vastly different, the resulting psychometric curves are nearly identical, meaning that the psychometric curve alone cannot be used to distinguish them.

2.2 A model for decision-making with multiple strategies

Recognizing the limitations of the classic lapse model, we propose to analyze perceptual decision-making behavior using a framework based on Hidden Markov models (HMMs) with Bernoulli generalized linear model (GLM) observations [12, 22]. The resulting “GLM-HMM” framework, also known as an input-output HMM [5], allows for an arbitrary number of states, which can persist for an extended number of trials and exhibit different dependencies on the stimulus and other covariates.

A GLM-HMM has two basic pieces: an HMM governing the distribution over latent states, and a set of state-specific GLMs, specifying the decision-making strategy employed in each state (see Fig. 1). For a GLM-HMM with K latent states, the HMM has a $K \times K$ transition matrix A specifying the probability of transitioning from any state to any other,

$$p(z_t = k \mid z_{t-1} = j) = A_{jk}, \quad (3)$$

where z_{t-1} and z_t indicate the latent state at trials $t - 1$ and t , respectively. The “Markov” property of

the HMM is that the state on any trial depends only on the state from the previous trial, and the “hidden” property refers to the fact that states are latent or hidden from external observers. For completeness, the HMM also has a distribution over initial states, given by a K -element vector π whose elements sum to 1, giving $p(z_1 = k) = \pi_k$.

To describe the state-dependent mapping from inputs to decisions, the GLM-HMM contains K independent Bernoulli GLMs, each defined by a weight vector specifying how inputs are integrated in that particular state. The probability of a rightward choice ($y_t = 1$) given the input vector \mathbf{x}_t and the latent state z_t is given by

$$p(y_t = 1 \mid \mathbf{x}_t, z_t = k) = \frac{1}{1 + e^{-\mathbf{x}_t \cdot \mathbf{w}_k}}, \quad (4)$$

where $\mathbf{w}_k \in \mathbb{R}^M$ denotes the GLM weights for latent state $k \in \{1, \dots, K\}$. The full set of parameters for a GLM-HMM can therefore be denoted $\Theta \equiv \{\pi, A, \{\mathbf{w}_k\}_{k=1}^K\}$, consisting of an initial probability vector π , a transition matrix A , and a set of state-specific GLM weights $\{\mathbf{w}_k\}_{k=1}^K$.

It is worth noting that the classic lapse model described in Eq. 1 and Eq. 2 corresponds to a restricted 2-state GLM-HMM. If we consider state 1 to be “engaged” and state 2 to be the “lapse” state, then the state-1 GLM has weights $\mathbf{w}_1 = \mathbf{w}$, and the state-2 GLM has all weights set to 0 except the bias weight, which is equal to $-\log(\gamma_l/\gamma_r)$. The transition matrix has identical rows, with probability $1 - (\gamma_r + \gamma_l)$ of going into state 1 and probability $(\gamma_r + \gamma_l)$ of going into state 2 at the next trial, regardless of the current state. This ensures that the probability of a lapse on any given trial is stimulus-independent and does not depend on the previous trial’s state. Fig. 1a-c shows an illustration of the classic lapse model formulated as a 2-state GLM-HMM.

However, there is no general reason to limit our analyses to this restricted form of the GLM-HMM. By allowing the model to have more than two states, multiple states with non-zero stimulus weights, and transition probabilities that depend on the current state, we obtain a model family with a far richer set of dynamic decision-making behaviors. Fig. 1d shows an example GLM-HMM with three latent states, all of which have high probability of persisting for multiple trials. Intriguingly, the psychometric curve arising from this model (Fig. 1f) is indistinguishable from that of classic lapse model. Thus, the psychometric curve, which is simply the average probability of a choice given the stimulus, cannot provide insight into the dynamics of decision-making across trials.

2.3 Mice switch between multiple strategies during visual decision-making

To examine whether animals employ multiple strategies during decision-making, we fit the GLM-HMM to behavioral data from two binary perceptual decision-making tasks (see Methods 4.1). First, we fit the GLM-HMM to choice data from 37 mice performing a visual detection decision-making task developed in [10] and adopted by the International Brain Laboratory (IBL) [35]. During the task, a sinusoidal grating with contrast between 0 and 100% appeared either on the left or right side of the screen (Fig. 2a). The mouse had to indicate this side by turning a wheel. If the mouse turned the wheel in the correct direction, it received a water reward; if incorrect, it received a noise burst and an additional 1-second timeout. We analyzed choice data from animals with at least 3000 trials of data (across multiple sessions) after they had successfully learned the task (see Methods 4.4).

We modeled the animals’ decision-making strategies using a GLM-HMM with four inputs: (1) the (signed) stimulus contrast, where positive values indicate a right-side grating and negative values indicate a left-side grating; (2) a constant offset or bias; (3) the animal’s choice on the previous trial; and (4) the stimulus side on the previous trial. A large weight on the animal’s previous choice gives rise to a strategy known as “persistence” in which the animal makes the same choice many times in a

row, regardless of whether it receives a reward. A large weight on the previous stimulus side, which we refer to as the “win-stay, lose-switch” regressor, gives rise to the well-known strategy in which the animal repeats a choice if it was rewarded, and switches choices if it was not. Note that for the IBL task in question, bias and trial-history dependencies were sub-optimal, meaning that the maximal-reward strategy was to have a large weight on the stimulus and zero weights on the other three inputs.

To determine the number of different strategies underlying decision-making behavior, we fit GLM-HMMs with varying numbers of latent states. Note that the 1-state model is simply a standard Bernoulli GLM, while the classic lapse model (Eq. 1, Eq. 2) is a constrained version of the 2-state model. We found that a 3-state GLM-HMM substantially outperformed models with fewer states, including the classic lapse model. The states of the fitted model were readily interpretable, and tended to persist for many trials in a row.

Figures 2 and 3 show results for an example mouse. For this animal, the multi-state GLM-HMM outperformed both the standard (1-state) GLM and the classic lapse model, both in test log-likelihood and percent correct, with the improvement approximately leveling off at 3 latent states (Fig. 2b-c). Note that the test set for this mouse contained 900 trials. Log-likelihood increases of 0.13 bits/trial and 0.09 bits/trial for the 3-state model over the 1-state model and classic lapse models, respectively, meant that the data were $(2^{0.13})^{900} \approx 1.7 \times 10^{35}$ and $(2^{0.09})^{900} \approx 2.4 \times 10^{24}$ times more likely under the 3-state model.

The transition matrix for the fitted 3-state model describes the transition probabilities between three different states, each of which corresponds to a different decision-making strategy (Fig. 2d). Large entries along the diagonal of this matrix, ranging between 0.94 and 0.98, indicate a high probability of remaining in the same state for multiple trials. The other set of inferred parameters were the GLM weights, which define how the animal makes decisions in each state (Fig. 2e). One of these GLMs (“state 1”) had a large weight on the stimulus and negligible weights on other inputs, giving rise to high-accuracy performance on the task (Fig. 2f). The other two GLMs (“state 2” and “state 3”), by comparison, had small smaller weights on the stimulus, and relatively large bias weights.

We can visualize the decision-making strategies associated with these states by plotting the corresponding psychometric curves (Fig. 2g), which show the probability of a rightward choice as a function of the stimulus, conditioned on both previous choice and reward. The steep curve observed in state 1, which exhibited near-perfect performance on high-contrast stimuli and little dependence on previous choice or reward, led us to call it the ‘engaged’ state. By comparison, the psychometric curves for states 2 and 3 reflected large leftward and rightward biases, respectively. They also had relatively large dependence on previous choice and reward, as indicated by the gap between solid and dashed lines. While this mouse had an overall accuracy of 80%, it achieved 90% accuracy in engaged state, compared to only 60% and 58% accuracy in the two biased states (Fig. 2f).

To gain insight into the temporal fluctuations in decision-making behavior, we used the fitted 3-state model to compute the posterior probability over the mouse’s latent state across trials (Fig. 3). These probabilities, formally given by $p(z_t = k \mid \{\mathbf{x}_t, y_t\}_{t=1}^T)$, reflect the experimenter’s ability to infer the animal’s state on any given trial from the entire sequence of choices and inputs during a session. Fig. 3a shows state probabilities for three example sessions (Note that we examined only the first 90 trials of each session, when the stimulus statistics were stationary; see Methods Sec. 4.4 for details). Contrary to the predictions of the classic lapse model, strategies persisted for many trials at a time. Remarkably, the most probable state often had probability close to 1, indicating that we can be highly confident about the mouse’s internal state given the observed data.

If we assign each trial to its most probable state, this mouse spent approximately 69% of all trials (out of

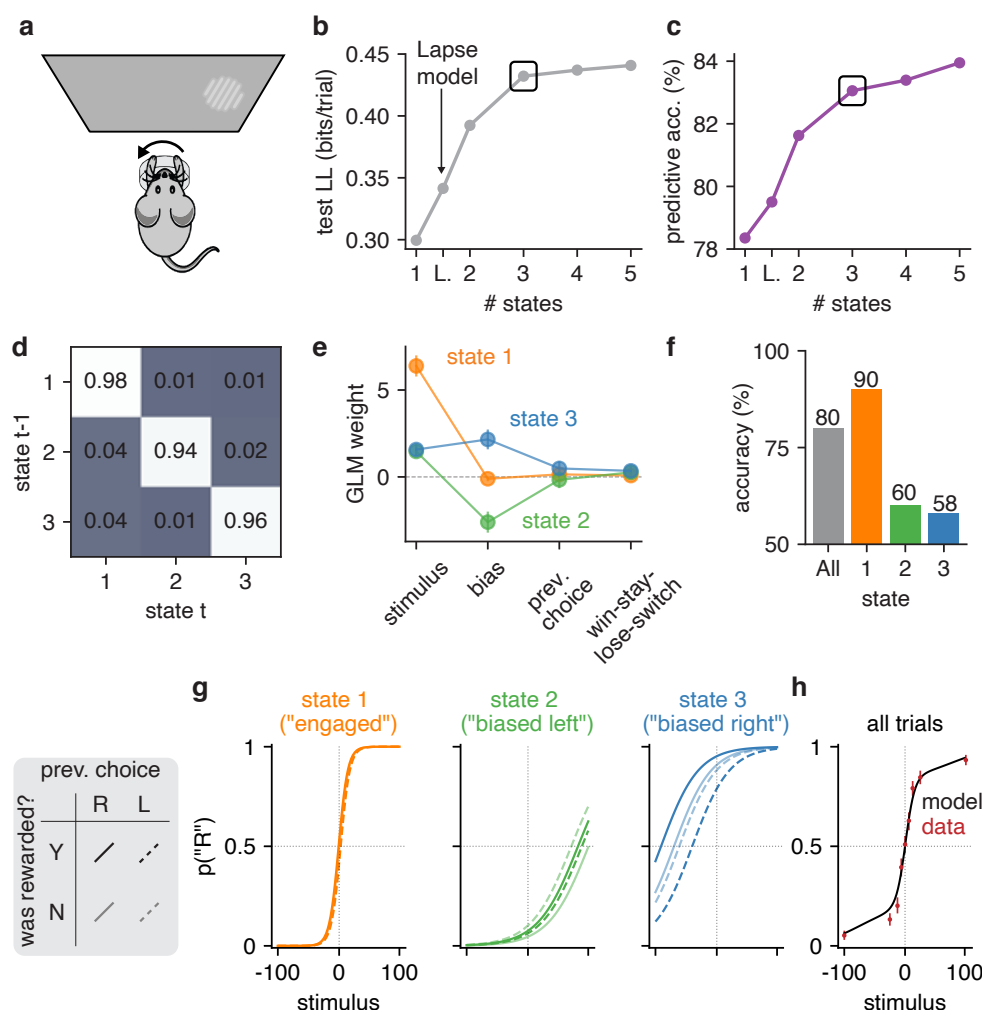


Figure 2: GLM-HMM analysis of choice behavior of an example IBL mouse. (a) Schematic for visual decision-making task, which involved turning a wheel to indicate whether a Gabor patch stimulus appeared on the left or right side of the screen [10, 35]. (b) Model comparison between GLM-HMMs with different numbers of latent states, as well as classic lapse model (labeled 'L') using 5-fold cross-validation. Normalized log-likelihood in units of bits/trial was computed relative to a 'null' Bernoulli coin flip model (see Methods). A black rectangle highlights the log-likelihood for the 3-state model, which we used for all subsequent analyses. (c) Test set predictive accuracy for each model, indicating the percent of held-out trials where the model successfully predicted the mouse's true choice. (d) Inferred transition matrix for best-fitting 3 state GLM-HMM for this mouse. Large entries along the diagonal indicate a high probability of remaining in the same state. (e) Inferred GLM weights for the 3-state model. State 1 weights have a large weight on the stimulus, indicating an "engaged" or high-accuracy state. In states 2 and 3, the stimulus weight is small, and the bias weights give rise to large leftward (state 2) and rightward (state 3) biases. (f) Overall accuracy of this mouse (grey), and accuracy for each of the three states. (g) Psychometric curve for each state, conditioned on previous reward and previous choice. (h) The standard psychometric curve for all choice data from this mouse, which can be seen to arise from the mixture of the three per-state curves shown in (g). Empirical data are shown on top of the fit curve, as are 95% confidence intervals.

5040 total trials over 56 sessions) in the engaged state, compared to 15% and 16% of trials in the biased left and rightward states (Fig. 3c). Moreover, the mouse changed state at least once within a session in roughly 71% of all 90-trial sessions, and changed multiple times in 59% of sessions (Fig. 3d). This rules out the possibility that the states merely reflect the use of different strategies on different days. Rather,

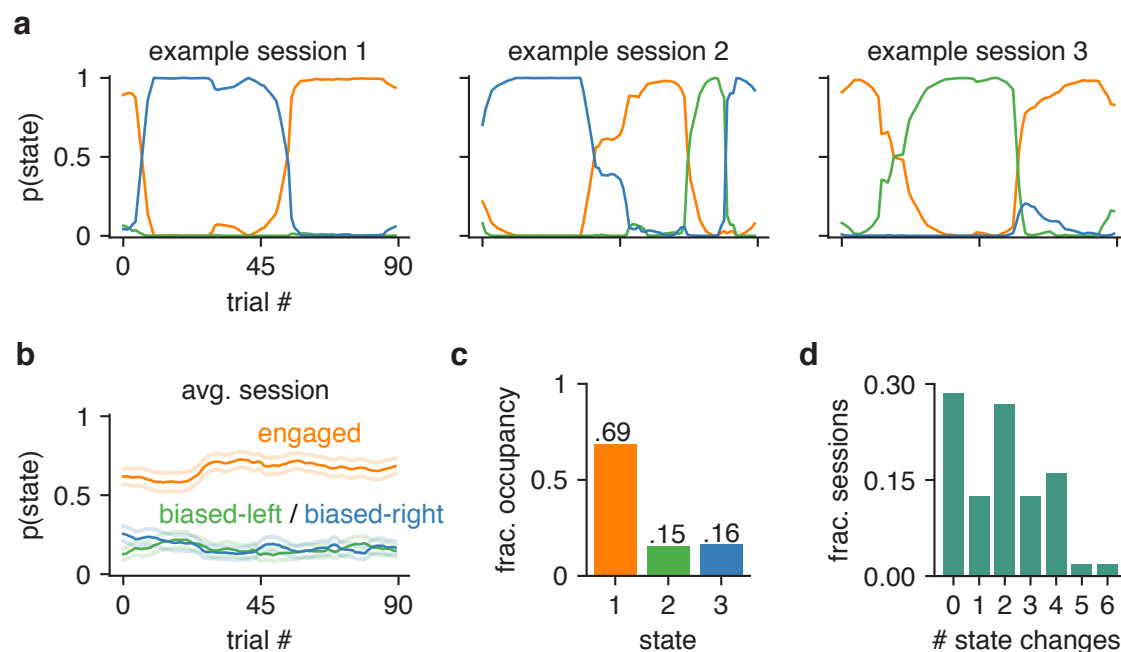


Figure 3: **Inferred state dynamics for example IBL mouse.** (a) Posterior state probabilities for three example sessions, revealing high levels of certainty about the mouse's internal state, and showing that states typically persisted for many trials in a row. (b) Average trajectories of state probabilities within a session, computed over 56 sessions. Error bars indicate ± 1 standard error of the mean. (c) Fractional occupancies for the three states across all trials. For this analysis, we assigned each trial to its most likely state and then counted the fraction of trials assigned to each state. Overall, the mouse spent 69% of all trials in the engaged state, and 31% of trials in one of the two biased states. (d) Histogram showing the number of inferred state changes per session, for all 56 sessions of data for this mouse. Only 29% of sessions involved the mouse persisting in a single state for the entire 90-trial session.

the mouse tended to remain in an engaged, high-performance state for tens of trials at a time, with lapses arising predominantly during interludes when it adopted a left-biased or right-biased strategy for multiple trials in a row. The multi-state GLM-HMM thus provides a very different portrait of mouse decision-making behavior than the basic GLM or lapse model.

2.4 State-based strategies are consistent across mice

To assess the generality of these findings, we fit the GLM-HMM to the choice data from 37 mice in the IBL dataset [35] (Fig. 4). We found that the results shown for the example animal considered above were broadly consistent across animals. Specifically, we found that the 3-state GLM-HMM strongly outperformed the basic GLM and classic lapse model in cross-validation for all 37 mice (Fig. 4a and Fig. S5). On average, it predicted mouse choices with 4.2% higher accuracy than the basic GLM (which had an average prediction accuracy of 78%), and 2.8% higher accuracy than the classic lapse model (Fig. 4b). Furthermore, for one animal the improvement in prediction accuracy for the 3 state GLM-HMM was as high as 12% relative to the basic GLM, and 7% relative to the classic lapse model.

Although performance continued to improve slightly with four and even five latent states, we will focus our analyses on the 3-state model for reasons of simplicity and interpretability. Supplementary figure S8 provides a full analysis of 4-state model, showing that it tended to divide the engaged state from the

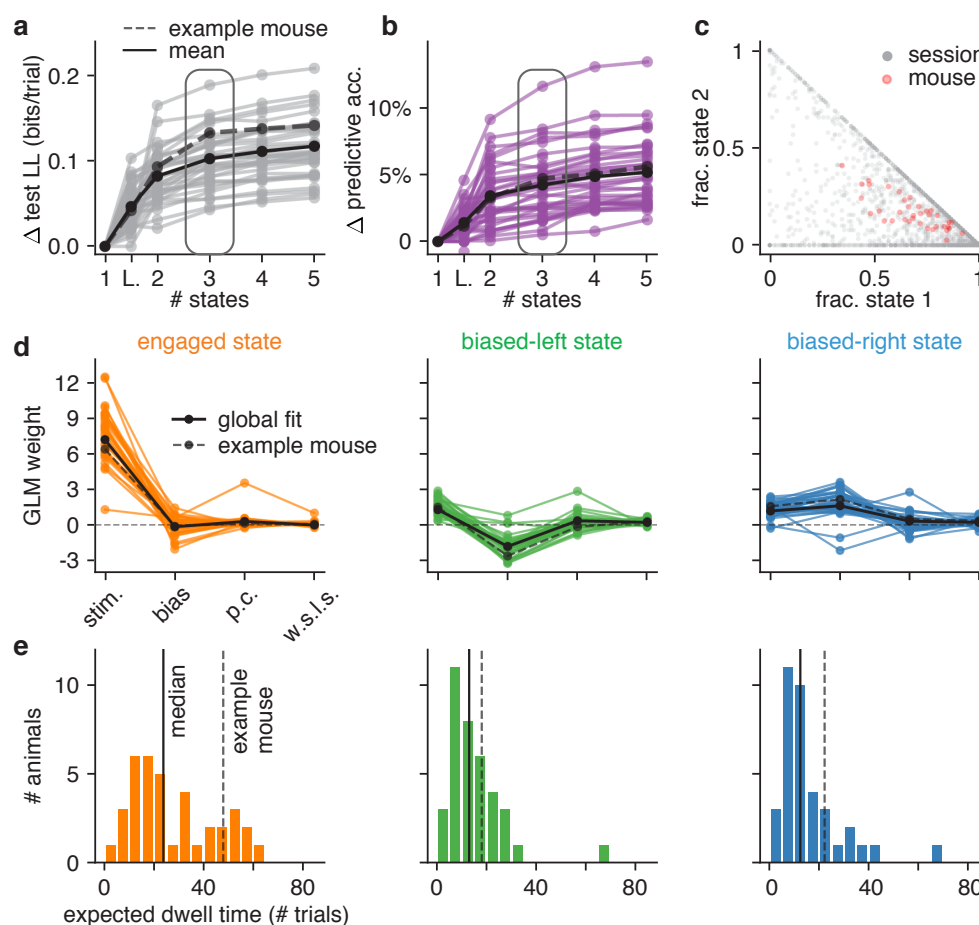


Figure 4: Model fits to full population of 37 IBL mice. (a) Change in test log-likelihood as a function of number of states relative to a (1-state) GLM, for each mouse in the population. The classic lapse model, a restricted form of the 2-state model, is labeled ‘L’. Each trace represents a single mouse. Solid black indicates the mean across animals, and the dashed line indicates the example mouse from Figs. 2 and 3. The rounded rectangle highlights performance of the 3-state model, which we selected for further analyses. (b) Change in predictive accuracy relative to a basic GLM for each mouse, indicating the percentage improvement in predicting choice. (c) Grey dots correspond to 2017 individual sessions across all mice, indicating the fraction of trials spent in states 1 (engaged) and 2 (biased left). Points at the vertices (1,0), (0,1), or (0,0) indicate sessions with no state changes, while points along the sides of the triangle indicate sessions that involved only 2 of the 3 states. Red dots correspond the same fractional occupancies for each of the 37 mice, revealing that the engaged state predominated, but that all mice spent time in all 3 states. (d) Inferred GLM weights for each mouse, for each of the three states in the 3-state model. The solid black curve represents a global fit using pooled data from all mice (see Algorithm 1); the dashed line is the example mouse from Fig. 2 and Fig. 3. (e) Histogram of expected dwell times across animals in each of the three states, calculated from the inferred transition matrix for each mouse.

3-state model into two “sub-states” that differed slightly in accuracy.

Fits of the 3-state GLM-HMM exhibited remarkable consistency across mice, with a majority exhibiting states that could be classified as “engaged”, “biased-left”, and “biased-right” (Fig. 4d). (See Methods section 4.1.5 for details about the alignment of states across mice). While we plot inferred transition matrices for all 37 mice in supplementary figure S6, here we used the diagonal elements of each matrix to compute an expected dwell time for each animal in each state (Fig. 4e). This revealed a median dwell time, across animals, of 24 trials for the engaged state, versus 13 and 12 trials for the biased-

left and biased-right states, respectively. Thus, mice tended to remain engaged for longer periods than they remained in either biased state, though the average duration of biased states still departed dramatically from the assumptions of the classic lapse model. For context, for the classic lapse model and a lapse rate of 20%, the expected dwell time in the lapse state is just $\frac{1}{1-0.2} = 1.25$ trials. We analyzed the distribution of state dwell times inferred from data and found they were well approximated by a geometric distribution, matching the theoretical distribution of data sampled from a Hidden Markov Model (Fig. S9 and Fig. S10).

Finally, we examined the fraction of trials per session that mice spent in each of the three states (Fig. 4c). To do so, we used the fitted model parameters to compute the posterior probabilities over state, and assigned each trial to its most likely state. The resulting “fractional occupancies” revealed that the median mouse spent 69% of its time in the engaged state, with the best mice exceeding 90% engagement. Moreover, the majority of sessions (83% of 2017 sessions across all animals) involved a switch between two or more states; in only 17% of sessions did a mouse remain in the same state for an entire 90-trial session.

2.5 The GLM-HMM captures the statistics of real data

In order to understand why the GLM-HMM outperformed the classic lapse model, we compared the statistics of real data with data simulated from the two models. The GLM-HMM incorporates long temporal dependencies in strategy, which may induce particular temporal patterns in the observed choice data. Hence, a natural statistic to examine is choice run-length, where a run refers to a sequence of consecutive trials in which the mouse makes the same repeated choice (Fig. 5a). Figure 5b shows the distribution of choice run-lengths observed in the 181,530 trials of real mouse data (red), as well as the choice run-lengths that would arise if the mice performed the task perfectly (grey). Under perfect performance, only 2% of trials would occur in runs of length 10 or greater, whereas 19% of trials occur in such runs in the real data. A good model of mouse decision-making behavior should be able to capture the heavy tail of this run-length distribution.

We examined the distribution of simulated run-lengths generated by three different fitted models: (1) the classic lapse model with only a bias and stimulus weight; (2) the classic lapse model with the usual 4 inputs (stimulus, bias, previous choice and win-stay-lose-switch); (3) the full 3-state GLM-HMM. The stimulus-only lapse model had the largest discrepancy with real data, with a pronounced under-prediction of long runs (Fig. 5c). The full lapse model performed better, confirming that the trial-history regressors were instrumental in producing longer run-lengths. However, the full 3-state GLM-HMM performed far better than either lapse model, generating a run-length distribution that was closest to the real data.

We also examined a related statistic, given by the number of runs lasting more than 5 trials (Fig. 5d). The real data contained 6,111 runs longer than 5 trials. This was consistent with datasets simulated from the fitted GLM-HMM, but far greater than the number produced by either of the two lapse models. Thus, compared to the classic lapse model, the GLM-HMM was far better able to account for the temporal distribution of mouse choices, in particular the tendency to produce long sequences of repeated choices.

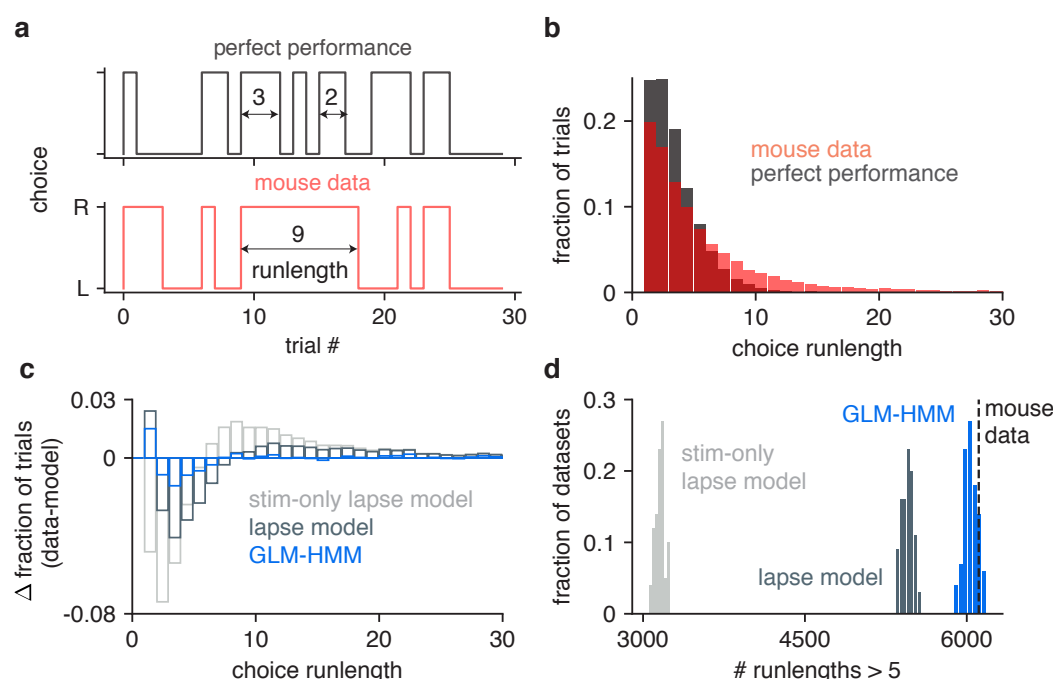


Figure 5: Simulated data from GLM-HMM captures statistics of real choice data. (a) Definition of choice run-length. Shown are the choices that an IBL mouse made over the course of 30 trials (red, bottom), as well as the choices it should have made during that same time course if the mouse performed the task perfectly (grey, top). Choice run-length is defined as the number of trials during which a mouse repeated the same decision (example choice run-lengths of 2, 3 and 9 trials are highlighted). (b) Red: fraction of trials in choice run-lengths of between 1 and 30 trials when calculated from all trials for all mice. Grey: distribution of choice run-lengths that would have been obtained if IBL mice performed the task perfectly. (c) Difference in choice run-length distribution for simulated data (from three different models) compared to the red distribution shown in (b). Models used to simulate data were a lapse model with only stimulus intensity and bias regressors, a lapse model that also included history regressors (previous choice and win-stay-lose-switch), and a 3 state GLM-HMM (also with history regressors). We simulated 100 example choice sequences from each model and calculated the mean histogram of choice run-lengths across the 100 simulations. This was then subtracted from the red histogram of (b). (d) Number of choice run-lengths with more than 5 trials for each model simulation used in (c). In the 181,530 trials of real choice data, there were 6111 run-lengths lasting more than 5 trials (as shown with the dashed line). When we simulated choice data according to each of the models shown in (c), we found that only the GLM-HMM could generate simulations with as many run-lengths lasting more than 5 trials as in the real data (15/100 simulations had 6111 or more run-lengths lasting more than 5 trials for the GLM-HMM compared to 0/100 for both of the lapse models).

2.6 Data provide evidence for discrete, not continuous, states

The GLM-HMM describes perceptual decision-making in terms of discrete states that persist for many trials in a row before switching. However, the model's dramatic improvement over classic models does not guarantee that the states underlying decision-making are best described as discrete. One could imagine, for example, that a continuous state governing the animal's degree of engagement drifts gradually over time, and that the GLM-HMM simply divides these continuous changes into discrete clusters. To address this possibility, we fit the data with PsyTrack, a psychophysical model with continuous latent states [55, 56]. The PsyTrack model describes sensory decision-making using an identical Bernoulli GLM, but with dynamic weights that drift according to a Gaussian random walk (see Methods sec. 4.3).

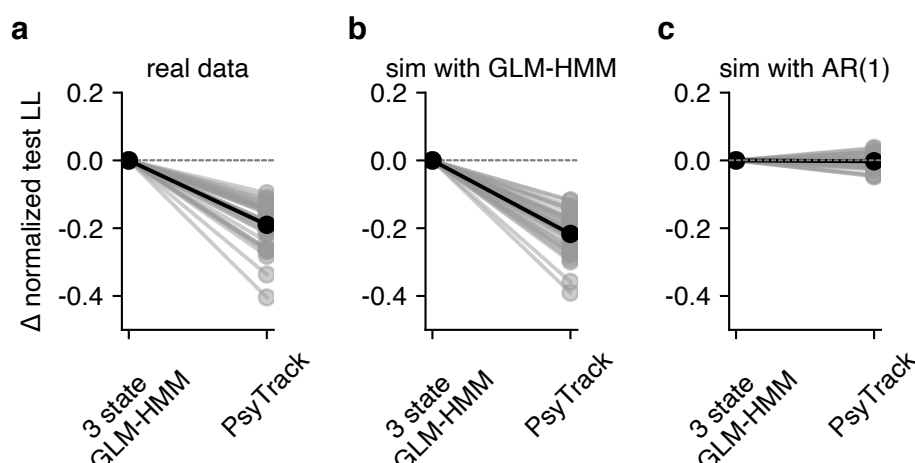


Figure 6: Mice have discrete—not continuous—decision-making states. (a) Cross-validation performance of the 3 state GLM-HMM compared to PsyTrack [55, 56] for all 37 mice studied (each individual line is a separate mouse; black is the mean across animals). Test loglikelihood is considerably higher for all mice for the GLM-HMM than for PsyTrack, which assumes that the mouse uses a set of smoothly evolving GLM weights to make its decision at each trial. (b) As a sanity check, we simulated datasets from a 3 state GLM-HMM with the parameters for each simulation chosen as the best fitting parameters for a single mouse. We then fit the simulated data both with PsyTrack and with the 3 state GLM-HMM in order to check that the 3 state GLM-HMM best described the data. (c) We did the opposite and fit PsyTrack to the animals' data and then generated data according to an AR(1) model with parameters specified using the PsyTrack fits (see section 4.3 for full details). By performing cross-validation on the simulated data, we confirmed that we could use model comparison to distinguish between discrete and continuous decision-making behavior in choice data.

Although PsyTrack has previously been used to characterize slow changes that arise over the course of learning, here we used it to assess whether the decision-making in well-trained mice is better described using weights that drift continuously or that switch between three discrete states. For all 37 mice in our dataset, the 3-state GLM-HMM achieved substantially higher test log-likelihood than the PsyTrack model (Fig. 6a). Model selection also correctly identified simulated data from the GLM-HMM, whereas datasets simulated from a matched first-order autoregressive model had roughly equal log-likelihood under the two models (Fig. 6b-c).

We also used the GLM-HMM to look for evidence of discrete state changes in the raw behavioral time series data itself. Fig. S11 shows the resulting plots, revealing that an identified state change from engaged to one of the two biased states corresponded to a drop of approximately 20% in accuracy, with accuracy approximately constant during the 5 trials before and after the switch. This suggests that—at least for the state changes inferred using the fitted model—that performance did not appear to descend gradually over the 10 trials spanning the state change, but rather appeared consistent with a discrete drop in accuracy. We performed a similar analysis for datasets simulated from a GLM-HMM and from a continuous autoregressive model (see Section 4.3). Both models exhibited step-like changes in accuracy after an inferred state change. However, only the GLM-HMM produced the large jump in accuracy found in real data.

2.7 Mice switch between multiple strategies in a second task

To ensure that our findings were not specific to the IBL task or training protocol, we examined a second mouse dataset with a different sensory decision-making task. Odoemene et al. [48] trained mice to report whether the flash rate of an LED was above or below 12Hz by making a right or left nose poke (Fig. 7a). Once again, we found that the multi-state GLM-HMM provided a far more accurate description of mouse decision-making than a basic GLM or the classic lapse model (Fig. 7b). Although the performance of the 3-state and 4-state models was similar, we focused on the 4-state model because—in addition to having slightly higher test log-likelihood for a majority of animals—the 4-state model balanced simplicity and interpretability, with each state in the 4-state model corresponding to a distinct behavioral strategy. (See supplementary figures S17 and S18 for a comparison to 3-state and 5-state fits). The 4-state model exhibited an average improvement of 0.025 bits/trial over the classic lapse model, making the test dataset approximately 1×10^{18} times more likely under the GLM-HMM than to the lapse model.

Fig. 7d shows the inferred GLM weights associated with each state in the 4-state GLM-HMM, while Fig. 7e shows the associated psychometric curves, conditioned on previous choice and previous reward. Based on these curves, we labeled the four states as (from left to right): “engaged”, “biased left”, “biased right” and “win-stay”. The combination of stimulus and choice history weights for this fourth state gave rise to a large separation between psychometric curves conditioned on previous reward; the resulting strategy could be described as “win-stay” because the animal tended to repeat a choice if it was rewarded. (However, it did *not* tend to switch if a choice was unrewarded). Accuracy was highest in the engaged state (92%) and lowest in the biased-left (67%) and biased right states (77%), and took an intermediate value in the win-stay state (83%).

Similar to the IBL dataset, the identified states tended to persist for many trials in a row before switching. As before, we used the diagonal entries of the inferred transition matrices to compute the expected dwell time for each animal for each state, producing median expected dwell times of 63, 34, 49 and 13 trials, respectively (Fig. 7f). These dwell times were nevertheless much shorter than the length of a session, which lasted 650 trials on average in this experiment. This indicates that the mice in the Odoemene et al. [48] task, like the IBL mice, typically switched strategies multiple times per session.

Finally, we used the fitted model to examine the temporal evolution of latent states within a session. Figure 7g shows the average posterior state probabilities over the first 200 trials in a session for two example mice and the average over all mice (Fig. S19 shows average posterior state probabilities for each individual mouse in the cohort separately). These trajectories reveal that mice typically began a session in one of the two biased states, and had a low probability of entering the engaged state within the first 50 trials: mice used these initial trials of a session to “warm-up” and gradually improve their performance [34]. This represents a substantial departure from the IBL mice, the majority of which had a high probability of being engaged from the very first trial, and had relatively flat average trajectories over the first 90 trials of a session (Fig. 3 and Fig. S7). We also examined whether the effects of fatigue or satiety could be observed in the average state probabilities at the end of sessions, but did not find consistent patterns across animals (Fig. S20).

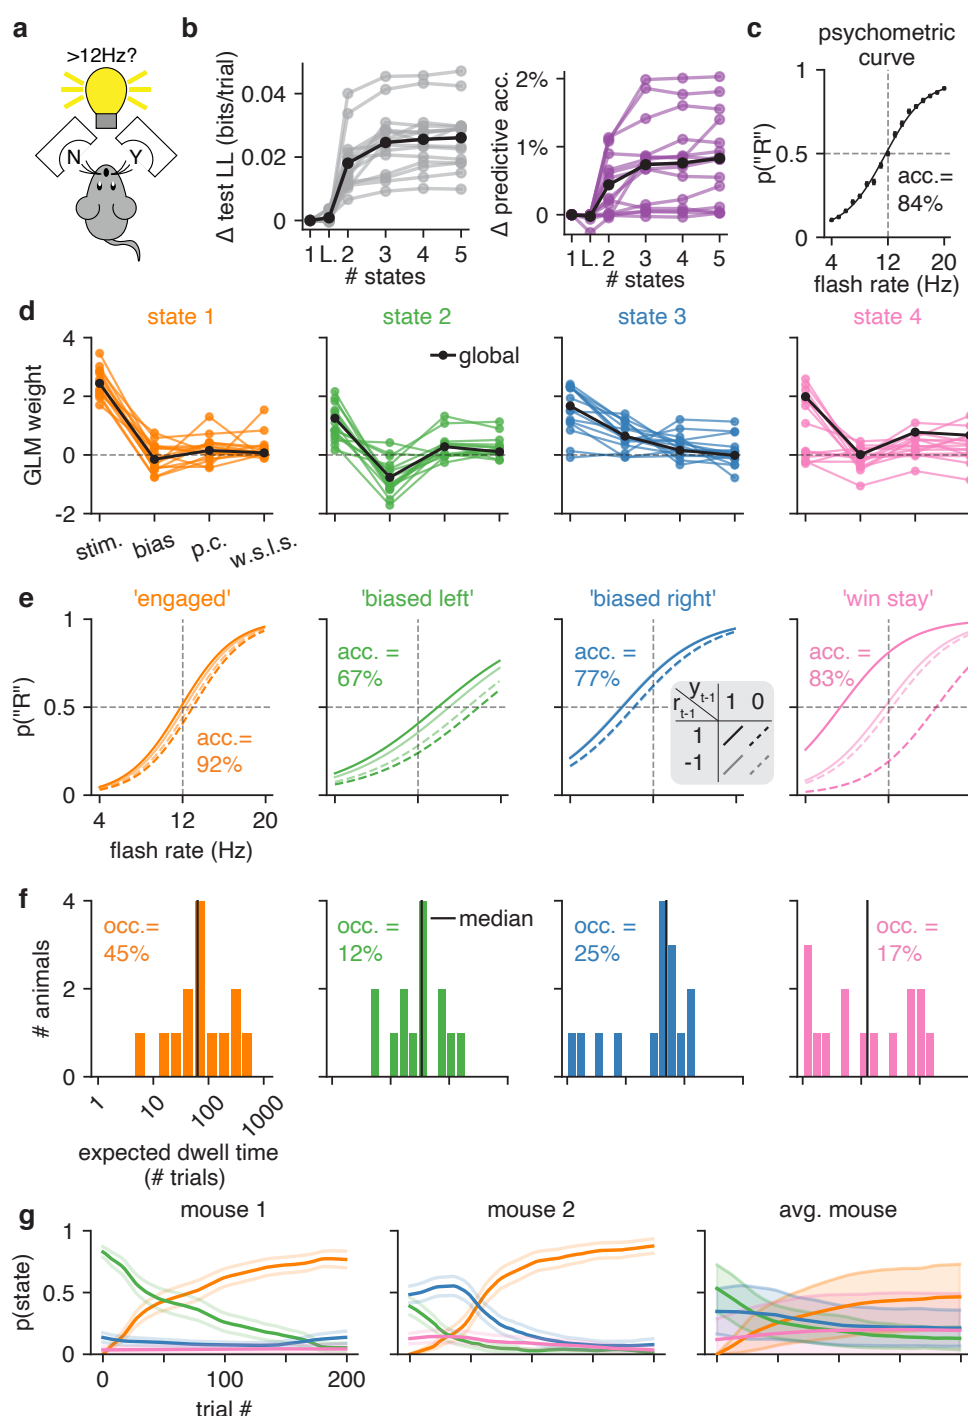


Figure 7: **GLM-HMM application to second mouse dataset**

(a) Mice in the Odoemene et al. [48] task had to indicate if the flash rate of light pulses from an LED panel was above or below 12Hz by nose poking to the right or left.

(b) Left: test set log-likelihood for 15 mice in this dataset. Black is the mean across animals. Right: change in test set predictive accuracy relative to a GLM for each mouse in the dataset. Black is the mean across all animals.

(c) The psychometric curve when all data from all mice are concatenated together.

(d) The retrieved weights for a 4 state GLM-HMM for each mouse. The covariates influencing choice in each of the four states are the stimulus frequency, the mouse's bias, its choice on the previous trial ('p.c') and the product of reward and choice on the previous trial ('w.s.l.s.'). Black indicates the global fit across all mice.

Figure 7 (*previous page*):

(e) The probability that the mouse went rightward in each of the four states as a function of the stimulus, previous choice and reward on the previous trial (each of the four lines corresponds to a different setting of previous choice and previous reward; the four lines are calculated using the global fit weights shown in d). We also report animals' task accuracy (labeled 'acc.') when in each of the four states.

(f) The expected dwell times for each mouse in each state, as obtained from the inferred transition matrices (for the full transition matrices for all mice, see Fig. S16). Black indicates the median dwell time across animals. We report the fractional occupancies (labeled 'occ.') of the four states across all mice.

(g) Left and middle: average (across 20 sessions) posterior state probabilities for the first 200 trials within a session for two example mice; error bars are standard errors. Right: posterior state probabilities for first 200 trials when averaged across all mice. Error bars indicate standard deviation across mice.

2.8 External correlates of engaged and disengaged states

One powerful feature of the GLM-HMM is the fact that it can be used to identify internal states from binary decision-making data alone. However, it is natural to ask whether these states manifest themselves in other observable aspects of mouse behavior. In other words, does a mouse behave differently when it is in the engaged state than in a biased state, above and beyond its increased probability of making a correct choice? To address this question, we examined response times and violation rates, two observable features that we have did not incorporate into the models.

Previous literature has revealed that humans and monkeys commit more errors on long duration trials than short duration trials [32, 53, 61]. We, thus, looked at the distributions of response times for the engaged state, compared to for the disengaged states (the biased left and biased right states), for mice performing the IBL task. Within the IBL task, response time is the time from when the stimulus appears on the screen to when the animal receives feedback on the outcome of its decision (when it receives a reward, or hears an auditory cue to indicate an error trial). The median response time across trials and across IBL mice was just 0.34s, but it was not uncommon for trials to last much longer (up to tens of seconds). We show response time Q-Q plots for each of the 37 IBL mice in Fig. 8a. The engaged and disengaged response time distributions were statistically different for all 37 mice (Komogorov-Smirnov tests reject the null hypothesis with $p < 0.05$ for all 37 mice). Examining the Q-Q plots, it is clear that the most extreme response times for each animal were typically associated with the lower accuracy disengaged states.

Figure 8b shows the difference in the 90th quantile response times for the disengaged and engaged states for each mouse. For the majority of mice, 90th quantile response times were longer for the disengaged states compared to for the engaged state. The median difference in the 90th quantile response time, across all mice, was 0.95s (shown in blue) and this was statistically different from 0s (we calculate 95% bootstrap confidence intervals).

We also examined the difference in violation rates for each animal in the Odoemene et al. [48] dataset for the disengaged (states 2, 3 and 4) compared to engaged states (state 1), and plot these in Fig. 8c. Violation trials are those where the mouse did not make a decision within the task-specific response period after the appearance of the stimulus. The mean violation rate across all mice and all trials in this dataset was 21% (much higher than in the IBL dataset, where the violation rate was less than 1%), and we found that, across all mice, the violation rate was 3.2% higher in the disengaged states compared to in the engaged state (shown in blue). Given that no information about response times or violation rates was used for to train the GLM-HMM, these analyses provide a useful external validation for our interpretation of the GLM-HMM's latent states as corresponding to 'engaged' and 'disengaged'

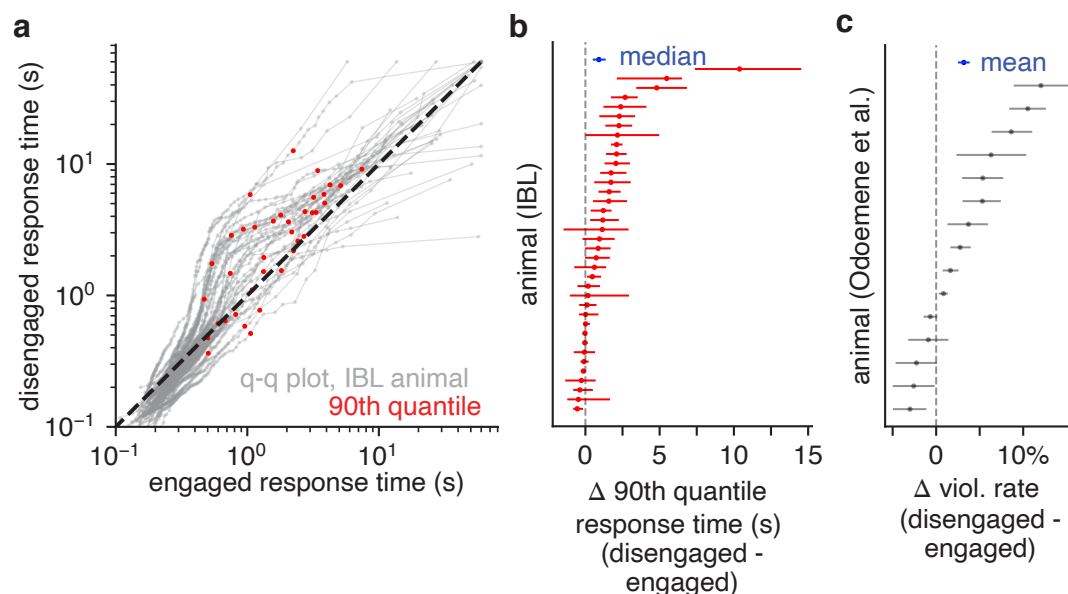


Figure 8: **Behavioral correlates for GLM-HMM states.** (a) Q-Q plots for response time distributions associated with engaged and disengaged (biased left/rightward) states for IBL mice. Each curve is an individual animal, and the red dots indicate the 90th quantile response times. For all 37 mice, a KS-test indicated that the engaged and disengaged response time distributions were statistically different. Furthermore, as can readily be observed from the Q-Q plots, the longest response times typically occurred in the disengaged state. (b) Difference in the 90th quantile response time for the engaged and disengaged states for each IBL animal, as well as 95% bootstrap confidence intervals. Blue is the median difference across mice (0.95s), as well as the 95% bootstrap confidence interval. (c) Violation rate differences for mice in the Odoemene et al. [48] dataset when in the engaged (state 1) and disengaged states (states 2, 3, 4). Error bars are 95% bootstrap confidence intervals; blue indicates the mean difference in violation rate across all mice (3.2%).

behavior.

3 Discussion

In this work, we used the GLM-HMM framework to identify hidden states from perceptual decision-making data. Mouse behavior in two different perceptual decision-making tasks [35, 48] was far better described by a GLM-HMM with sustained engaged and disengaged or biased states. Unlike the classic lapse model, these states alternated on the timescale of tens to hundreds of trials. Additionally we found that, for most of the IBL mice, the slowest response times were associated with the disengaged states, and that mice performing the Odoemene et al. [48] task had higher violation rates in the less engaged states. These behavioral correlates for the GLM-HMM's latent states provide independent validation for our interpretation of the retrieved states as corresponding to 'engaged' and 'disengaged' behavior.

While we found similarities in the strategies pursued by mice performing the different visual detection tasks, we also found some differences. In particular, we found that mice performing the Odoemene et al. [48] task often "warmed up" to the task [34], only reaching their high accuracy, engaged state after one hundred or so trials. In comparison, mice performing the IBL task often started the session in their engaged states. This ability to infer the state or strategy employed by an animal at different times

during a session will be useful for characterizing differences in performance across sessions and across animals. It will also provide a powerful tool for neuroscientists studying the neural mechanisms that support decision-making, as different strategies may well rely on different circuits or different patterns of neural activity [18, 19, 21, 33, 62, 69].

Although we found evidence of warm-up behavior at the start of sessions in the Odoemene et al. [48] mice, we were somewhat surprised to find few signatures of satiation or fatigue toward the end of a session. This might be due to the fact that sessions were typically of fixed duration, and may have ended before mice had a chance to grow satiated or fatigued. One future direction will be to apply the GLM-HMM to experiments with longer sessions, where it may be useful for detecting changes in behavior reflecting satiety or fatigue.

Another exciting future direction will be to explore how the GLM-HMM framework contributes to the wider discussion on the observed differences in lapse rates across species [60, 61]. While the lapse rates of rodents performing perceptual decision-making tasks are often as high as 10-20% [35, 48, 49], lapse rates for non-human primates and humans performing these tasks are often much lower [60, 61]. Could these differences be due to different species using different strategies during these tasks? We look forward to applying the model to data from rats, non-human primates, and humans to explore this problem further.

In future work, we will aim to make explicit the connection between the ‘engaged’ and ‘disengaged’ strategies identified by model and existing measures of arousal and engagement in the literature. Identifying the relationship between the GLM-HMM’s hidden states and pupil diameter, low-frequency LFP oscillations, spontaneous neuronal firing, noise correlations and the action of neuromodulators [14, 31, 54, 64] will be a priority.

That discrete states underpin mouse choice behavior may also call for new normative models to explain why mice may develop these states to begin with [1]. The existence of disengaged and engaged states could reflect explore-exploit behavior [15, 16, 50], optimal learning (e.g., [4, 23, 26, 47]), or could simply indicate incomplete learning of the task. Another promising direction will be to replace the model’s fixed transition matrix with a generalized linear model allowing external covariates to modulate the probability of state changes [12]). This would allow us to identify the factors that influence state changes (e.g., a preponderance of unrewarded choices) and, potentially, seek to control such transitions.

While there are many avenues for future research, we believe that these results call for a significant rethinking of rodent perceptual decision-making behavior and the methods for analyzing it. Indeed, standard analysis methods do not take account of the possibility that an animal makes abrupt changes in decision-making strategy multiple times per session. We feel that the ability to infer internal states from choice behavior will open up new directions for data analysis and provide new insights into a previously inaccessible dimension of perceptual decision-making behavior.

4 Methods

4.1 Inference of GLM-HMM parameters

4.1.1 GLM-HMM objective function

We fit the parameters of the GLM-HMM, $\Theta \equiv \{\pi, A, \{\mathbf{w}_k\}_{k=1}^K\}$, to choice data using Maximum A Posteriori (MAP) estimation via the Expectation Maximization (EM) algorithm [17]. This algorithm iteratively

maximizes the log-posterior of the parameters given the data, given (up to an unknown constant) by:

$$f(\Theta|\mathbf{y}, \{\mathbf{x}\}_{t=1}^T) \equiv \log \sum_s \sum_{\mathbf{z}_s} p(\mathbf{y}_s, \mathbf{z}_s | \{\mathbf{x}_{s,t}\}_{t=1}^T, \Theta) + \log p(\Theta), \quad (5)$$

where s indexes the session in which the data was collected (out of S sessions), and the sum over \mathbf{z}_s is over all possible state allocations for trials in that session. The prior distribution over the model parameters $p(\Theta)$ was given by:

$$p(\Theta) \equiv p(W)p(A)p(\pi) = \left(\prod_{j=1}^K \mathcal{N}(\mathbf{w}_j | 0, \sigma^2 I) \right) \left(\prod_{j=1}^K \text{Dirichlet}(A_j | \alpha) \right) \text{Dirichlet}(\pi | \alpha_\pi). \quad (6)$$

Here, $W \equiv [\mathbf{w}_1 \dots \mathbf{w}_K]$ represents the matrix formed by concatenating the vectors of per-state GLM weights, and $\mathbf{w}_j \in \mathbb{R}^M$ denotes the GLM weights for state j , where $M = 4$ was the number of inputs or covariates (including the bias). For each weight vector, we used an independent zero-mean Gaussian prior with variance σ^2 . For the $(K \times K)$ transition matrix A , we placed an independent Dirichlet prior distribution over each row A_i . The Dirichlet is controlled by a shape parameter α , giving:

$p(A_j | \alpha) = \frac{1}{B(\alpha)} \prod_{i=1}^K (A_{ji})^{\alpha-1}$. We also placed a Dirichlet prior over the initial state distribution with $\alpha_\pi = 1$, which corresponds to a flat prior. In order to select the prior hyperparameters for the transition matrix and the weights for a given dataset, we performed a grid search for $\sigma \in \{0.5, 0.75, 1, 2, 3\}$ and $\alpha \in \{1, 2\}$ and selected the set of hyperparameters that resulted in the best performance on a held-out validation set. For IBL mice, the prior hyperparameters selected were $\sigma = 2$ and $\alpha = 2$. While for mice in the Odoemene et al. [48] dataset, the best hyperparameters were $\sigma = 0.75$ and $\alpha = 2$.

4.1.2 Expectation Maximization Algorithm

We used the Expectation Maximization (EM) algorithm [17, 45] to maximize the objective function, $f(\Theta|\mathbf{y}, \{\mathbf{x}\}_{t=1}^T)$, given in (Eq. 5) with respect to the GLM-HMM parameters. The E-step of the EM algorithm involves forming the Expected Complete log-likelihood (ECLL) using the forward-backward algorithm [3]. The ECLL is a lower bound on $f(\Theta|\mathbf{y}, \{\mathbf{x}\}_{t=1}^T)$ [7, 17, 45], which is then maximized with respect to Θ during the M-Step. Concretely, the Expected Complete log-likelihood is:

$$\begin{aligned} \text{ECLL}(\Theta) &\equiv \sum_s \sum_{\mathbf{z}_s} p(\mathbf{z}_s | \mathbf{y}_s, \{\mathbf{x}_{s,t}\}_{t=1}^T, \Theta^{\text{old}}) \log p(\mathbf{z}_s, \mathbf{y}_s | \Theta, \{\mathbf{x}_{s,t}\}_{t=1}^T) + \log p(\Theta) \\ &= \sum_{s=1}^S \left[\sum_{k=1}^K \gamma_{s,1,k} \log \pi_k + \sum_{t=1}^T \sum_{j=1}^K \sum_{k=1}^K \xi_{s,t,j,k} \log A_{jk} + \sum_{t=1}^T \sum_{k=1}^K \gamma_{s,t,k} \log p(y_{s,t} | z_{s,t} = k, \mathbf{x}_{s,t}, \mathbf{w}_k) \right] \\ &\quad + \log p(\Theta) \end{aligned} \quad (7)$$

where we substituted the definition of the joint distribution,

$$p(\mathbf{y}, \mathbf{z} | \{\mathbf{x}\}_{t=1}^T, \Theta) = p(z_1) p(y_1 | z_1, \mathbf{x}_1) \prod_{t=2}^T p(z_t | z_{t-1}) p(y_t | z_t, \mathbf{x}_t) \quad (8)$$

in order to get to the second line. $\gamma_{s,t,k} \equiv p(z_{s,t} = k | \mathbf{y}_s, \{\mathbf{x}_{s,t}\}_{t=1}^T, \Theta^{\text{old}})$ and $\xi_{s,t,j,k} \equiv p(z_{s,t+1} = k | z_{s,t} = j, \mathbf{y}_s, \{\mathbf{x}_{s,t}\}_{t=1}^T, \Theta^{\text{old}})$ are the quantities estimated through the forward-backward algorithm, as we will describe below; $p(y_{s,t} | z_{s,t} = k, \mathbf{x}_{s,t}, \mathbf{w}_k)$ is obtained from Eq. 4.

4.1.3 Expectation Step

During the E-Step of the EM algorithm, the quantities $\{\gamma_{s,t,k}\}$ and $\{\xi_{s,t,j,k}\}$ were computed using the forward-backward algorithm [3] at the current setting of the GLM-HMM parameters, Θ^{old} . In particular, the forward-backward algorithm involves calculating $a_{s,t,k} \equiv p(\mathbf{y}_{s,[0:t]}, z_{s,t} = k | \{\mathbf{x}_{t'}\}_{t'=1}^t)$ and $b_{s,t,k} \equiv p(\mathbf{y}_{s,[t+1:T]} | z_{s,t} = k, \{\mathbf{x}_{t'}\}_{t'=t+1}^T)$ so as to be able to form $\gamma_{s,t,k}$ and $\xi_{s,t,j,k}$ as follows:

$$\begin{aligned} \gamma_{s,t,k} &\equiv p(z_{s,t} = k | \mathbf{y}_s, \{\mathbf{x}_{s,t}\}_{t=1}^T, \Theta^{\text{old}}) \\ &= \frac{p(\mathbf{y}_{s,[0:t]}, z_{s,t} = k | \{\mathbf{x}_{s,t'}\}_{t'=1}^t, \Theta^{\text{old}}) p(\mathbf{y}_{s,[t+1:T]} | z_{s,t} = k, \{\mathbf{x}_{s,t'}\}_{t'=t+1}^T, \Theta^{\text{old}})}{p(\mathbf{y}_s | \{\mathbf{x}_{s,t}\}_{t=1}^T, \Theta^{\text{old}})} \\ &= \frac{a_{s,t,k} b_{s,t,k}}{\sum_{k=1}^K a_{s,T,k}}. \end{aligned} \quad (9)$$

Similarly,

$$\xi_{s,t,j,k} = \frac{a_{s,t,j} A_{jk} b_{s,t+1,k} p(y_{s,t+1} | z_{t+1} = k, \mathbf{x}_{t+1}, \mathbf{w}_k)}{\sum_{k=1}^K a_{s,T,k}} \quad (10)$$

where, once again, $p(y_{s,t+1} | z_{t+1} = k, \mathbf{x}_{t+1}, \mathbf{w}_k)$ comes from (Eq. 4).

The forward-backward algorithm is a recursive algorithm and calculates the forward probabilities, $\{a_{s,t,k}\}$, as follows:

$$a_{s,1,k} = \pi_k p(y_{s,1} | z_{s,1} = k, \mathbf{x}_{s,1}, \mathbf{w}_k) \quad (11)$$

and, for $1 < t \leq T$:

$$a_{s,t,k} = \sum_{j=1}^K a_{s,t-1,j} A_{jk} p(y_{s,t} | z_{s,t} = k, \mathbf{x}_{s,t}, \mathbf{w}_k) \quad (12)$$

Similarly, the backward probabilities are calculated recursively as follows:

$$b_{s,T,k} = 1 \quad (13)$$

and, for $t \in \{T-1, \dots, 1\}$:

$$b_{s,t,j} = \sum_{k=1}^K b_{s,t+1,k} A_{jk} p(y_{s,t+1} | z_{s,t+1} = k, \mathbf{x}_{s,t+1}, \mathbf{w}_k). \quad (14)$$

The forward-backward algorithm enables evaluation of the sum over all possible state allocations in Eq. 5 and Eq. 7 in linear time in the number of trials, rather than in exponential time.

4.1.4 Maximization Step

After running the forward-backward algorithm, we formed the ECLL as in Eq. 7 and then maximized it with respect to the GLM-HMM parameters, Θ . For the initial state distribution π and the transition matrix A , this resulted in the closed-form updates:

$$\pi_k^{\text{new}} = \frac{\sum_{s=1}^S \gamma_{s,1,k}}{\sum_{s=1}^S \sum_{j=1}^K \gamma_{s,1,j}} \quad (15)$$

$$A_{jk}^{\text{new}} = \frac{\alpha - 1 + \sum_{s=1}^S \sum_{t=2}^T \xi_{s,t,j,k}}{K(\alpha - 1) + \sum_{s=1}^S \sum_{t=2}^T \sum_{k=1}^K \xi_{s,t,j,k}}. \quad (16)$$

For the GLM weights, there was no such closed form update, but a Bernoulli GLM falls into the class of functions mapping external inputs to HMM emission probabilities considered in [22], so we know that the ECLL is concave in the GLM weights. As such, we were able to numerically find the GLM weights that maximize (not just locally but globally) the ECLL using the BFGS algorithm [9, 24, 28, 58] as implemented by the scipy optimize function in python [65].

4.1.5 Comparing states across animals and GLM-HMM parameter initialization

In Fig. 4 and Fig. 7, we show the results from fitting a single GLM-HMM to each animal; however it is nontrivial to map the retrieved states across animals to one another. As such, we employed a multistage fitting procedure that allowed us to make this comparison, and we detail this procedure in Algorithm 1. In the first stage, we concatenated the data for all animals in a single dataset together (for example, in the case of the IBL dataset, this would be the data for all 37 animals). We then fit a GLM (a 1 state GLM-HMM) to the concatenated data using Maximum Likelihood estimation. We used the fit GLM weights to initialize the GLM weights of a K -state GLM-HMM that we again fit to the concatenated dataset from all animals together (to obtain a “global fit”). We added Gaussian noise with $\sigma_{\text{init}} = 0.2$ to the the GLM weights, so that the initialized states were distinct, and we initialized the transition matrix of the K -State GLM-HMM as $0.95 \times \mathbb{1} + \mathcal{N}(0, \Sigma_{\text{trans.}})$ where $\Sigma_{\text{trans.}} \in \mathbb{R}^{K^2 \times K^2}$ and $\Sigma_{\text{trans.}} = 0.05 \times \mathbb{1}$. We then normalized this so that that rows of the transition matrix added up to 1, and represented probabilities. While the EM algorithm is guaranteed to converge to a local optimum in the log probability landscape of Eq. 5, there is no guarantee that it will converge to the global optimum [57]. Correspondingly, for each value of K , we fit the model 20 times using 20 different initializations.

In the next stage of the fitting procedure, we wanted to obtain a separate GLM-HMM fit for each animal, so we initialized a model for each animal with the GLM-HMM global fit parameters from all animals together (out of the 20 initializations, we chose the model that resulted in the best training set log-likelihood). We then ran the EM algorithm to convergence; it is these recovered parameters that are shown in Fig. 4 and Fig. 7. By initializing each individual animal’s model with the parameters from the fit to all animals together, it was no longer necessary for us to permute the retrieved states from each animal so as to map semantically similar states to one another.

We note that the initialization scheme detailed above is sufficiently robust so as to allow recovery of GLM-HMM parameters in various parameter regimes of interest. In particular, we simulated datasets from a GLM-HMM with the global fit parameters for both the IBL and Odoemene et al. [48] datasets, as well as a global fit lapse model. We show the results of these recovery analyses in Fig. S3 and Fig. S4.

4.2 Assessing Model Performance

4.2.1 Cross Validation

There are two ways in which to perform cross-validation when working with Hidden Markov Models. Firstly, it is possible to hold out entire sessions of choice data for assessing test-set performance. That is, when fitting the model, the objective function in Eq. 5 and the ECLL in Eq. 7 are modified to

Algorithm 1 Multistage GLM-HMM fitting procedure

```

1: Fit GLM (1 state GLM-HMM) to all data from all animals
2: Fit global GLM-HMM to all data from all animals:
3: for  $K \in \{2, \dots, 5\}$  do
4:   for  $\text{init.} \in \{1, \dots, 20\}$  do
5:     Initialize K-state GLM-HMM using noisy GLM weights
6:     Run EM algorithm until convergence
7:   end for
8: end for
9: Fit separate GLM-HMM to each animal by initializing with global fit:
10: for each individual animal do
11:   for  $K \in \{2, \dots, 5\}$  do
12:     Initialize K-state GLM-HMM using best global GLM-HMM parameters for this K
13:     Run EM algorithm until convergence
14:   end for
15: end for

```

only include 80% of sessions (since we use 5-fold cross-validation throughout this work); and the log-likelihood of the held-out 20% of sessions is calculated using the fit parameters, and a single run of the forward-pass on the held-out sessions:

$$\text{LL}_{\text{test}} \equiv \sum_{s \in S \setminus S'} \log p(\mathbf{y}_s | \{\mathbf{x}_{s,t}\}_{t=1}^T, \Theta') = \log \sum_{s \in S'} \sum_{k=1}^K a_{s,T,k} \quad (17)$$

where $S \setminus S'$ is the set of held out sessions, and Θ' is the set of GLM-HMM parameters obtained by fitting the model using the trials from S' .

The second method of performing cross-validation involves holding out 20% of trials within a session. When fitting the model, the third term in the ECLL is modified so as to exclude these trials and is now $\sum_{t \in T'} \sum_{k=1}^K \gamma_{s,t,k} \log p(y_{s,t,k} | z_{s,t} = k, \mathbf{x}_{s,t} \mathbf{w}_k)$, where T' is the set of trials to be used to fit the model. Furthermore, the calculation of the posterior state probabilities, $\gamma_{s,t,k}$ and $\xi_{s,t,j,k}$, is also modified so as to exclude the test set choice data. In particular, $\gamma_{s,t,k}$ is now $p(z_{s,t} | \{y_{s,t'}\}_{t' \in T'}, \{\mathbf{x}_{s,t'}\}_{t' \in T'}, \Theta^{\text{old}})$ and similarly $\xi_{s,t,j,k}$ is now $p(z_{s,t+1} = k | z_{s,t} = j, \{y_{s,t'}\}_{t' \in T'}, \{\mathbf{x}_{s,t'}\}_{t' \in T'}, \Theta^{\text{old}})$. The method of calculating these modified posterior probabilities is as detailed in Eq. 9 and Eq. 10, but now the calculation of the forward and backward probabilities, $a_{s,t,k}$ and $b_{s,t,k}$ in Eq. 11, Eq. 12, Eq. 13 and Eq. 14 is modified so that, on trials that are identified as test trials, the $p(y_{s,t} | z_{s,t} = k, \mathbf{x}_{s,t}, \mathbf{w}_k)$ term in these equations is replaced with 1.

In Fig. 2, Fig. 4 and Fig. 7, we perform cross-validation by holding out entire sessions. We believed it would be harder to make good predictions on entire held out sessions, compared to single trials within a session, as we thought that mice would exhibit more variability in behavior across sessions compared to within sessions. When we compare the performance of the GLM-HMM against the PsyTrack model of [55] in Fig. 6, we use the second method of cross-validation so as to use the same train and test sets as PsyTrack (PsyTrack cannot make predictions on entire held-out sessions).

4.2.2 Normalized Log-likelihood

In Fig. 2, Fig. 4 and Fig. 7 we report the normalized log-likelihood of different models on held-out sessions. This is calculated as follows:

$$\text{NLL}(\Theta) = \frac{\text{LL}_{\text{test}} - \text{LL}_0}{n_{\text{test}} \log(2)} \quad (18)$$

where, for the GLM-HMM, LL_{test} is the test set log-likelihood as calculated in Eq. 17, and LL_0 is the log-likelihood of the same test set under a Bernoulli model of animal choice behavior. Specifically, this baseline model assumes that animals flip a coin on each trial so as to decide to go Right, and the probability of going Right is equal to the fraction of trials in the training set in which the animal chose to go to the Right. n_{test} is the number of trials in the test set, and is important to include since LL_{test} depends on the number of trials in the test set. Dividing by $\log(2)$ gives the Normalized log-likelihood the units of bits per trial. Clearly, larger values of the Normalized log-likelihood are better, with a value of 0 indicating that a model offers no improvement in prediction compared to the crude baseline model described above. However, even small values of normalized log-likelihood can indicate a large improvement in predictive power. For a test set size of $n_{\text{test}} = 500$, a normalized log-likelihood value of 0.01 indicates that the test data is 31.5 times more likely to have been generated with the GLM-HMM compared to the baseline model. For a test set of $n_{\text{test}} = 5000$, and the same value of NLL, the test set becomes 1×10^{15} times more likely under the GLM-HMM compared to the baseline model!

4.2.3 Predictive Accuracy

In Fig. 2, Fig. 4 and Fig. 7, we also report the predictive accuracy of the GLM-HMM. When calculating the predictive accuracy, we employ a method similar to the second method described above in section 4.2.1. In particular, we hold out 20% of trials and then obtain the posterior state probabilities for these trials, $t'' \in \{T \setminus T'\}$, as $\gamma_{s,t'',k} = p(z_{s,t''} | \{y_{s,t'}\}_{t' \in T'}, \{\mathbf{x}_{s,t'}\}_{t' \in T'}, \Theta)$, using the other 80% of trials (this latter set of trials being labeled T'). We then calculate the probability of the held-out choices being to go Right as:

$$p_{R,s,t''} \equiv p(y_{s,t''} = 1 | \{y_{s,t'}\}_{t' \in T'}, \{\mathbf{x}_{s,t'}\}_{t' \in T'}, \mathbf{x}_{s,t''}, \Theta) = \sum_{k=1}^k (\gamma_{s,t'',k}) p(y_{s,t''} = 1 | z_{s,t''} = k, \mathbf{x}_{s,t''}, \mathbf{w}_k) \quad (19)$$

We then calculate the predictive accuracy as:

$$\text{predictive accuracy} = \frac{\sum_{t'' \in T \setminus T'} \left(\left(\mathbb{1}(y_{s,t''} = 1) \mathbb{1}(p_{R,s,t''} > 0.5) \right) + \left(\mathbb{1}(y_{s,t''} = 0) \mathbb{1}(p_{R,s,t''} \leq 0.5) \right) \right)}{|T \setminus T'|} \quad (20)$$

4.3 Comparison with PsyTrack Model of Roy et al.

The PsyTrack model of Roy et al. [56] assumes an animal makes its choice at trial t according to

$$p(y_t = 1 | \mathbf{x}_t) = \frac{1}{1 + e^{-\mathbf{w}_t \cdot \mathbf{x}_t}} \quad (21)$$

where \mathbf{w}_t evolves according to

$$\mathbf{w}_t = \mathbf{w}_{t-1} + \eta_t \quad (22)$$

where $\eta_t \sim \mathcal{N}(0, \text{diag}(\sigma_1^2, \dots, \sigma_M^2))$ and $\mathbf{w}_0 \sim \mathcal{N}(0, \text{diag}(\sigma_{\text{init},1}^2, \dots, \sigma_{\text{init},M}^2))$. Specifically, the animal is assumed to use a set of slowly changing weights to make its decision on each trial.

In order to perform model comparison with the PsyTrack model of [55, 56], we utilized the code provided at <https://github.com/nicholas-roy/psytrack>.

4.3.1 Simulating Choice Data with AR(1) Model for Weights

While the GLM-HMM is a generative model that can be readily used to simulate choice data that resembles, in accuracy and in the resulting psychometric curves, the choice data of real animals, this is not true for the PsyTrack model of [55, 56]. Indeed, specifying only the $\{\sigma_{\text{init},m}^2\}$, $\{\sigma_m^2\}$ hyperparameters of that model and then generating weights according to Eq. 22 and choice data according to Eq. 21 will likely result in choice behavior that is vastly different from that of real animals (the PsyTrack model is underconstrained as a generative model). As such, so as to produce Fig. S11 we simulated smoothly evolving weights from an AR(1) model where the parameters of this model were obtained using the PsyTrack fits to real data. Specifically, we assumed that the probability of going rightward at time t was given by:

$$p(y_t = 1 | \mathbf{w}_t, \mathbf{x}_t) = \frac{1}{1 + e^{-\mathbf{w}_t \cdot \mathbf{x}_t}} \quad (23)$$

and we assumed that the weights evolved according to an AR(1) process as follows:

$$w_{m,t} = \alpha_m w_{m,t-1} + (1 - \alpha_m) \bar{w}_m^{\text{PsyTrack}} + \eta_{m,t} \quad (24)$$

where $w_{m,t}$ is the m th element of \mathbf{w}_t in Eq. 23 and $\bar{w}_m^{\text{PsyTrack}}$ is the average weight that an animal places on covariate m across all trials when fit with the PsyTrack model: $\bar{w}_m^{\text{PsyTrack}} = \frac{1}{T} \sum_t w_{t,m}^{\text{PsyTrack}}$. We obtained α_m by regressing $w_{t,m}^{\text{PsyTrack}}$ against $w_{t-1,m}^{\text{PsyTrack}}$ and taking the retrieved slope (after confirming that the retrieved slope had a magnitude of less than 1, so as to ensure that the weights did not diverge as $t \rightarrow \infty$). $\eta_{m,t} \sim \mathcal{N}(0, \tilde{\sigma}_m^2)$, where $\tilde{\sigma}_m^2 = (1 - \alpha_m^2) \sigma_m^2$ and σ_m^2 was obtained from the PsyTrack fit to an animal's real choice data. Finally, we set $\mathbf{w}_0 = 0$.

We simulated weight trajectories for thousands of trials for each animal, so that the AR(1) process reached the stationary regime for each covariate, and so that the mean, variance and autocovariance of each weight for each covariate were close to those returned by the PsyTrack fits to the real choice data.

4.3.2 Additional details about Figure S11

In order to produce supplementary Fig. S11, we aligned sets of 10 trials for which there was a state transition after trial 5 either out of the engaged state or into the engaged state. Animals were in the same state for all 5 trials prior to the transition and in the same state for all 5 trials after the transition (we excluded sequences where state switches occurred during the first 5 or latter 5 trials). Each grey line in Fig. S11 is the average accuracy across all eligible 10 trial sequences for a particular animal; we required that all animals had at least 30 sequences to be included in Fig. S11. As a result, 27 (out of 37) animals are shown in Fig. S11 and are used to compute the average curve across all animals.

4.4 Datasets studied

In this paper, we applied the GLM-HMM to two publicly available behavioral datasets associated with recent publications. Firstly, we studied the data associated with [35] that is made available via figshare at <https://doi.org/10.6084/m9.figshare.11636748>. We used the framework developed in [36] to access the data. We modeled the choice data for the 37 animals in this dataset which had more than 30 sessions of data during the ‘bias block’ regime. We focused on this regime because of the fact that mice, when they have reached this regime, understand the rules of the task and exhibit stationary behavior (see Fig. S1 for plots of accuracy against session identity for each animal, as well as Fig. S2 for the psychometric curves for these animals for the trials studied). For each session, we subset to the first 90 trials of data because, during these trials, the stimulus was equally likely to appear on the left or right of the screen. After the first 90 trials, the structure of the task changed and for a block of trials, the stimulus appeared on the left with a probability of either 80% or 20%; the block identity switched multiple times throughout a session, so that 80% and 20% blocks were interleaved. We subset to the animals with more than 30 sessions of data because we were able to confidently recover GLM-HMM and lapse model generative parameters when we simulated datasets with this number of trials: see Figures Fig. S3 and Fig. S4. As a sanity check to make sure that the recovered states and transitions were not a consequence of the animals we study having been exposed to bias blocks in earlier sessions, we obtained data for 4 animals that were never exposed to bias blocks (not included in the publicly released dataset) and fit the GLM-HMM to the choice data for these animals. In Fig. S12, we show that the retrieved states, dwell times and model comparison results for these animals look very similar to those shown in Fig. 4.

The second dataset that we studied was that associated with [48], with the data being made available at <https://doi.org/10.14224/1.38944>. Once again, we studied sessions after animals had learned the task (see Fig. S13 and Fig. S14). For this dataset, the retrieved states were less distinct compared to those for the IBL dataset, and as such, we required more trials to be able to recover the generative parameters in simulated data: see Fig. S3. We thus subset to the 15 animals with more than 20 sessions of data and 12,000 trials of data. Compared to the IBL dataset, where the violation rate across all animals’ data was less than 1% of trials (where a violation is where the animal chose not to respond), the violation rate across the 15 animals that we studied from this second dataset was 21%. Thus, it was important to develop a principled method for dealing with violation trials. We treated violation trials as trials with missing choice data, and we handled these trials in a similar way to how we handled test data when performing the second type of cross-validation described in section 4.2.1 above. That is, we modified the third term of the ECLL given in Eq. 7 to exclude violation trials, and we modified the definition of the posterior state probabilities for these trials to be $\gamma_{s,t,k} = p(z_{s,t} | \{y_{s,t'}\}_{t' \in T'}, \{\mathbf{x}_{s,t'}\}_{t' \in T'}, \Theta^{\text{old}})$ and $\xi_{s,t,j,k} = p(z_{s,t+1} = k | z_{s,t} = j, \{y_{s,t'}\}_{t' \in T'}, \{\mathbf{x}_{s,t'}\}_{t' \in T'}, \Theta^{\text{old}})$, where T' , rather than representing the training set data, is now the set of non-violation trials. The calculation of the forward and backward probabilities, $a_{s,t,k}$ and $b_{s,t,k}$, was modified so that, on violation trials, in Eq. 11, Eq. 12, Eq. 13 and Eq. 14, the $p(y_{s,t} | z_{s,t} = k, \mathbf{x}_{s,t}, \mathbf{w}_k)$ term was replaced with 1.

4.4.1 Forming the Design Matrix

Each of the models discussed in this paper (GLM-HMM, the classic lapse model, the PsyTrack model of Roy et al. [56]) were fit using a design matrix of covariates, $X \in \mathbb{R}^{T \times M}$, where T was the number of trials of choice data for a particular animal. A single row in this matrix was the vector of covariates, $\mathbf{x}_t \in \mathbb{R}^M$, influencing the animal’s choice at trial t . For all analyses presented in text, unless specified otherwise, $M = 4$.

For both tasks, the first column in the design matrix was the z-scored stimulus intensity. For the IBL task, we calculated the stimulus intensity as the difference in the value of the visual contrast on the right side of the screen minus the visual contrast on the left of the screen. This resulted in 9 different values for the ‘signed contrast’: $\{-100, -25, -12.5, -6.25, 0, 6.25, 12.5, 25, 100\}$. We then z-scored this difference quantity across all trials. For the Odoemene et al. [48] task, we subtracted the 12Hz threshold from the flash rate presented on each trial and then z-scored the resulting quantity.

For all trials, all animals and both tasks, the second column of the design matrix was set to 1, so as to enable us to capture the animal’s innate bias for going rightward or leftward. The third column in the design matrix was the animal’s choice on the previous trial $X_{t,3} \equiv 2y_{t-1} - 1$. Whereas $y_{t-1} \in \{0, 1\}$, $X_{t,3} \in \{-1, 1\}$. It is not strictly necessary to perform this scaling, but we did so to ensure that the range of values for $X_{:,1}$ and $X_{:,3}$ were more similar (which can be useful when performing parameter optimization). Finally, the fourth column in the design matrix was the win-stay-lose-switch covariate, which was calculated as $X_{:,4} \equiv r_{t-1} \times (2y_{t-1} - 1)$, where $r_{t-1} \in \{-1, 1\}$ was a binary variable indicating whether or not the animal was rewarded on the previous trial. Again $X_{:,4} \in \{-1, 1\}$.

4.5 Code availability

We contributed code to the Bayesian State Space Modeling framework of [42] and we use this code base to perform GLM-HMM inference. The code to analyze the resulting model fits and to produce the figures in this paper will be made available at <https://github.com/zashwood/glm-hmm>.

5 Acknowledgments

We are grateful to Miles Wells, Rebecca Terry and the Cortexplab at University College London for providing us with the data for the 4 mice plotted in Fig. S12. We are grateful to Scott Linderman for developing the beautiful Bayesian State Space Modeling framework of [42]: as described in our Methods section, we built our code on top of this framework. We thank members of the Pillow Lab, the International Brain Laboratory (IBL), and specifically the Behavior Analysis Working Group within the IBL for helpful feedback throughout the project. We thank Peter Dayan, Sebastian Bruijns and Liam Paninski for acting as the IBL Review Board for this paper. We thank Anne Urai and Emily Dennis for their feedback at various points during the project. We thank Abigail Russo and Matthew Whiteway for providing feedback on drafts of this manuscript. We thank Hannah Bayer for her help and advice as we were preparing to submit this paper. This work was supported by grants from the Simons Collaboration on the Global Brain (SCGB AWD543027; JWP), the NIH BRAIN initiative (NS104899 and R01EB026946; JWP), and a U19 NIH-NINDS BRAIN Initiative Award (5U19NS104648; JWP).

References

- [1] Ashwood, Z., Roy, N. A., Bak, J. H., and Pillow, J. W. (2020). Inferring learning rules from animal decision-making. *Advances in Neural Information Processing Systems*, 33:3442–3453.
- [2] Batty, E., Whiteway, M., Saxena, S., Biderman, D., Abe, T., Musall, S., Gillis, W., Markowitz, J., Churchland, A., Cunningham, J. P., Datta, S. R., Linderman, S., and Paninski, L. (2019). BehaveNet: nonlinear embedding and Bayesian neural decoding of behavioral videos. In Wallach, H., Larochelle,

- 670 H., Beygelzimer, A., Alché-Buc, F. d., Fox, E., and Garnett, R., editors, *Advances in Neural Informa-*
671 *tion Processing Systems 32*, pages 15706–15717. Curran Associates, Inc.
- 672 [3] Baum, L. E., Petrie, T., Soules, G., and Weiss, N. (1970). A Maximization Technique Occurring in the
673 Statistical Analysis of Probabilistic Functions of Markov Chains. *Annals of Mathematical Statistics*,
674 41(1):164–171.
- 675 [4] Behrens, T. E. J., Woolrich, M. W., Walton, M. E., and Rushworth, M. F. S. (2007). Learning the
676 value of information in an uncertain world. *Nature Neuroscience*, 10(9):1214–1221. Number: 9
677 Publisher: Nature Publishing Group.
- 678 [5] Bengio, Y. and Frasconi, P. (1995). An Input Output HMM Architecture. In Tesauro, G., Touretzky,
679 D. S., and Leen, T. K., editors, *Advances in Neural Information Processing Systems 7*, pages 427–
680 434. MIT Press.
- 681 [6] Berman, G. J. (2018). Measuring behavior across scales. *BMC Biology*, 16(1):23.
- 682 [7] Bishop, C. M. (2006). *Pattern recognition and machine learning*. springer.
- 683 [8] Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal
684 decision making: a formal analysis of models of performance in two-alternative forced-choice tasks.
685 *Psychological review*, 113(4):700. Publisher: American Psychological Association.
- 686 [9] Broyden, C. G. (1970). The Convergence of a Class of Double-rank Minimization Algorithms 1.
687 General Considerations. *IMA Journal of Applied Mathematics*, 6(1):76–90. Publisher: Oxford Aca-
688 demic.
- 689 [10] Burgess, C. P., Lak, A., Steinmetz, N. A., Zatzka-Haas, P., Bai Reddy, C., Jacobs, E. A. K., Linden,
690 J. F., Paton, J. J., Ranson, A., Schröder, S., Soares, S., Wells, M. J., Wool, L. E., Harris, K. D.,
691 and Carandini, M. (2017). High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in
692 Head-Fixed Mice. *Cell Reports*, 20(10):2513–2524.
- 693 [11] Busse, L., Ayaz, A., Dhruv, N. T., Katzner, S., Saleem, A. B., Schölvink, M. L., Zaharia, A. D.,
694 and Carandini, M. (2011). The Detection of Visual Contrast in the Behaving Mouse. *Journal of*
695 *Neuroscience*, 31(31):11351–11361.
- 696 [12] Calhoun, A. J., Pillow, J. W., and Murthy, M. (2019). Unsupervised identification of the internal
697 states that shape natural behavior. *Nature Neuroscience*, 22(12):2040–2049. Number: 12 Publisher:
698 Nature Publishing Group.
- 699 [13] Carandini, M. and Churchland, A. K. (2013). Probing perceptual decisions in rodents. *Nature*
700 *Neuroscience*, 16(7):824–831.
- 701 [14] Carter, M. E., Yizhar, O., Chikahisa, S., Nguyen, H., Adamantidis, A., Nishino, S., Deisseroth, K.,
702 and de Lecea, L. (2010). Tuning arousal with optogenetic modulation of locus coeruleus neurons.
703 *Nature Neuroscience*, 13(12):1526–1533. Number: 12 Publisher: Nature Publishing Group.
- 704 [15] Daw, N. D., O’Doherty, J. P., Dayan, P., Seymour, B., and Dolan, R. J. (2006). Cortical substrates
705 for exploratory decisions in humans. *Nature*, 441(7095):876–879.
- 706 [16] Dayan, P. and Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive,*
707 *Affective, & Behavioral Neuroscience*, 8(4):429–453. Publisher: Springer.

- [17] Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.
- [18] Ding, L. and Gold, J. I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cerebral Cortex (New York, N.Y.: 1991)*, 22(5):1052–1067.
- [19] Ding, L. and Gold, J. I. (2013). The basal ganglia’s contributions to perceptual decision making. *Neuron*, 79(4):640–649.
- [20] Duda, R. O., Hart, P. E., and Stork, D. G. (2012). *Pattern Classification*. John Wiley & Sons. Google-Books-ID: Br33IRC3PkQC.
- [21] Engel, T. A., Steinmetz, N. A., Gieselmann, M. A., Thiele, A., Moore, T., and Boahen, K. (2016). Selective modulation of cortical state during spatial attention. *Science*, 354(6316):1140–1144. Publisher: American Association for the Advancement of Science Section: Report.
- [22] Escola, S., Fontanini, A., Katz, D., and Paninski, L. (2011). Hidden Markov models for the stimulus-response relationships of multistate neural systems. *Neural computation*, 23(5):1071–1132.
- [23] Fiser, J., Berkes, P., Orbán, G., and Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3):119–130.
- [24] Fletcher, R. (1970). A new approach to variable metric algorithms. *The Computer Journal*, 13(3):317–322. Publisher: Oxford Academic.
- [25] Fründ, I., Wichmann, F. A., and Macke, J. H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision*, 14(7):9–9. Publisher: The Association for Research in Vision and Ophthalmology.
- [26] Gershman, S. J. (2020). Origin of perseveration in the trade-off between reward and complexity. *bioRxiv*, page 2020.01.16.903476. Publisher: Cold Spring Harbor Laboratory Section: New Results.
- [27] Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual review of neuroscience*, 30.
- [28] Goldfarb, D. (1970). A family of variable-metric methods derived by variational means. *Mathematics of Computation*, 24(109):23–26.
- [29] Gomez-Marin, A., Paton, J. J., Kampff, A. R., Costa, R. M., and Mainen, Z. F. (2014). Big behavioral data: psychology, ethology and the foundations of neuroscience. *Nature Neuroscience*, 17(11):1455–1462. Number: 11 Publisher: Nature Publishing Group.
- [30] Green, D. M. and Swets, J. A. (1966). *Signal detection theory and psychophysics*, volume 1. Wiley New York.
- [31] Harris, K. D. and Thiele, A. (2011). Cortical state and attention. *Nature Reviews Neuroscience*, 12(9):509–523. Number: 9 Publisher: Nature Publishing Group.
- [32] Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in Neuroscience*, 8. Publisher: Frontiers.

- [33] Hennig, J. A., Oby, E. R., Golub, M. D., Bahureksa, L. A., Sadtler, P. T., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Chase, S. M., and Yu, B. M. (2021). Learning is shaped by abrupt changes in neural engagement. *Nature Neuroscience*, pages 1–10. Publisher: Nature Publishing Group.
- [34] Iigaya, K., Fonseca, M. S., Murakami, M., Mainen, Z. F., and Dayan, P. (2018). An effect of serotonergic stimulation on learning rates for rewards apparent after long intertrial intervals. *Nature Communications*, 9(1):2477. Number: 1 Publisher: Nature Publishing Group.
- [35] International Brain Laboratory, T., Aguillon-Rodriguez, V., Angelaki, D. E., Bayer, H. M., Bonacchi, N., Carandini, M., Cazes, F., Chapuis, G. A., Churchland, A. K., Dan, Y., Dewitt, E. E., Faulkner, M., Forrest, H., Haetzel, L. M., Hausser, M., Hofer, S. B., Hu, F., Khanal, A., Krasniak, C. S., Laranjeira, I., Mainen, Z. F., Meijer, G. T., Miska, N. J., Mrcic-Flogel, T. D., Murakami, M., Noel, J.-P., Pan-Vazquez, A., Sanders, J. I., Socha, K. Z., Terry, R., Urai, A. E., Vergara, H. M., Wells, M. J., Wilson, C. J., Witten, I. B., Wool, L. E., and Zador, A. (2020a). A standardized and reproducible method to measure decision-making in mice. *bioRxiv*, page 2020.01.17.909838.
- [36] International Brain Laboratory, T., Bonacchi, N., Chapuis, G., Churchland, A., Harris, K. D., Hunter, M., Rossant, C., Sasaki, M., Shen, S., Steinmetz, N. A., Walker, E. Y., Winter, O., and Wells, M. (2020b). Data architecture for a large-scale neuroscience collaboration. *bioRxiv*. Publisher: Cold Spring Harbor Laboratory _eprint: <https://www.biorxiv.org/content/early/2020/02/23/827873.full.pdf>.
- [37] Johnson, R. E., Linderman, S., Panier, T., Wee, C. L., Song, E., Herrera, K. J., Miller, A., and Engert, F. (2020). Probabilistic Models of Larval Zebrafish Behavior Reveal Structure on Many Scales. *Current Biology*, 30(1):70–82.e4.
- [38] Klein, S. A. (2001). Measuring, estimating, and understanding the psychometric function: A commentary. *Perception & psychophysics*, 63(8):1421–1455.
- [39] Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., and Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron*, 93(3):480–490.
- [40] Lak, A., Hueske, E., Hirokawa, J., Masset, P., Ott, T., Urai, A. E., Donner, T. H., Carandini, M., Tonegawa, S., Uchida, N., and Kepecs, A. (2020). Reinforcement biases subsequent perceptual decisions when confidence is low, a widespread behavioral phenomenon. *eLife*, 9:e49834. Publisher: eLife Sciences Publications, Ltd.
- [41] Law, C.-T. and Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nature Neuroscience*, 12(5):655–663. Number: 5 Publisher: Nature Publishing Group.
- [42] Linderman, S., Antin, B., and Zoltowski, D. (2020). SSM: Bayesian learning and inference for state space models.
- [43] Lindsay, B. G. (1995). Mixture Models: Theory, Geometry and Applications. *NSF-CBMS Regional Conference Series in Probability and Statistics*, 5:i–163. Publisher: Institute of Mathematical Statistics.
- [44] Lueckmann, J.-M., Macke, J. H., and Nienborg, H. (2018). Can Serial Dependencies in Choices and Neural Activity Explain Choice Probabilities? *Journal of Neuroscience*, 38(14):3495–3506. Publisher: Society for Neuroscience Section: Research Articles.

- 784 [45] McLachlan, G. J. and Krishnan, T. (2007). *The EM Algorithm and Extensions*. John Wiley & Sons.
785 Google-Books-ID: NBawzaWoWa8C.
- 786 [46] McLachlan, G. J. and Peel, D. (2004). *Finite Mixture Models*. John Wiley & Sons. Google-Books-
787 ID: c2_fAox0DQoC.
- 788 [47] Mendonça, A. G., Drugowitsch, J., Vicente, M. I., DeWitt, E. E. J., Pouget, A., and Mainen, Z. F.
789 (2020). The impact of learning on perceptual decisions and its implication for speed-accuracy trade-
790 offs. *Nature Communications*, 11(1):2757. Number: 1 Publisher: Nature Publishing Group.
- 791 [48] Odoemene, O., Pisupati, S., Nguyen, H., and Churchland, A. K. (2018). Visual Evidence Accumu-
792 lation Guides Decision-Making in Unrestrained Mice. *Journal of Neuroscience*, 38(47):10143–10155.
- 793 [49] Pinto, L., Koay, S. A., Engelhard, B., Yoon, A. M., Deverett, B., Thiberge, S. Y., Witten, I. B., Tank,
794 D. W., and Brody, C. D. (2018). An Accumulation-of-Evidence Task Using Visual Pulses for Mice
795 Navigating in Virtual Reality. *Frontiers in Behavioral Neuroscience*, 12. Publisher: Frontiers.
- 796 [50] Pisupati, S., Chartarifsky-Lynn, L., Khanal, A., and Churchland, A. K. (2021). Lapses in perceptual
797 decisions reflect exploration. *eLife*, 10:e55490.
- 798 [51] Prins, N. (2012). The psychometric function: The lapse rate revisited. *Journal of Vision*, 12(6):25–
799 25.
- 800 [52] Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice
801 decision tasks. *Neural computation*, 20(4):873–922. Publisher: MIT Press.
- 802 [53] Ratcliff, R., Smith, P. L., Brown, S. D., and McKoon, G. (2016). Diffusion Decision Model: Current
803 Issues and History. *Trends in cognitive sciences*, 20(4):260–281.
- 804 [54] Reimer, J., Froudarakis, E., Cadwell, C. R., Yatsenko, D., Denfield, G. H., and Tolia, A. S. (2014).
805 Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron*, 84(2):355–
806 362.
- 807 [55] Roy, N. A., Bak, J. H., Akrami, A., Brody, C., and Pillow, J. W. (2018). Efficient inference for time-
808 varying behavior during learning. In *Advances in Neural Information Processing Systems*, pages
809 5695–5705.
- 810 [56] Roy, N. A., Bak, J. H., International Brain Laboratory, Akrami, A., Brody, C. D., and Pillow, J. W.
811 (2021). Extracting the dynamics of behavior in sensory decision-making experiments. *Neuron*,
812 109(4):597–610.e6.
- 813 [57] Salakhutdinov, R., Roweis, S. T., and Ghahramani, Z. (2003). Optimization with EM and
814 expectation-conjugate-gradient. In *Proceedings of the 20th International Conference on Machine*
815 *Learning (ICML-03)*, pages 672–679.
- 816 [58] Shanno, D. F. (1970). Conditioning of quasi-Newton methods for function minimization. *Mathemat-*
817 *ics of Computation*, 24(111):647–656.
- 818 [59] Sharma, A., Johnson, R., Engert, F., and Linderman, S. (2018). Point process latent variable
819 models of larval zebrafish behavior. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-
820 Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages
821 10919–10930. Curran Associates, Inc.

- [60] Sheppard, J. P., Raposo, D., and Churchland, A. K. (2013). Dynamic weighting of multisensory stimuli shapes decision-making in rats and humans. *Journal of Vision*, 13(6):4–4. Publisher: The Association for Research in Vision and Ophthalmology.
- [61] Shevinsky, C. A. and Reinagel, P. (2019). The Interaction Between Elapsed Time and Decision Accuracy Differs Between Humans and Rats. *Frontiers in Neuroscience*, 13. Publisher: Frontiers.
- [62] Steinmetz, N. A., Zatka-Haas, P., Carandini, M., and Harris, K. D. (2019). Distributed coding of choice, action and engagement across the mouse brain. *Nature*, 576(7786):266–273. Number: 7786 Publisher: Nature Publishing Group.
- [63] Urai, A. E., de Gee, J. W., Tsetsos, K., and Donner, T. H. (2019). Choice history biases subsequent evidence accumulation. *eLife*, 8:e46331. Publisher: eLife Sciences Publications, Ltd.
- [64] Vinck, M., Batista-Brito, R., Knoblich, U., and Cardin, J. A. (2015). Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron*, 86(3):740–754.
- [65] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Jarrod Millman, K., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C., Polat, İ., Feng, Y., Moore, E. W., Vand erPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and Contributors, S. . . (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.
- [66] Wichmann, F. A. and Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics*, 63(8):1293–1313.
- [67] Wiltischko, A. B., Johnson, M. J., Iurilli, G., Peterson, R. E., Katon, J. M., Pashkovski, S. L., Abaira, V. E., Adams, R. P., and Datta, S. R. (2015). Mapping Sub-Second Structure in Mouse Behavior. *Neuron*, 88(6):1121–1135. Publisher: Elsevier.
- [68] Yu, S.-Z. (2010). Hidden semi-Markov models. *Artificial Intelligence*, 174(2):215–243.
- [69] Zhou, Y. and Freedman, D. J. (2019). Posterior parietal cortex plays a causal role in perceptual and categorical decisions. *Science*, 365(6449):180–185. Publisher: American Association for the Advancement of Science Section: Report.

Supplementary Information

“Mice alternate between discrete strategies during perceptual decision-making”

Ashwood *et al*, 2021

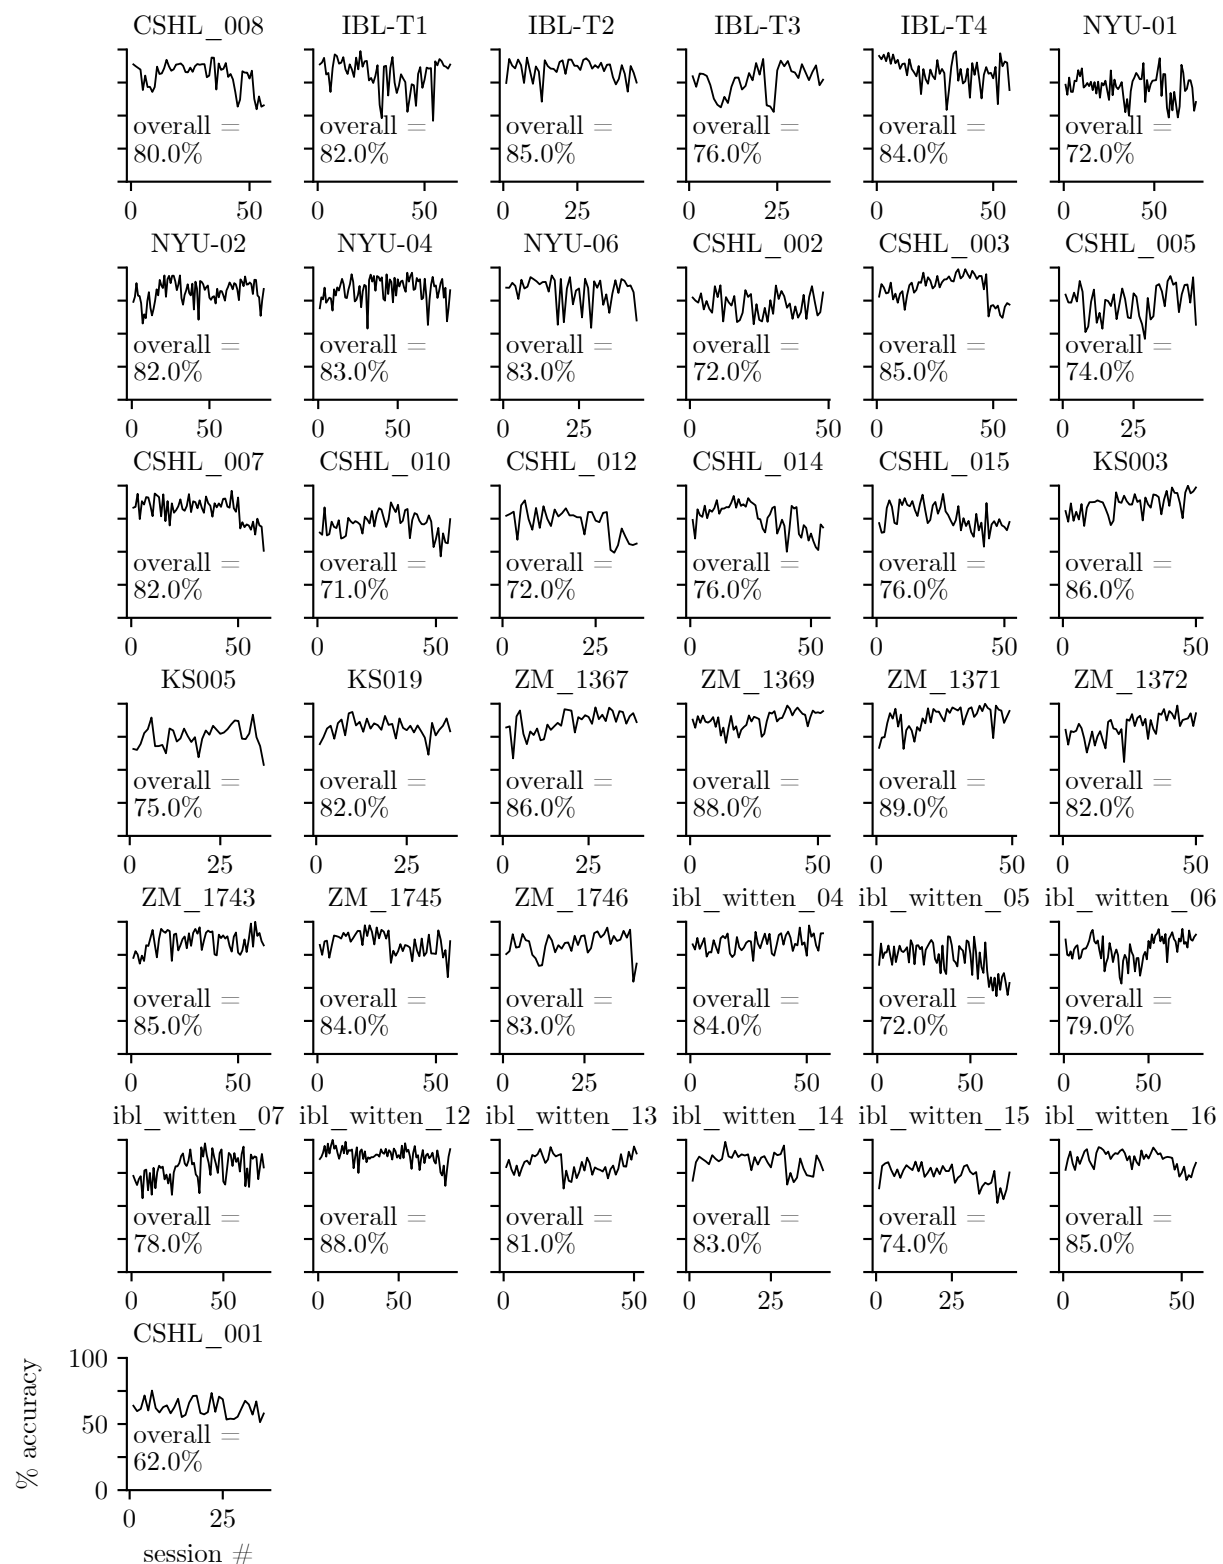


Figure S1: **Raw behavioral data for IBL animals: Accuracy across sessions.** We plot the accuracy of IBL animals [35] across sessions as evidence that mice have learned the task, and that their choice behavior has reached stationarity. We also report the overall accuracy of each animal when aggregated across sessions. The example animal studied in Fig. 2 and Fig. 3 is 'CSHL_008'.

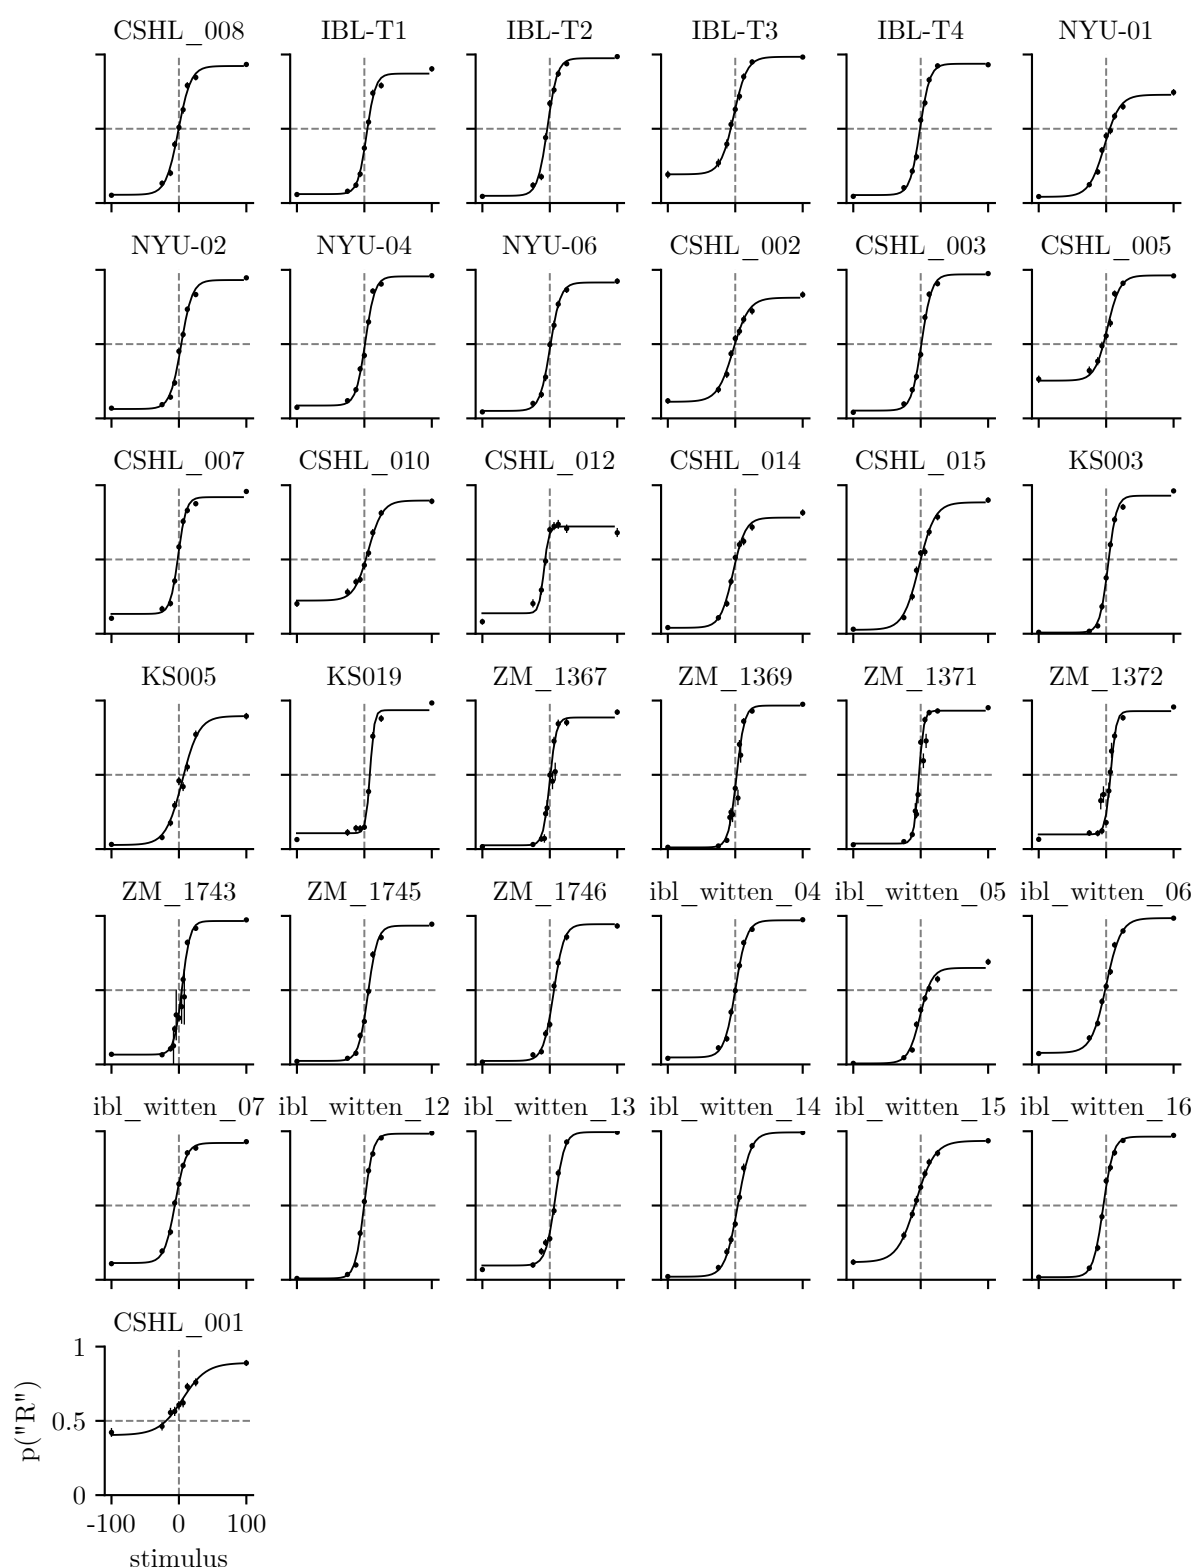


Figure S2: **Raw behavioral data for IBL animals: Psychometric curves.** We plot the psychometric curves for the 37 IBL animals [35] whose choice data we study. The example animal that we study in Fig. 2 and Fig. 3 is 'CSHL_008'. Animals are ordered in the same way that they are in Fig. S1 when considering row-major order, so plots can be compared across the two figures. In each plot, we also show each animal's empirical choice probabilities; error bars are 68% confidence intervals.

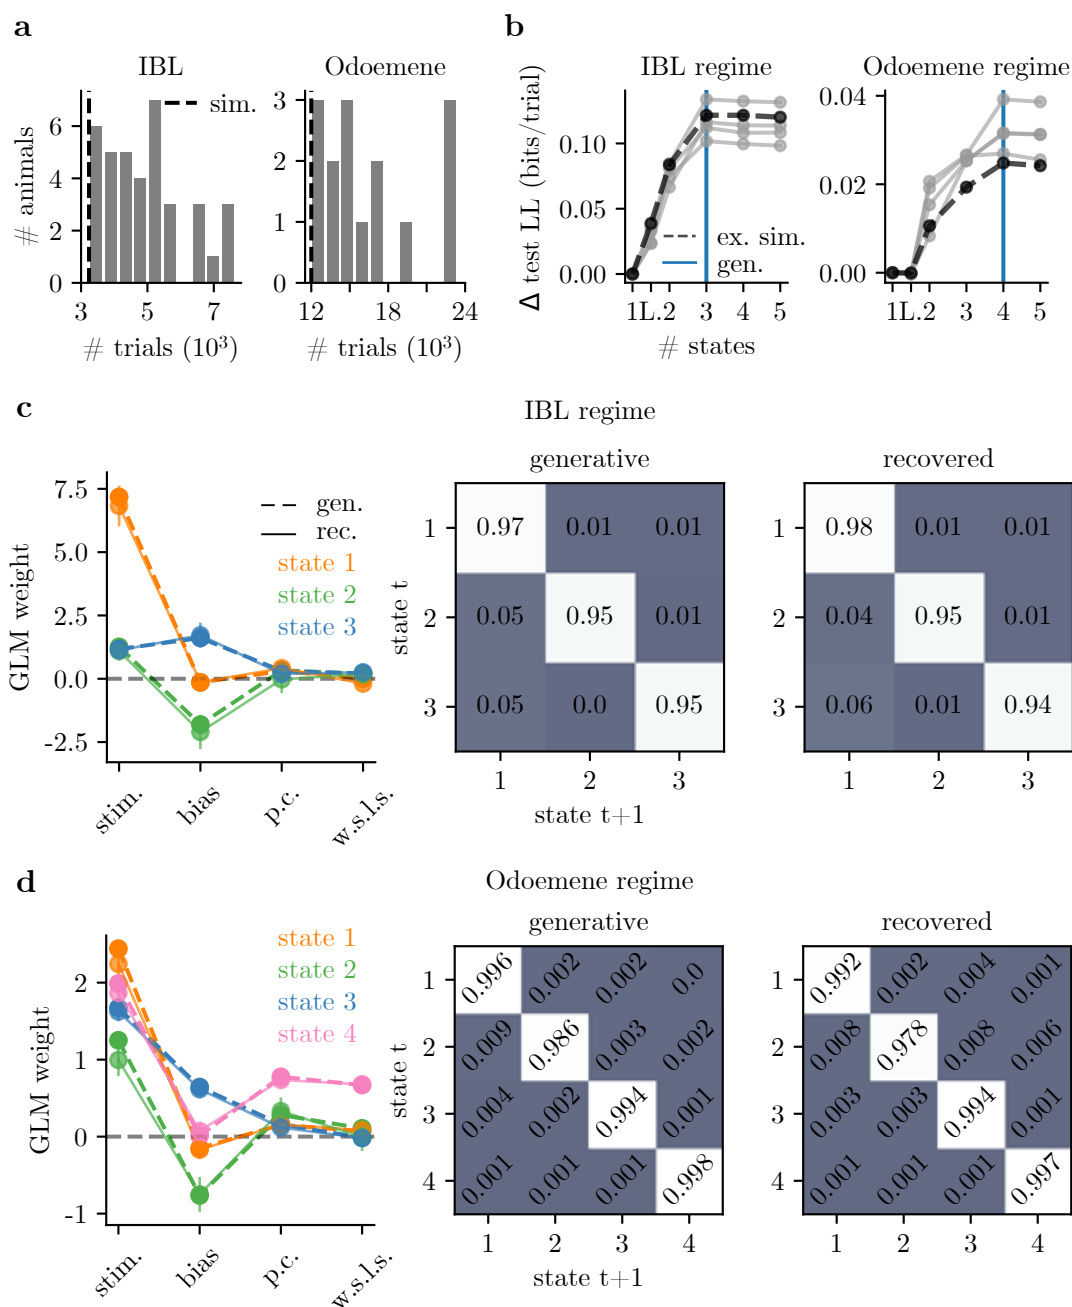


Figure S3: GLM-HMM Recovery Analysis 1. For dataset sizes comparable to those of real animals, we can recover the IBL and Odoemene global parameters in simulated data. **(a)** Dataset sizes for each of the 37 IBL animals studied (left) and 15 mice from Odoemene et al. (right). The dashed vertical line indicates the number of trials that we used in simulation data in panels b, c and d (3240 for the IBL parameter regime and 12000 for the Odoemene regime simulation). **(b)** Normalized test loglikelihood for each of 5 simulations is maximized at 3 states (blue vertical line) after we simulate according to the (IBL regime) parameters shown in panel c. Similarly, in the right panel, normalized test loglikelihood is maximized at 4 states when we simulate choice data with the (Odoemene regime) parameters shown in panel d. The thick black line marked as ‘ex. sim.’ (example simulation) indicates the simulation whose generative and recovered parameters we show in panels c and d. **(c)** Left: the generative and recovered GLM weights (for the simulation marked as ‘ex. sim.’ in panel b) when we simulate choice data in the IBL parameter regime. Middle and right: the generative and recovered transition matrices. **(d)** The generative and recovered parameters in the Odoemene et al. parameter regime.

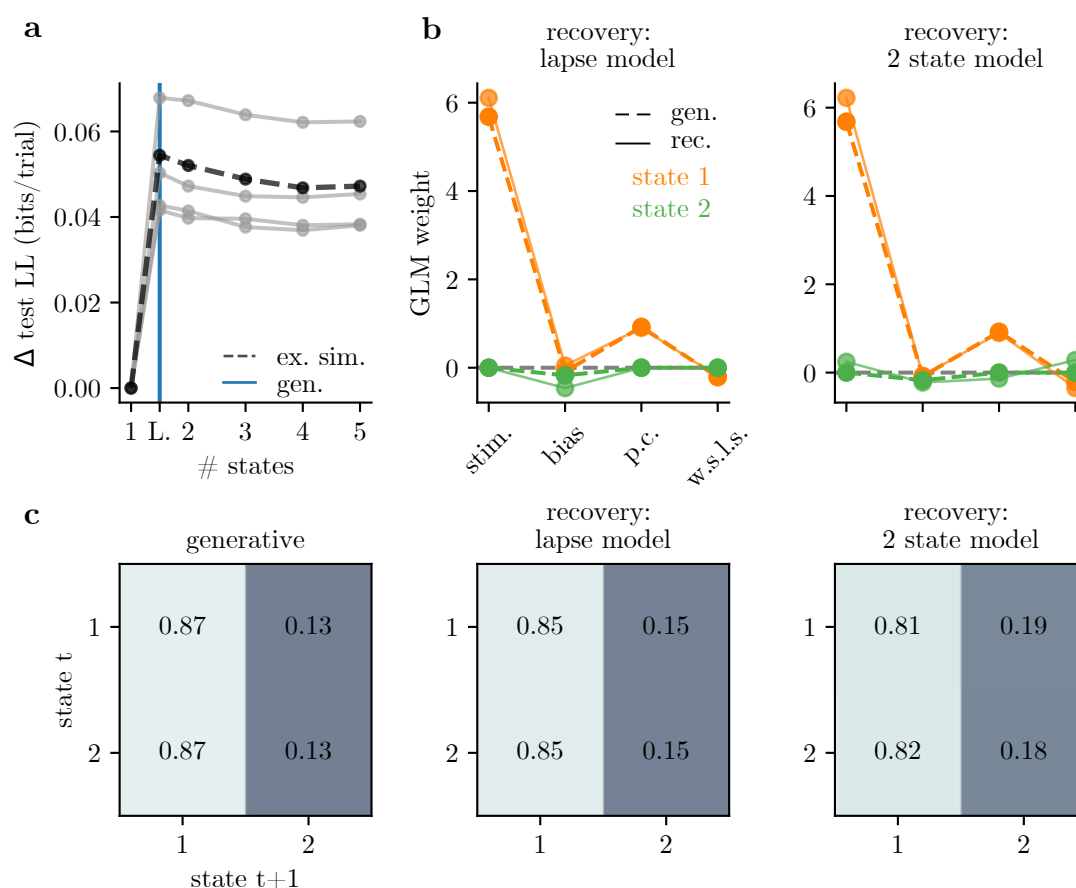


Figure S4: GLM-HMM Recovery Analysis 2: We can recover lapse behavior. **(a)** We simulate 5 datasets, each with 3240 trials, according to the best fitting lapse model for IBL animals. We then fit these simulated datasets with GLM-HMMs, as well as a lapse model (a constrained 2 state GLM-HMM). The normalized test loglikelihood is highest for the lapse model in all simulations, indicating that lapse behavior can be distinguished from the long-enduring multi-state behavior that best described the real data. The thick black line marked as 'ex. sim.' (example simulation) indicates the simulation whose generative and recovered parameters we show in panels b and c. **(b)** Left: the generative and recovered weights when recovery is with a lapse model. Right: the generative weights are the same as in the left panel, but we now recover with an unconstrained 2 state GLM-HMM (thus the stimulus, previous choice and w.s.l.s. weights for the second state can be non-zero) **(c)** The generative (left) transition matrix and the recovered transition matrices when we recover with a lapse model (middle) and an unconstrained 2 state GLM-HMM (right). While the lapse model and 2 state GLM-HMM results don't perfectly agree, if mice were truly lapsing, the transition matrix would not have the large entries on the diagonals that we observe in the real data.




Figure S5: **Model comparison for all IBL animals** Rather than plotting curves on top of one another as in Fig. 4a, we plot each individual animal's normalized test set loglikelihood in a grid. We do not show model comparison results for the example animal studied in Fig. 2 and Fig. 3 since we show the curve for that particular animal there, but animals are otherwise ordered in the same way (according to row-major order) that they are in Fig. S1 and Fig. S2, so plots can be compared across these figures.




Figure S6: **Retrieved transition matrices for all IBL animals** We do not plot the transition matrix for the example animal studied in Fig. 2 as the retrieved transition matrix is shown there. Animals are ordered in the same way as in other supplemental figures so plots can be compared across figures.




Figure S7: **Average posterior state probabilities for all IBL animals.** Average posterior state probabilities across all sessions for each individual animal. Each session lasts 90 trials. Animals are ordered in the same way as in other supplemental figures so plots can be compared across figures.

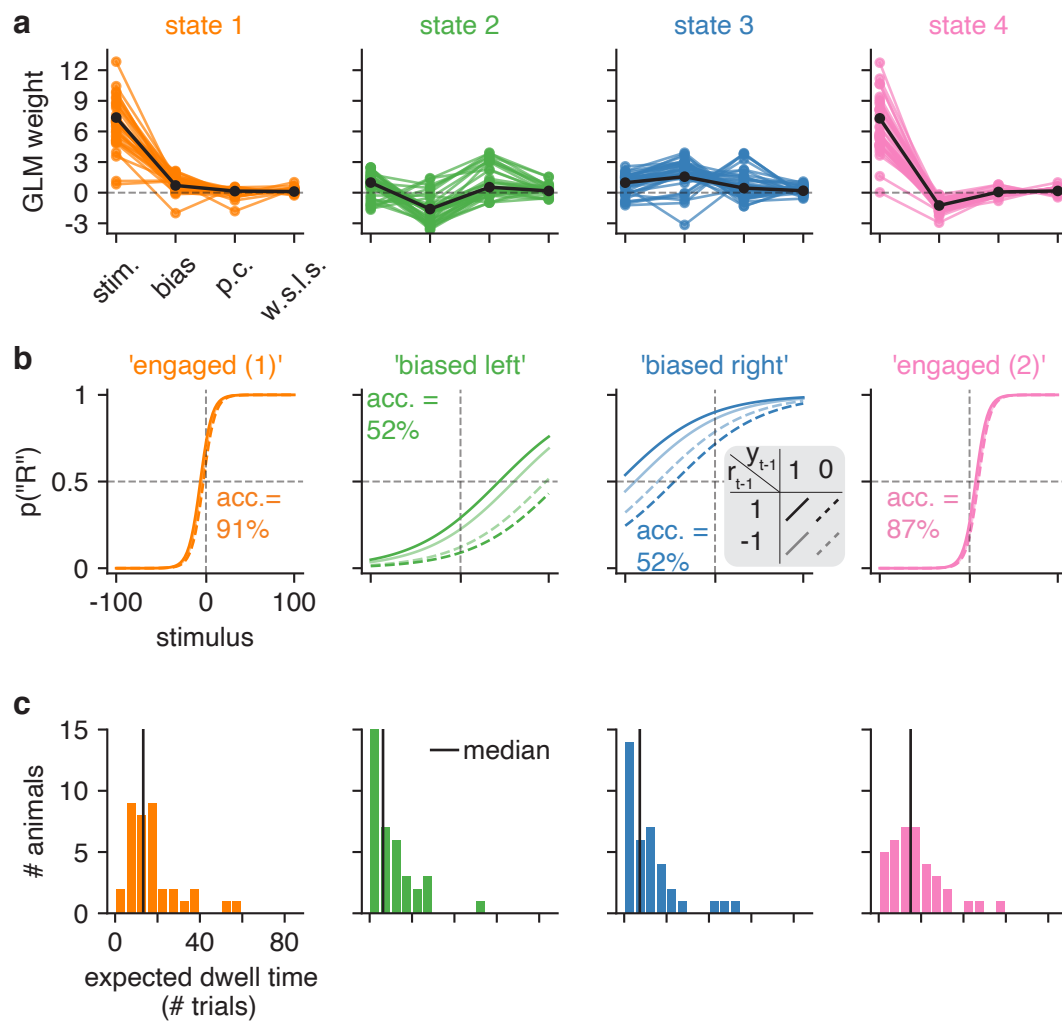


Figure S8: **GLM-HMM 4 state fits to IBL data.** (a) Retrieved GLM weights for each of the 37 IBL animals for each state of the 4 state GLM-HMM, as well as the global weights (black) for the fit when all data from all animals are concatenated. (b) An alternative method of plotting the weights in panel a: we plot, as a function of the stimulus intensity, as well as the animal's reward and choice on the previous trial, the probability that the animal goes rightward at the current trial. (c) The expected dwell time for each animal in each of the 4 states, as obtained from the best-fitting transition matrices for each animal.

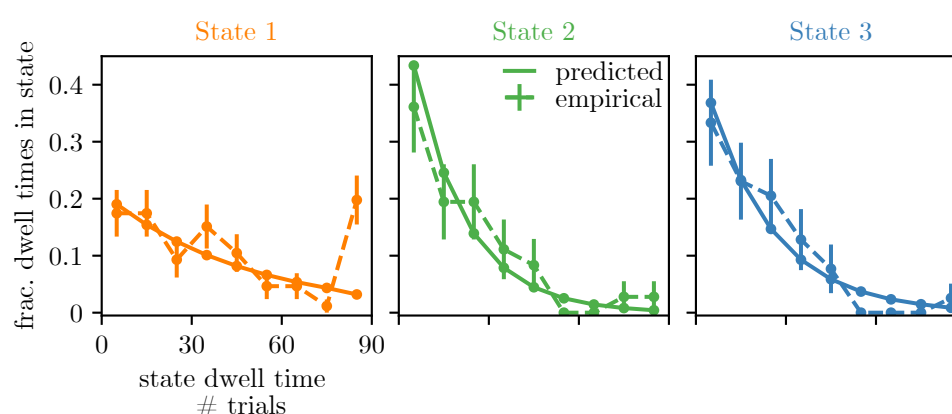


Figure S9: **Retrieved state dwell times are approximately geometrically distributed.** With the solid line, we show the predicted dwell times (according to the retrieved transition matrix) in each of the three states for the example animal of Fig. 2 and Fig. 3. Predicted dwell times can be obtained from the transition matrix as $p(\text{dwell time} = t) = (1 - A_{kk})A_{kk}^{t-1}$ because state dwell times in the Hidden Markov Model are geometrically distributed. We then use the posterior state probabilities to assign states to trials in order to calculate the dwell times that are actually observed in the real data (shown with the dashed line); we also show the 68% confidence intervals associated with these empirical probabilities. We find that the empirical dwell times for the biased leftward and rightward states seem to be geometrically distributed. For the engaged state, because there are some entire sessions (each session is 90 trials) during which the animal is engaged, we see that the empirical dwell times associated with this state are not as well described by a geometric distribution. A future direction may be to allow non-geometrically distributed state dwell times by replacing the Hidden Markov Model with e.g. the Hidden semi-Markov Model [68].

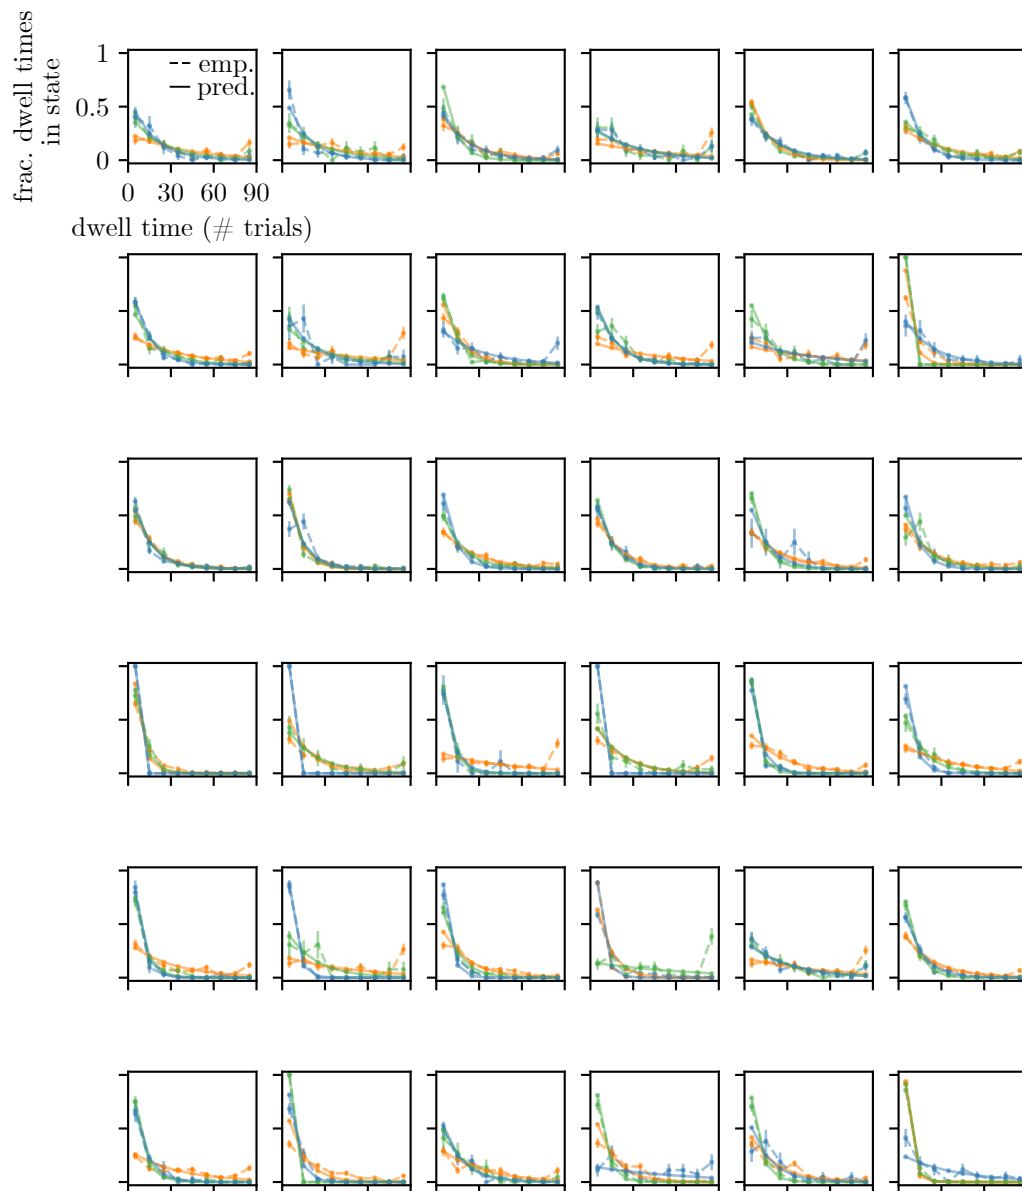


Figure S10: **State dwell times are approximately geometrically distributed for all IBL animals.** We plot the predicted (“pred.”; shown with the solid line) and empirical (“emp.”; shown with the dashed lines) state dwell time probabilities for each IBL animal and each state (excluding the example animal since the state dwell time probabilities for the example animal are shown in Fig. S9)). Colors map to states in the usual way (orange is state 1, green is state 2, blue is state 3). The empirical state dwell times are obtained by using the posterior state probabilities to assign state labels to trials, and the predicted state dwell time probabilities are obtained from the transition matrix according to $p(\text{dwell time} = t) = (1 - A_{kk})A_{kk}^{t-1}$.

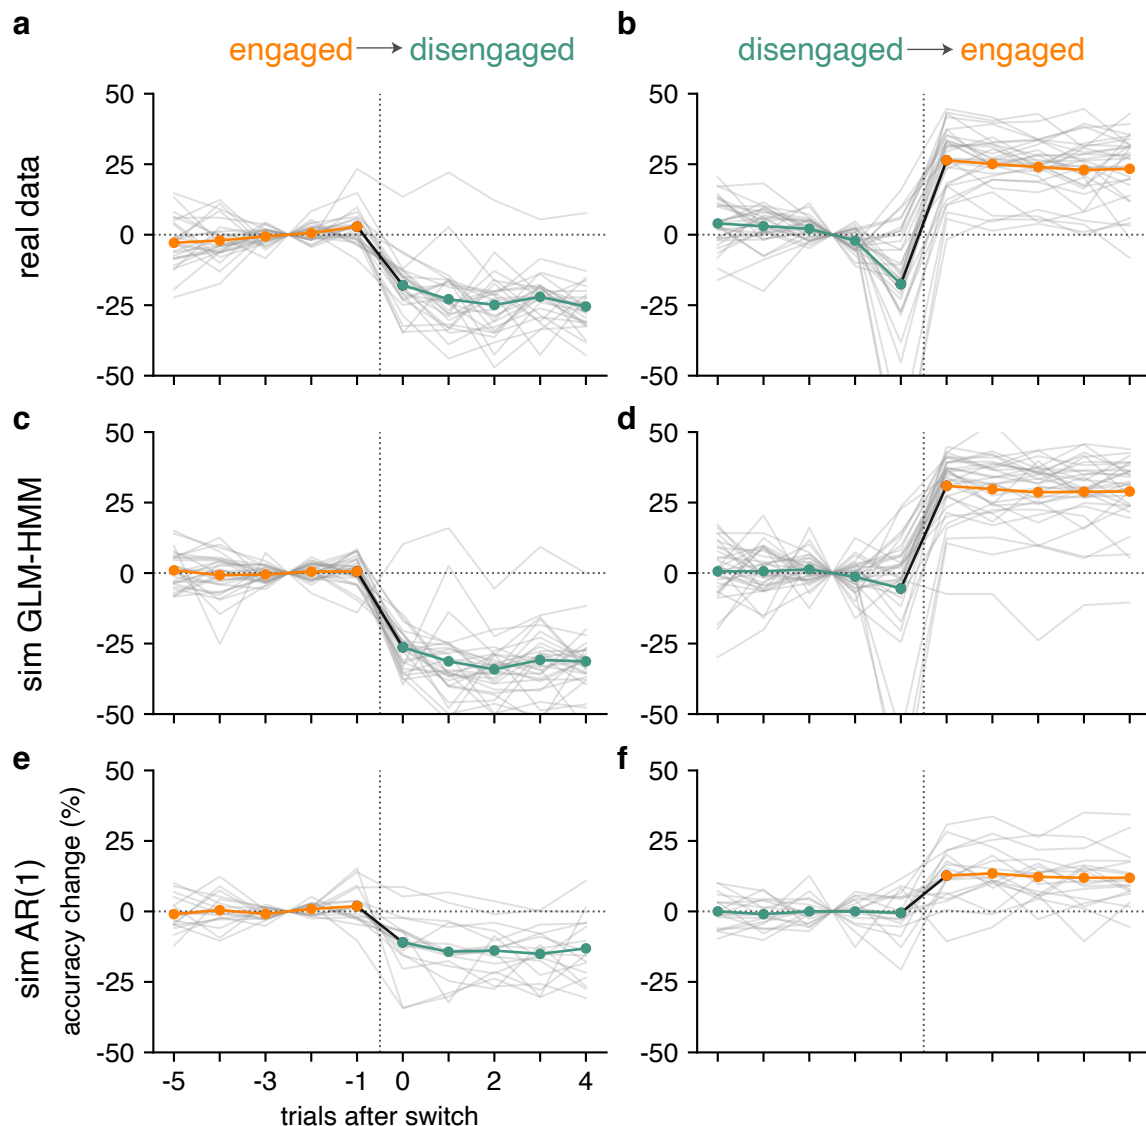


Figure S11: Evidence for discrete decision-making states in raw behavioral data. (a), (b) Average accuracy across 10 trials before and after a state change at trial 5 out of the engaged state (panel a) or to the engaged state (panel b). Each grey line is the average accuracy relative to the mean accuracy for the three trials before the state change for a single animal, while the colored line represents the average across all animals. For further details, see Methods 4.3. (c), (d) Choice data is simulated from a GLM-HMM using the best fitting parameters for each mouse (each grey line represents choice data simulated with a different animal's 3 state GLM-HMM fit). We then identified state changes out of and to the engaged state, and show the change in accuracy before and after these state transitions. (e), (f) We simulated choice data using weights that smoothly evolve according to an AR(1) model as described in Section 4.3 using the Psytrack fits for each animal. We then fit a 3 state GLM-HMM to the simulated choice data, and identified the state with the largest weight on the stimulus as the engaged state. We show the change in accuracy for each animal's simulated choice data (each grey line represents data simulated from a different animal's PsyTrack model fit) when switching out of the engaged state (panel e) and when switching to the engaged state (panel f). By comparing the plots in c and d to those in e and f, and observing that the change in accuracy is, on average, 27% when switching out of the engaged state for data simulated from the 3 state GLM-HMM compared to just 13% for data simulated from the PsyTrack-inspired model, we have verified that we can use changes in accuracy across the state change boundary to disambiguate behavior based on discrete decision-making strategies from behavior based on continuously evolving strategies. Furthermore, the plots in a and b resemble the plots in c and d more closely than the plots in e and f providing additional evidence that real mice use discrete decision-making strategies.

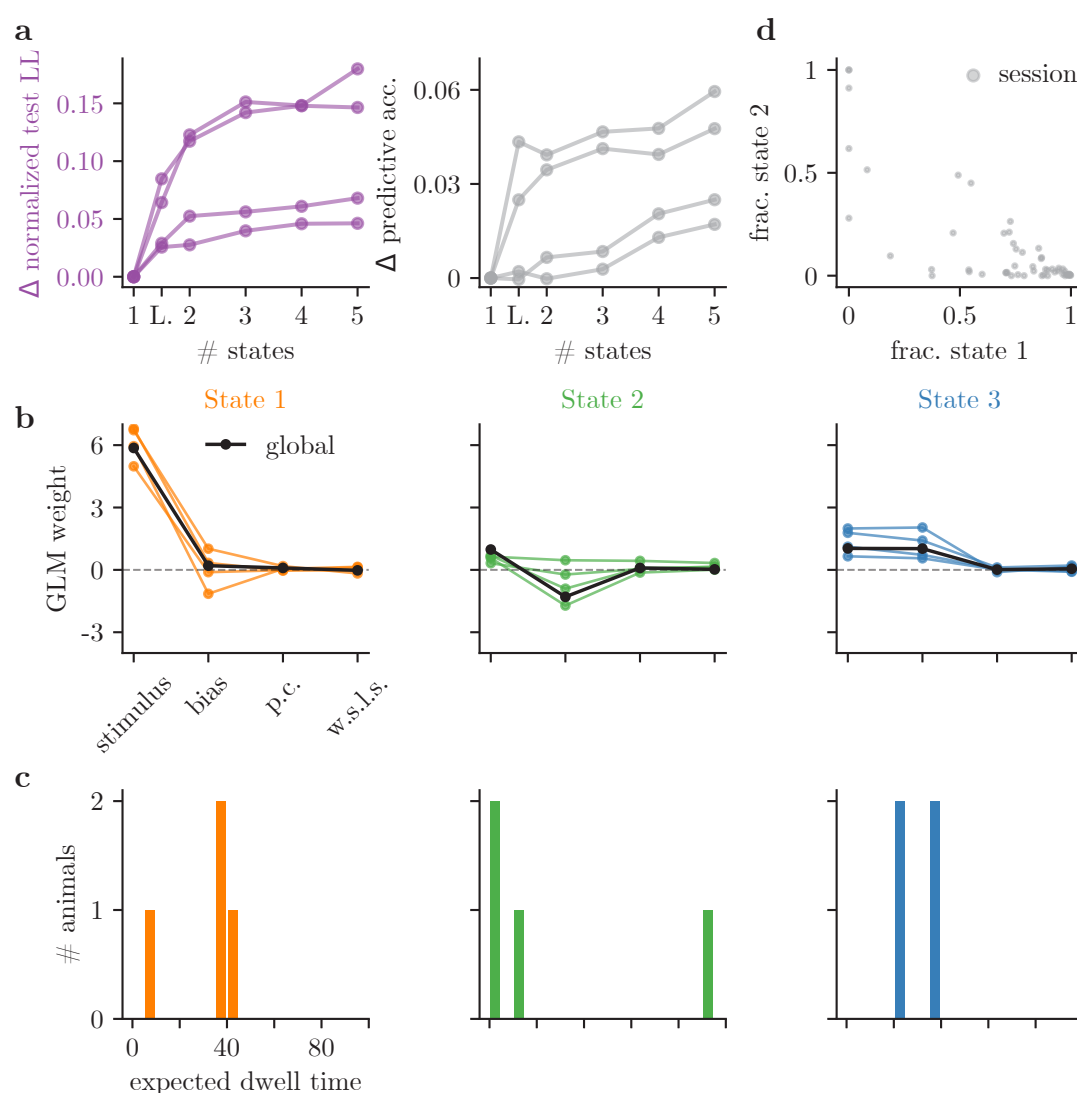


Figure S12: **GLM-HMM application to 4 mice not exposed to bias blocks in IBL task.** We confirm that mice that have never been exposed to bias blocks in the IBL task continue to show state-dependent decision-making. This is a sister figure to Fig. 4, and each panel can be interpreted in the same way as in Fig. 4.




Figure S13: **Accuracy across sessions for Odoemene et al. [48] animals.** We plot the accuracy across sessions for Odoemene et al. animals as evidence that the animals' choice behavior has reached stationarity.

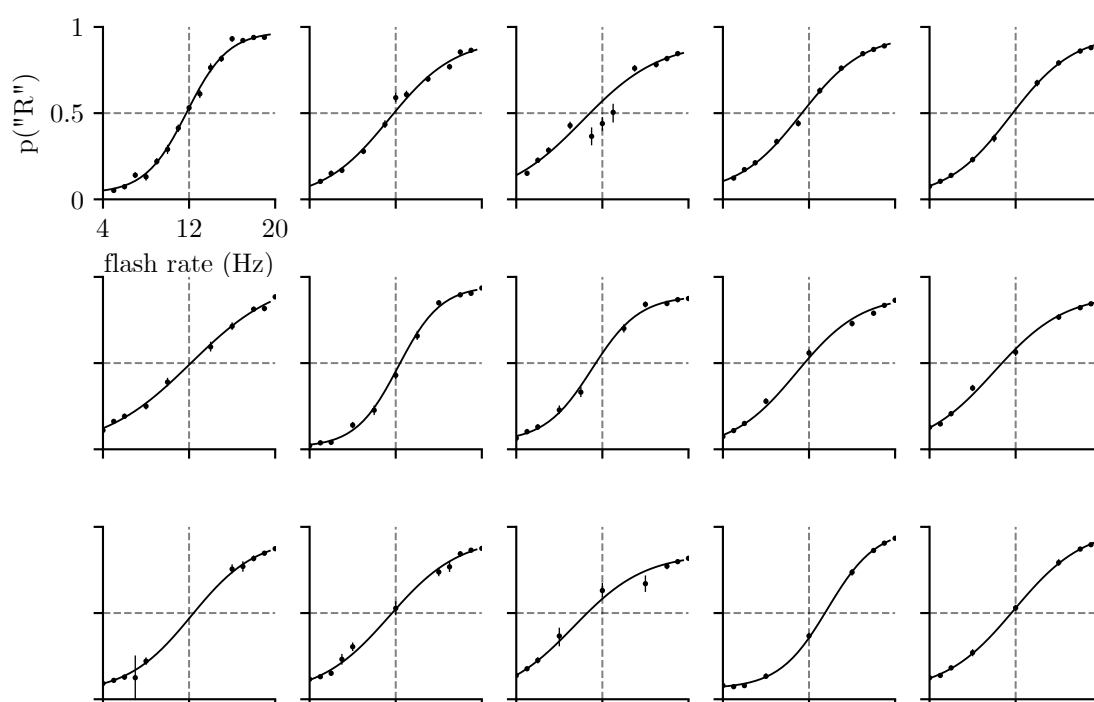


Figure S14: **Psychometric curves for Odoemene et al. [48] animals.** We plot the psychometric curves for each of the 15 animals whose choice data we study. We also show each animal's empirical choice probabilities; error bars are 68% confidence intervals. Animals are ordered as in Fig. S13.

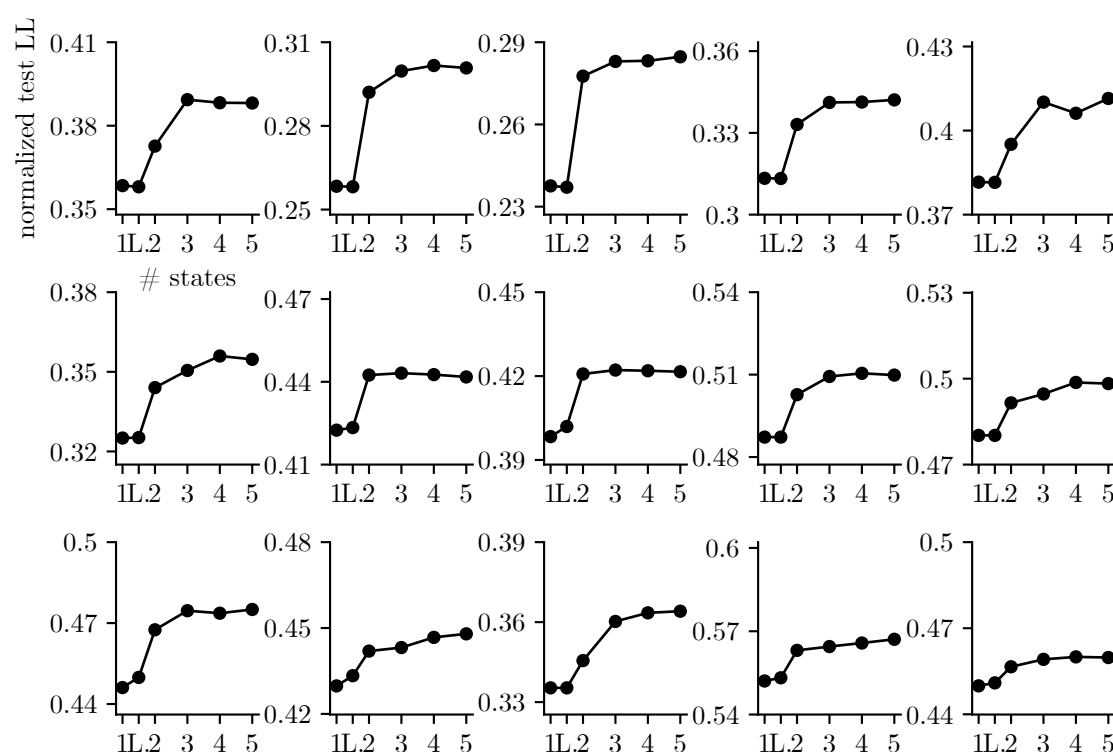


Figure S15: **Model Comparison for all Odoemene et al. animals.** Rather than plotting curves on top of one another as in Fig. 7b, we plot each individual animal's normalized test set loglikelihood in a grid. Animals are ordered in the same way that they are in Fig. S13 and Fig. S14, so plots can be compared across figures.

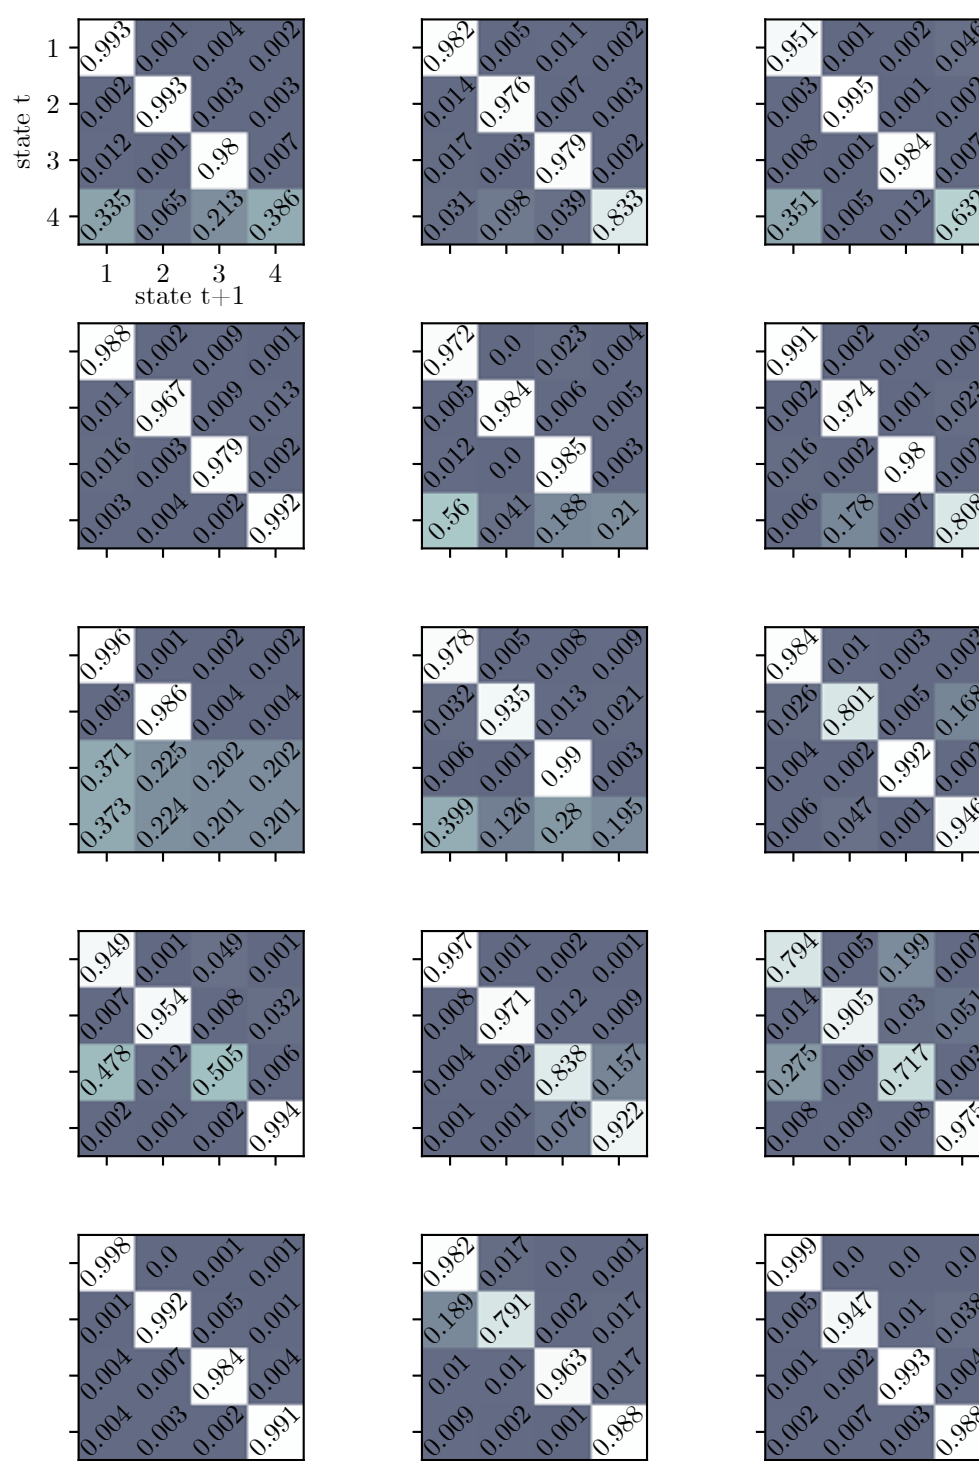


Figure S16: **Retrieved transition matrices for Odoemene et al. animals** Each individual transition matrix is the best fitting transition matrix for 1 of the 15 Odoemene et al. mice that we study. Row-major order has animals ordered in the same way that they are in Fig. S15.

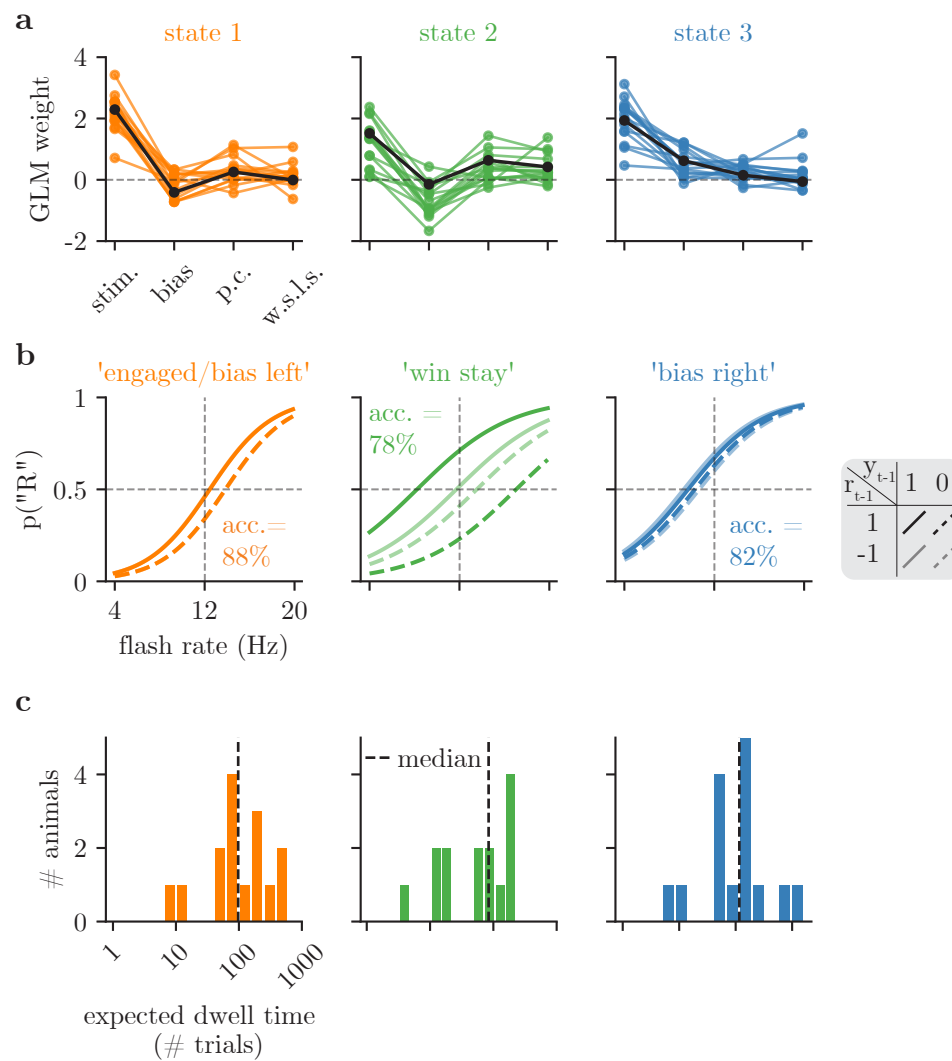


Figure S17: **GLM-HMM 3 state fits to Odoemene et al. data.** When 3 state GLM-HMMs are fit to the Odoemene et al. data, the engaged and bias left states are merged to form a single (mostly engaged) state (accuracy remains high at 88%), while the win-stay and bias right states are largely unchanged. This is a sister figure to panels d, e, and f of Fig. 7, where the panels can be interpreted in the same way that they are there.

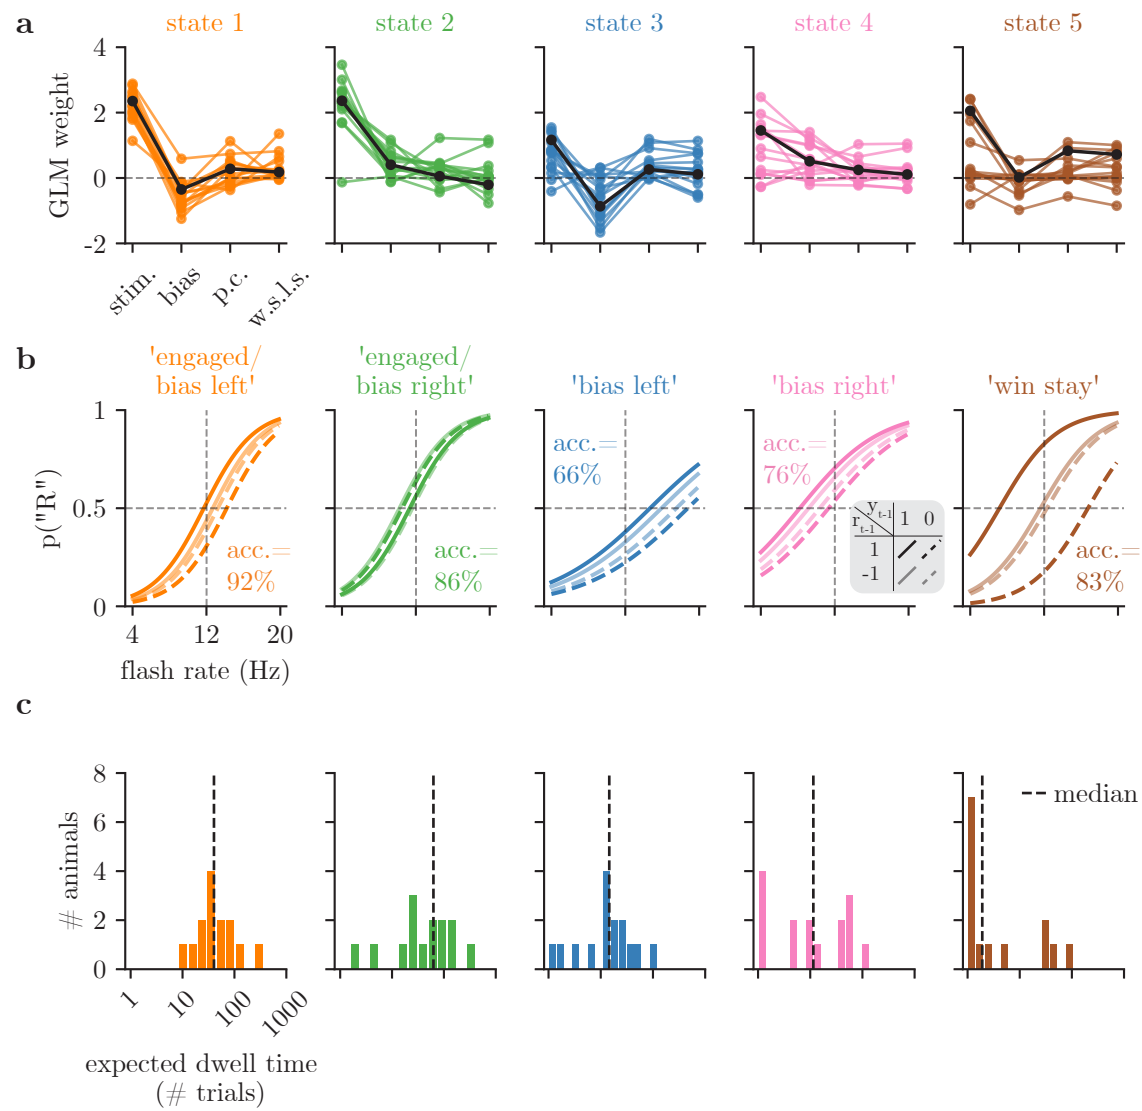


Figure S18: **GLM-HMM 5 state fits to Odoemene et al. data.** When 5 state GLM-HMMs are fit to the Odoemene et al. data, the engaged state is split into an engaged/bias left and engaged/bias right state, while the win-stay, bias right and bias left states are largely unchanged. This is a sister figure to panels d, e, and f of Fig. 7, where the panels can be interpreted in the same way that they are there.




Figure S19: **GLM-HMM captures ‘warm-up’ effect for Odoemene et al. animals.** Average (across 20 sessions) posterior state probabilities for the first 200 trials of a session for each animal in the Odoemene et al. dataset. Orange corresponds to the engaged state, green to the biased left, blue to the biased right and pink to the win-stay state from Fig. 7. Error bars represent standard errors.




Figure S20: **GLM-HMM posterior state probabilities at end of session.** Average (across 20 sessions) posterior state probabilities for the last 200 trials of a session for each animal in the Odoemene et al. dataset. Orange corresponds to the engaged state, green to the biased left, blue to the biased right and pink to the win-stay state from Fig. 7. Error bars represent standard errors.