Supplementary figure 1

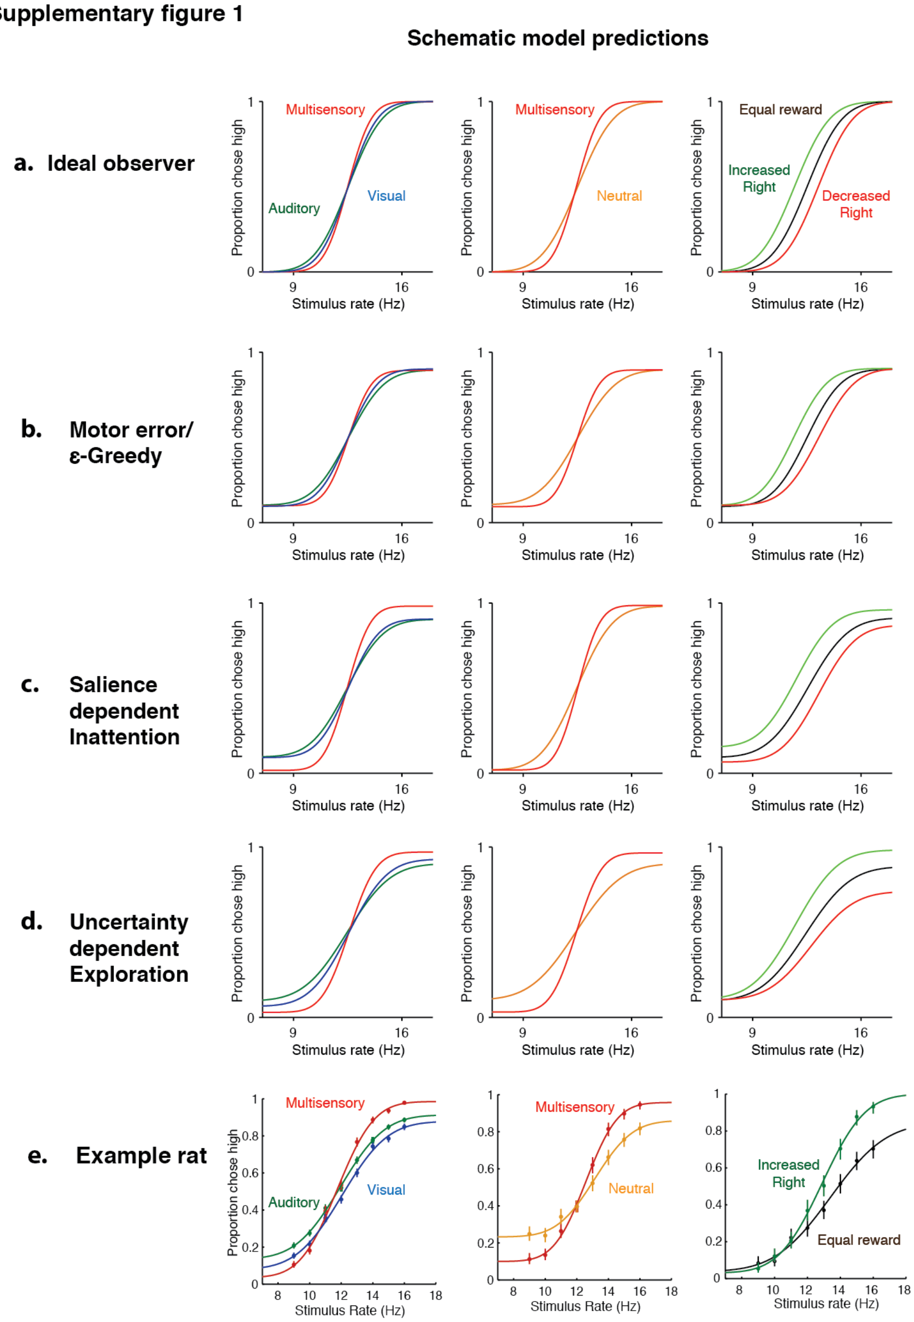

Supplementary Figure 1: Uncertainty-dependent exploration is the only model that accounts for behavioral data from all three manipulations Columns: data/predictions for three experimental manipulations. Left: unisensory vs. multisensory. Middle: matched vs. neutral. Right: Asymmetric reward. a-d: Four candidate models. (a) Ideal observer model predicts no lapses and only changes in sensitivity/bias across conditions. (b) Fixed motor error model predicts a constant rate of lapses across conditions in addition to changes in sensitivity/bias predicted from the ideal observer. (c) Inattention model predicts that the overall lapse rate (sum of lapses on both sides) depends on the level of bottom-up attentional salience, allowing for different rates for unisensory and multisensory. It also predicts that the lapse rate on neutral trials should be equal to that on multisensory trials, and that manipulating rightward reward should affect both lapse rates. (d) Uncertainty-dependent exploration model predicts that overall lapse rate depends on the level of exploratoriness and hence uncertainty associated with that condition, allowing for different lapse rates on unisensory and multisensory trials. It also predicts that the lapse rate on neutral trials should be equal to that on auditory trials and manipulating rightward reward should only affect high rate lapses. (e) Data from an example rat on all three manipulations.

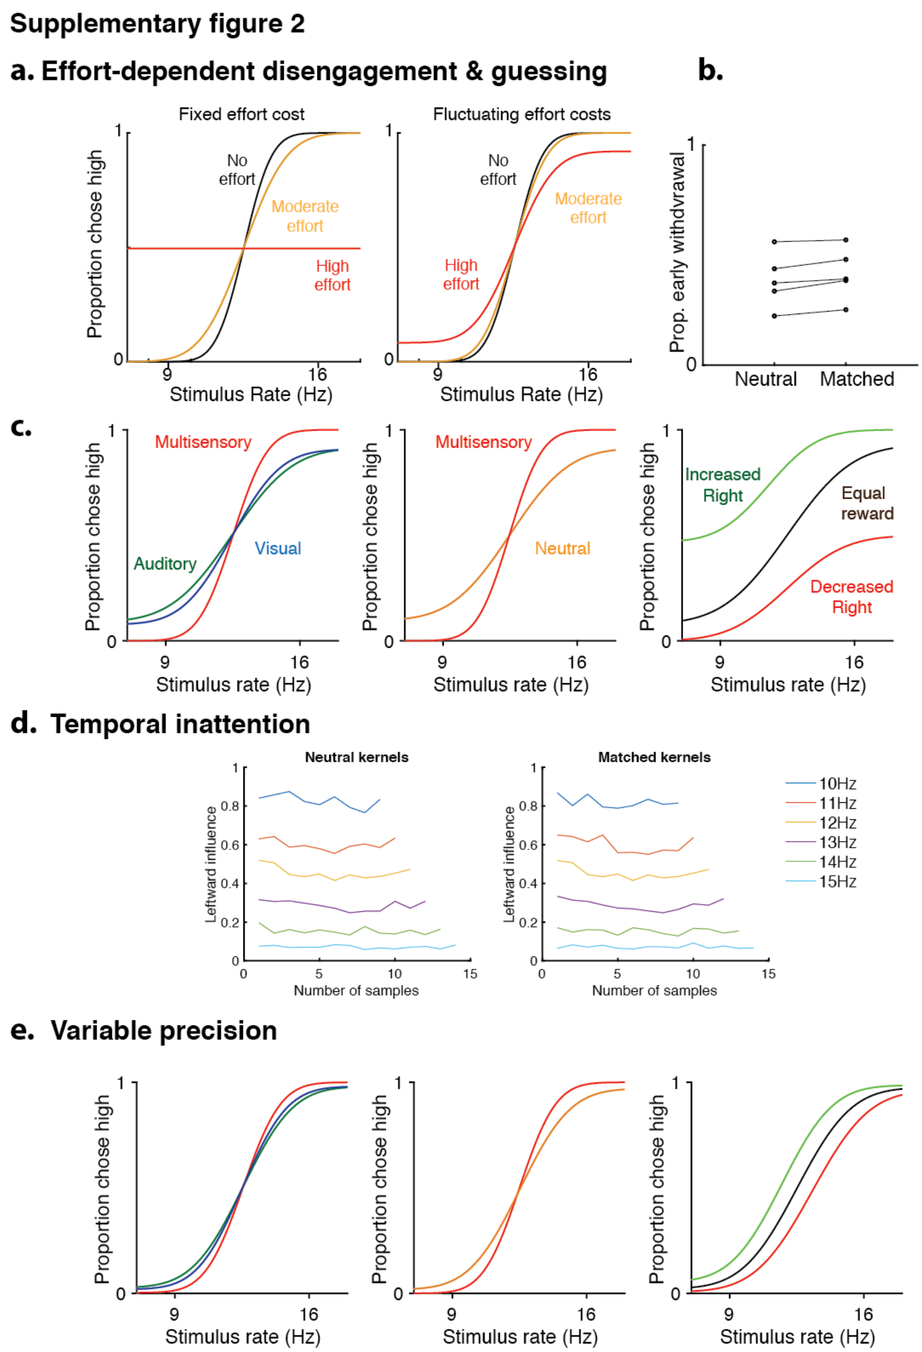

Supplementary Figure 2: Alternative models of inattentional lapses. Predictions of alternative models of lapses. (a) Effort-dependent disengagement model: In this model, there is an additional cost or mental effort to being engaged in the task which could vary with condition, and an additional random guessing action. If the net payoff of engagement is not greater than the average value of a guess, then it guesses randomly. Such a model does not produce lapses if the effort is fixed across trials (left), but could produce lapses if the effort fluctuates from trial to trial (center). (b) Proportion of trials on which the animal withdrew prematurely doesnt vary between matched and neutral trials, suggesting that rats are not disengaging preferentially on neutral trials. (c) Predictions of the effort-dependent disengagement model. The model accurately predicts increased lapses on unisensory trials (left panel, green/blue traces) and neutral multisensory trials (middle panel, orange trace). However, for asymmetric reward manipulations (right), the model fails to predict our behavioral observation (Fig. 4d) that only lapses on the manipulated side are affected. (d) Temporal inattention model: in this model, temporal weighting of evidence differs between matched and neutral trials. To test this, we compared psychophysical kernels on matched and neutral trials. The temporal dynamics of attention are unchanged between the two kinds of trials, arguing against the temporal inattention model. (e) Variable precision model: in this model, the sensory noise (or its inverse, precision) fluctuates from trial to trial. The model accurately predicts increased lapses on unisensory trials (left panel, green/blue traces) and neutral multisensory trials (middle panel, orange trace). However, for asymmetric reward manipulations (right), the model fails to predict our behavioral observation (Fig. 4d) that lapses only on the manipulated side are affected. Like other models of inattention, it predicts that manipulating reward on one side should affect both lapses.

Supplementary figure 3

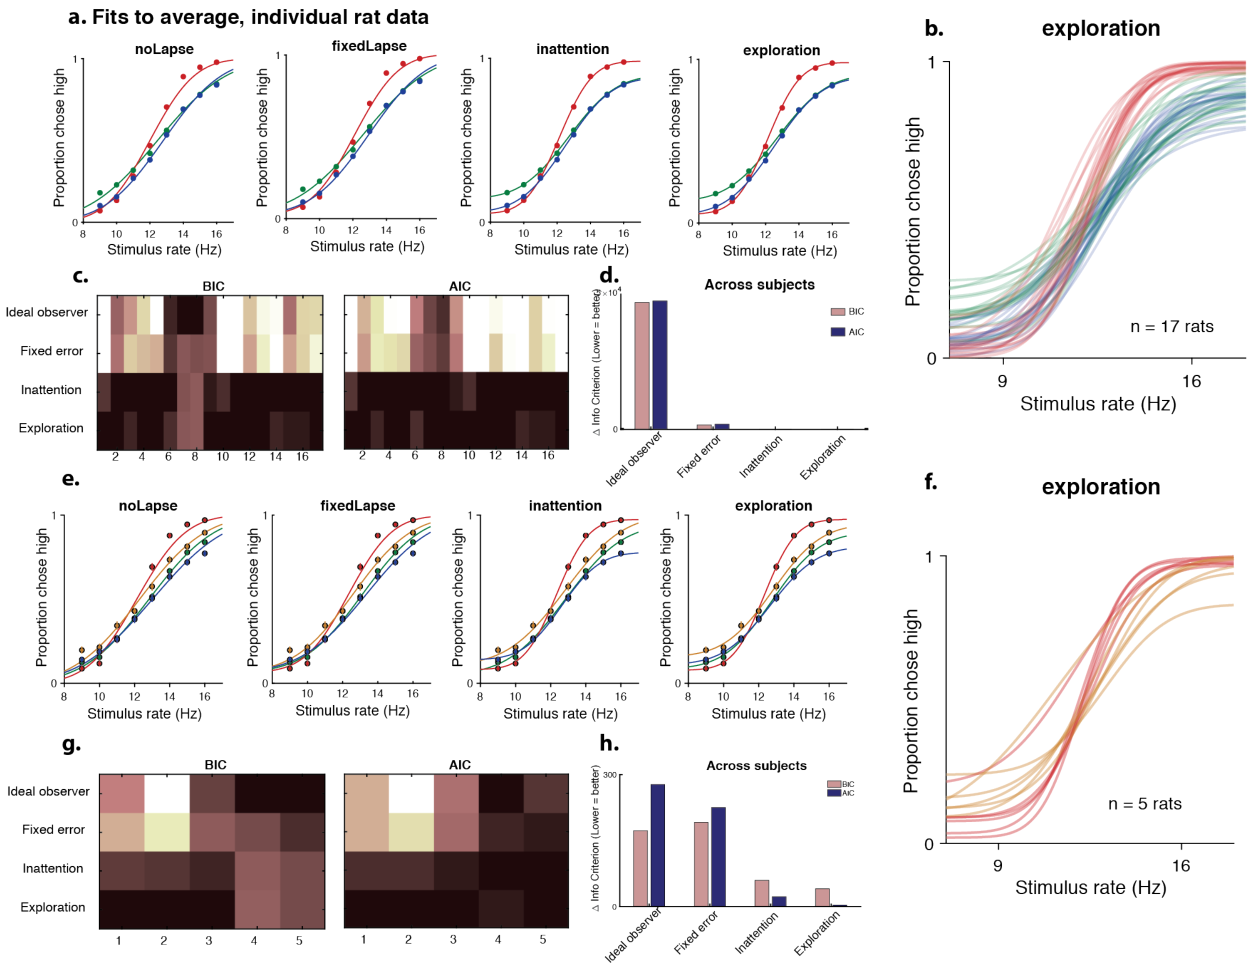

Supplementary Figure 3: Uncertainty guided exploration outperforms competing models for average and individual data (a) Fits of the four models to average rat data on unisensory (blue-visual, green-auditory) and multisensory (red) trials. (b) Exploration model fits to unisensory and multisensory data for 17 individual animals (c) Model comparison for individual animals using BIC (left), AIC (right). Darker colors are lower BICs/AICs, denoting a better fit. (d) Summed model comparison metrics across animals, showing that inattention and exploration models fit the data equally well, and much better than the ideal observer or fixed error models. (e) Fits of the four models to average data including neutral trials (orange) provide a stronger test of the inattention model. (f) Exploration model fits to multisensory data including neutral trials for 5 individual animals (g) Model comparison for individual animals. (h) Summed model comparison metrics across animals shows that the uncertainty-guided exploration model performs better than other models.

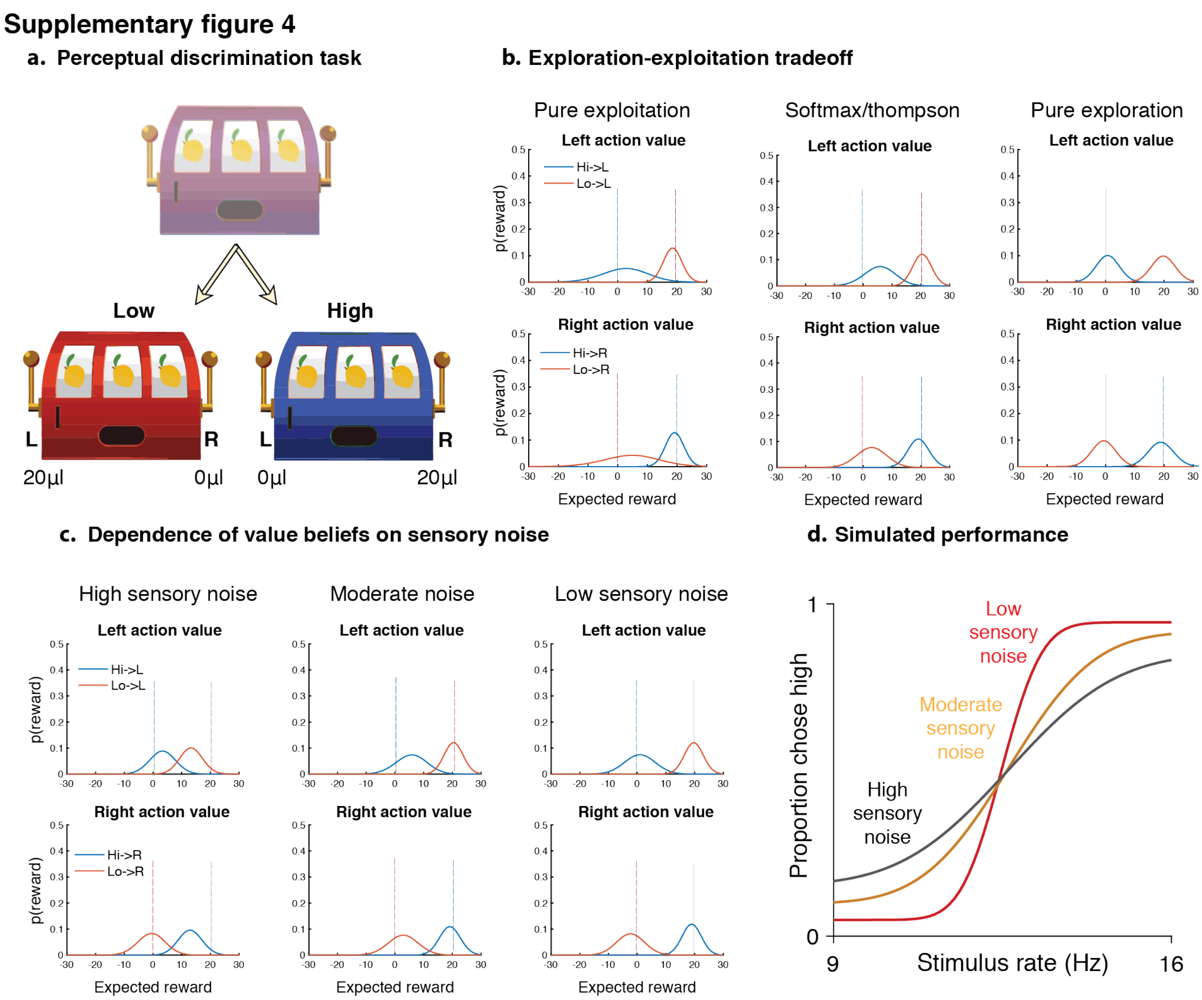

Supplementary Figure 4: Thompson sampling, which balances exploration and exploitation, predicts lapses that increase with perceptual noise Schematic illustrating the explore-exploit tradeoff in perceptual two-alternative tasks. (a) Formulation of perceptual decision making task as a partially observable contextual bandit. To solve this task, an observer needs to infer the true category of the stimulus (Low or High) based on noisy observations, and pick the best action given the true category (Left for Low, Right for High). This requires accurately learning the expected rewards from all 4 state-action pairs. (b) Beliefs about expected reward from different actions (L,R) performed in different states (Hi, Lo) showing different levels of uncertainty depending on policy. Beliefs are updated based on outcomes using a Bayesian update rule that takes into account uncertainty in state estimation. A greedy policy (left) that always picks the best action maximizes reward and learns well about the preferred state-action pairs (i.e. Lo-L and Hi-R) but has high uncertainty about the non-preferred pairs (Lo-R, Hi-L). A random policy (right) earns reward at chance, but learns equally well about all state-action pairs. Thompson sampling (center) implements a softmax decision rule that depends on the current uncertainties in each value, and balances immediately reward-maximizing decisions with decisions that reduce uncertainty, maximizing average reward in the long term. (c) Learnt beliefs about expected reward with Thompson sampling at various levels of perceptual uncertainty. High sensory noise (left) leads to large perceptual uncertainty, yielding highly overlapping belief distributions owing to a reduced ability to assign obtained rewards to one of the states. Lower levels of sensory noise (center, right) produce more separable beliefs. (d) Simulated performance over 2000 trials of the Bayesian learner shown above, under a Thompson sampling policy. As the sensory noise decreases (Black to Yellow to Red), the observer makes fewer exploratory choices owing to the more separable value beliefs, giving rise to lower lapse rates.

Supplementary figure 5

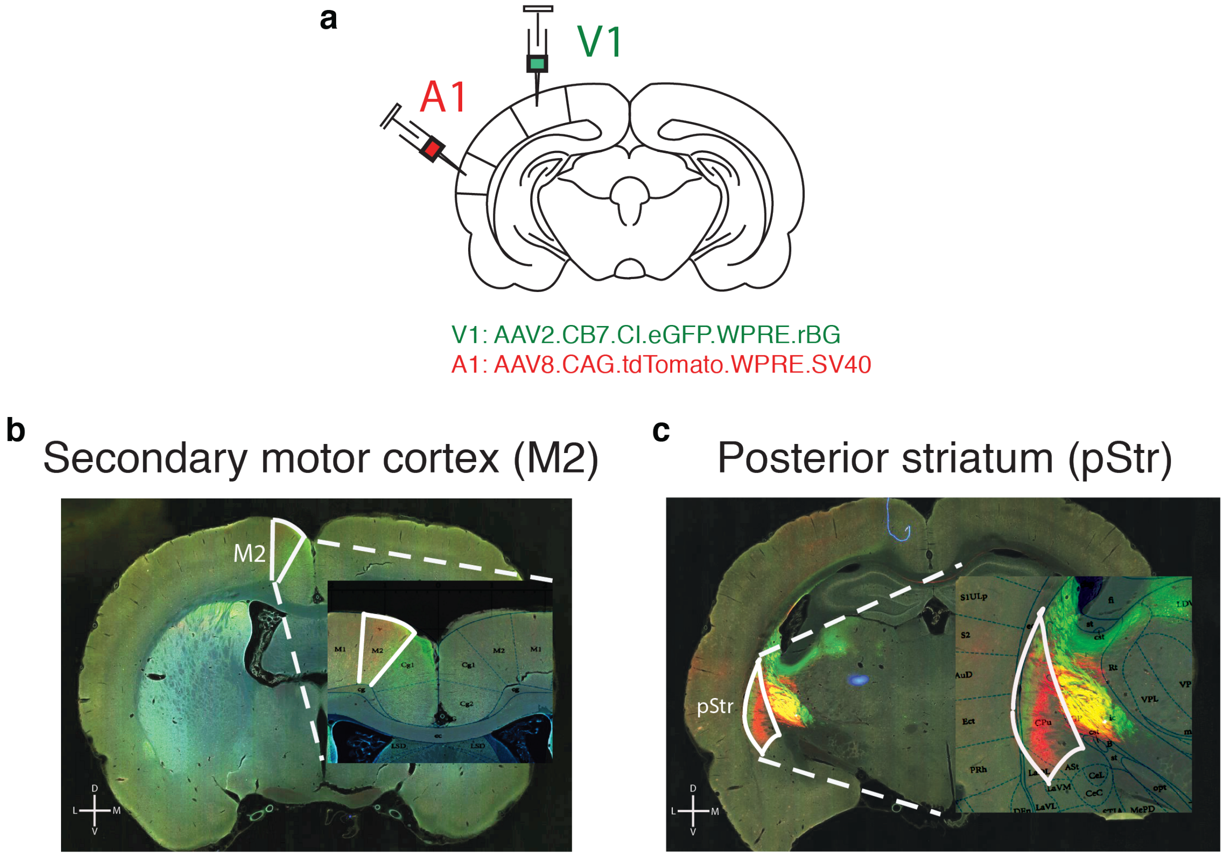

Supplementary Figure 5: pStr and M2 receive direct projections from visual and auditory cortex (a) Schematic of tracing experiments. AAV2.CB7.CI.EGFP.WPRE.RBG and AAV2.CAG.tdTomato.WPRE.SV40 constructs were injected unilaterally to primary visual (V1) and auditory (A1) cortices, respectively (V1 coordinates: 6.9 mm posterior to Bregma; 4.2 mm to the right of midline; A1 coordinates: 4.7 mm posterior to Bregma; 7 mm to the right of midline). (b) Secondary motor cortex (M2) receives inputs from V1 and A1 as shown by green and red fluorescence. (c) Posterior striatum (pStr) receives direct inputs from V1 and A1 as shown by green and red fluorescence. Yellow signal medial to pStr reflects overlapping passing fibers.

Supplementary figure 6

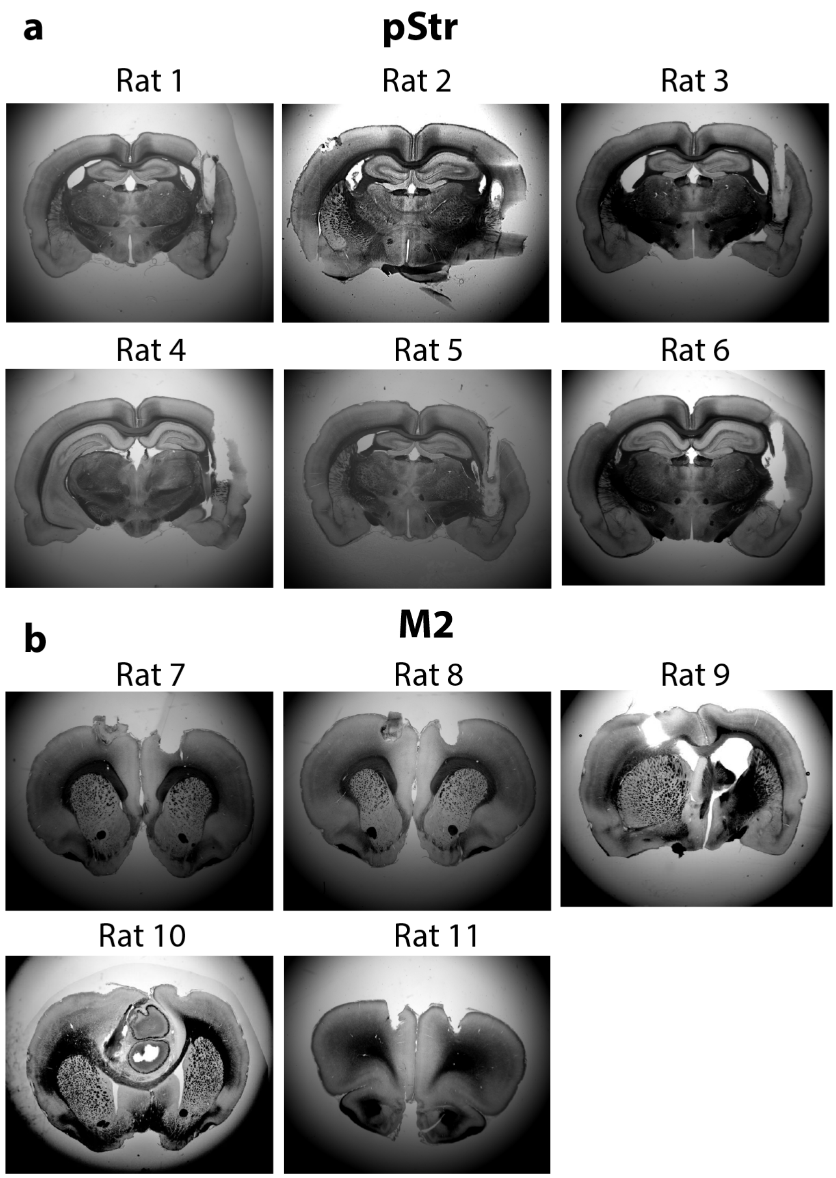

Supplementary Figure 6: Histological slices of implanted rats Representative coronal slices of all rats implanted with cannulae for muscimol inactivation experiments. (a) 6 rats were bilaterally implanted in posterior striatum (pStr). (b) 5 rats were implanted in secondary motor cortex (M2).

Supplementary figure 7

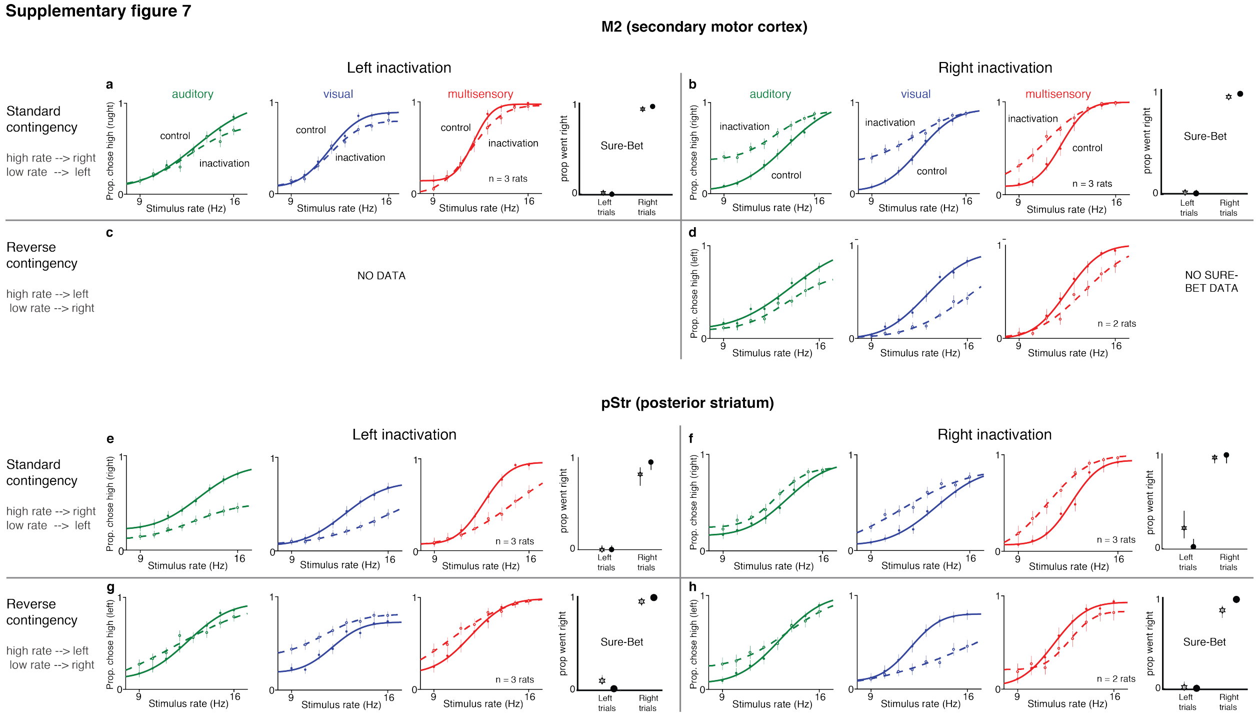

Supplementary Figure 7: Unilateral inactivation of M2 or pStr biases performance ipsilaterally and increases contralateral lapses Performance of the same rats shown in Figure 5b depicted as a function of the inactivated side (right or left) and the rate-contingency in which they were trained (standard or reverse). Standard contingency: high rate = go right, low rate = go left; reverse contingency: high rate = go left, low rate = go right. Each quadrant shows 4 plots: 3 psychometrics for rate discrimination trials and one for performance on sure-bet trials. auditory (green), visual (blue) and multisensory (red). (a)-(d) M2 inactivation. (e)-(h) pStr inactivation. (a), (d) Rats trained on the standard contingency and inactivated on the left hemisphere show increased lapses on the high rates (i.e., fewer rightward choices on high rates). No effect on sure-bet trials. (b), (f) Rats trained on the standard contingency and inactivated on the right hemisphere show increased lapses on the low rates (i.e., fewer leftward choices on low rates). No effect on sure-bet trials. (c), (g) Rats trained on the reverse contingency and inactivated on the left hemisphere show increased lapses on the low rates (i.e., fewer rightward choices on low rates). No effect on sure-bet trials. No data for this condition for M2 inactivation. (d), (h) Rats trained on the reverse contingency and inactivated on the right hemisphere show increased lapses on the high rates (i.e., fewer leftward choices on high rates). No effect on sure-bet trials for pStr inactivated animals; no data for M2 inactivated animals.

Supplementary figure 8

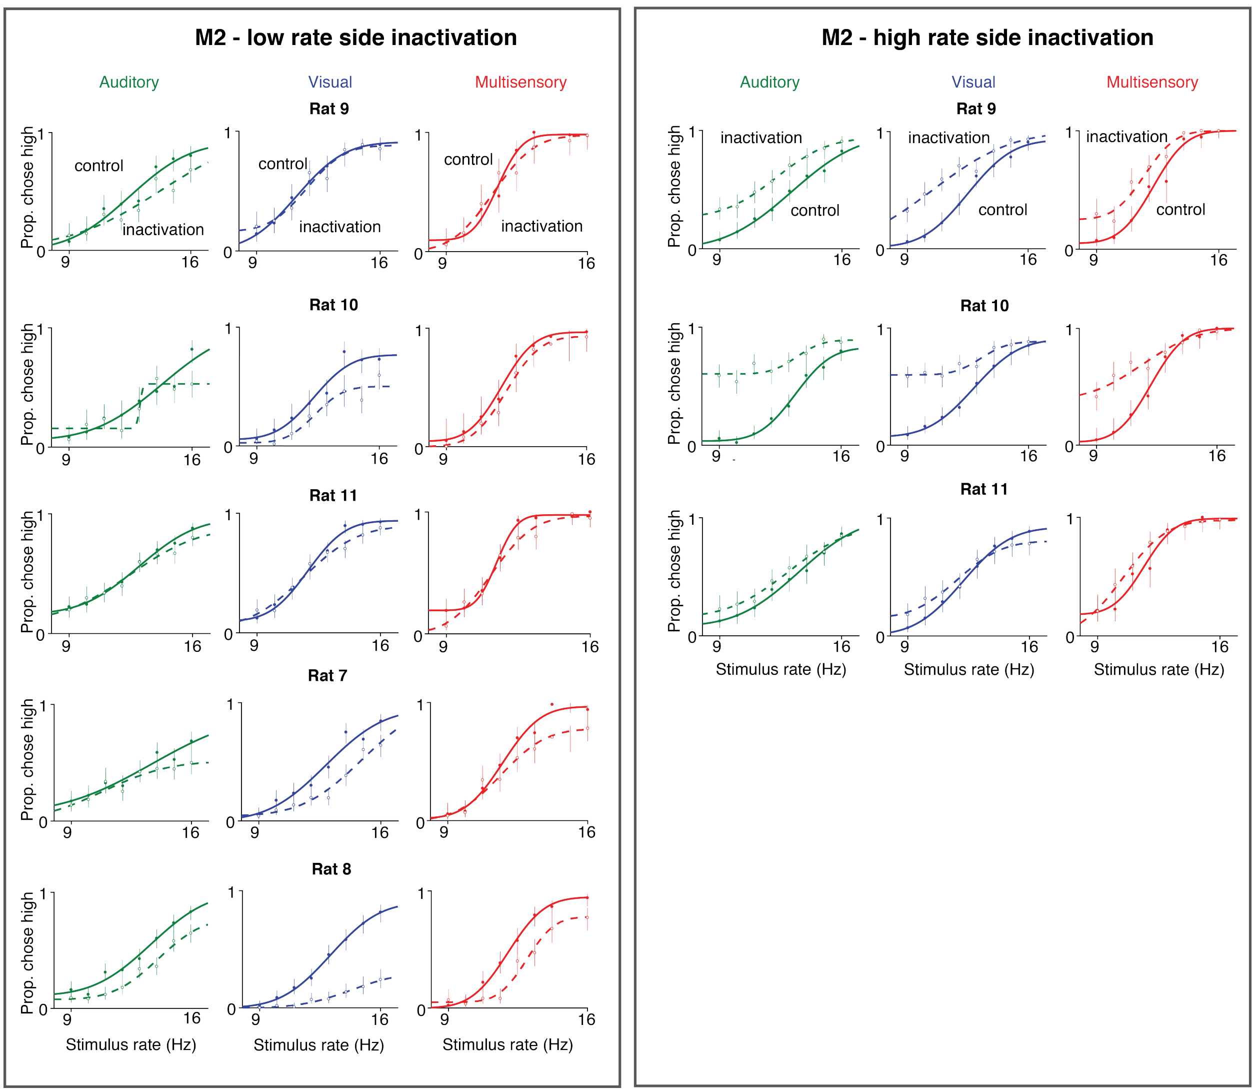

Supplementary Figure 8: Single rat performance following M2 inactivation Left: inactivation of the low-rate associated side. Rat shows increased lapses on high-rate trials on all sensory modalities. Right: inactivation of the high-rate associated side. Rat shows increased lapses on low-rate trials on all sensory modalities. Auditory (green), visual (blue) and multisensory (red).

Supplementary figure 9

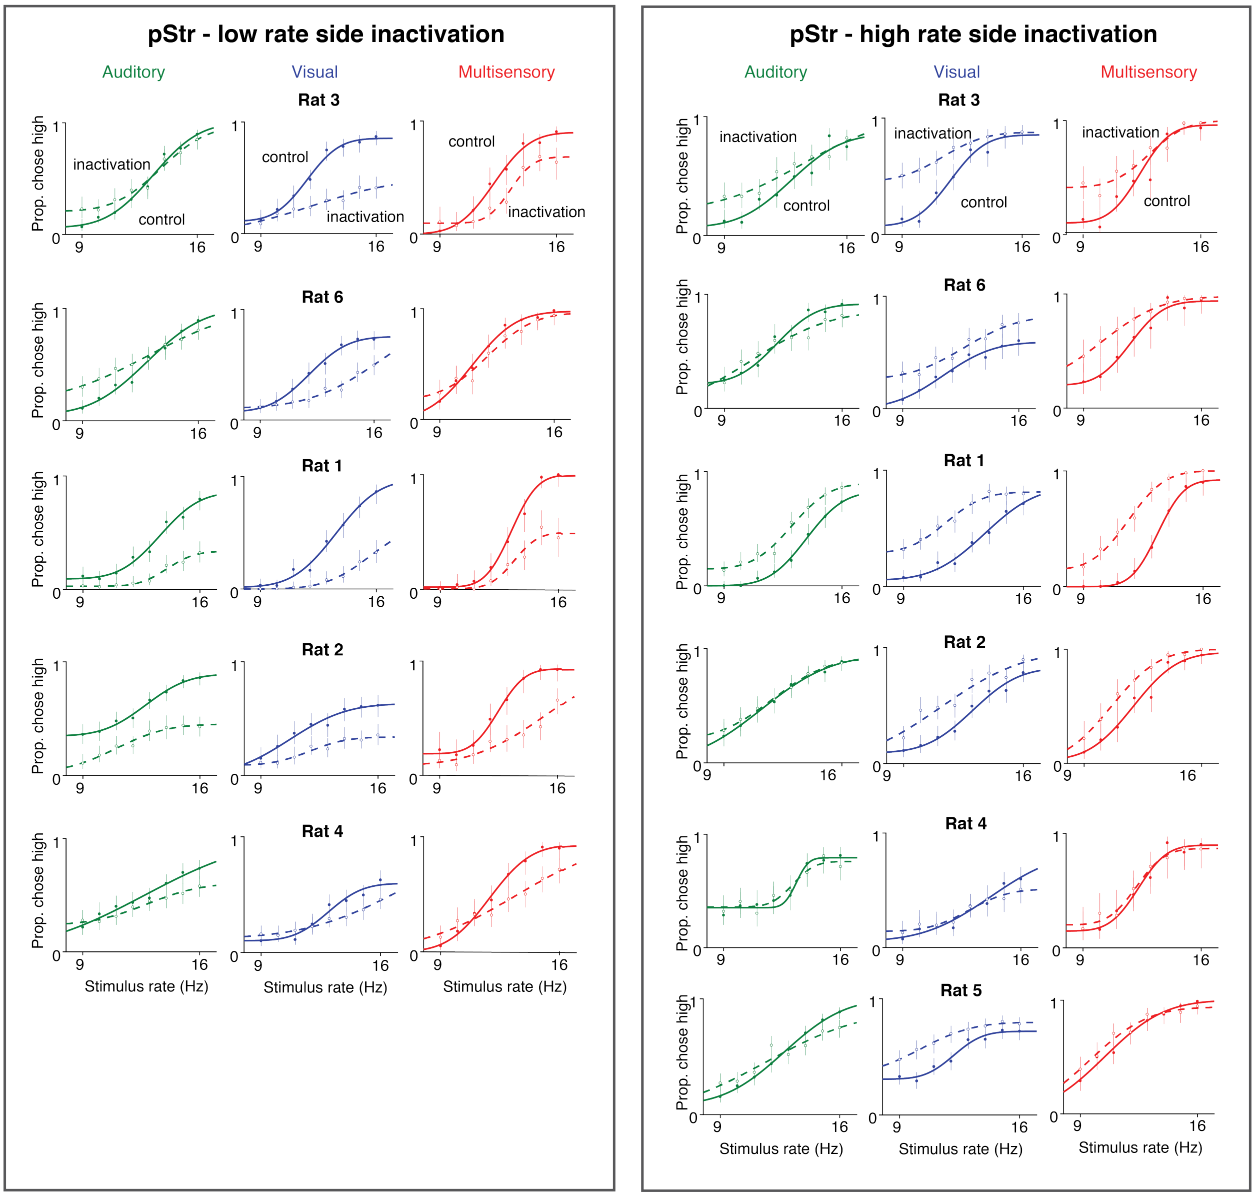

Supplementary Figure 9: Single rat performance following pStr inactivation Left: inactivation of the low-rate associated side. Rat shows increased lapses on high-rate trials on all sensory modalities. Right: inactivation of the high-rate associated side. Rat shows increased lapses on low-rate trials on all sensory modalities. Auditory (green), visual (blue) and multisensory (red).

Supplementary figure 10

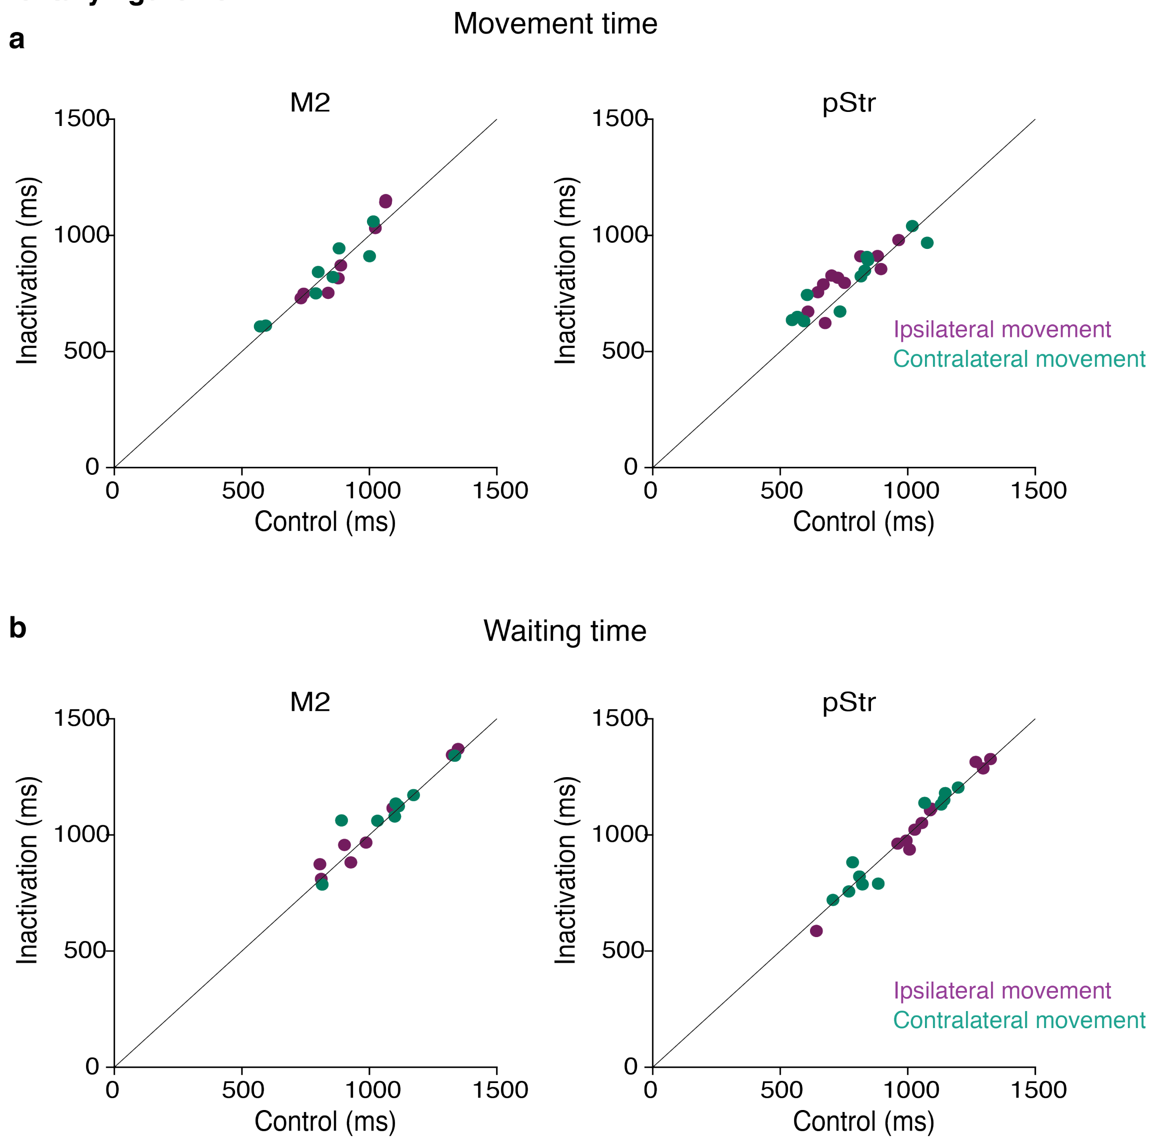

Supplementary Figure 10: No significant effect on movement parameters following muscimol inactivation (a) Mean movement times from the center port to the side ports were not significantly different following muscimol inactivation of M2 (left; p = 0.9554 for contralateral, 0.9852 for ipsilateral movements; n=5 rats) or pStr (right; p = 0.6629 for contra, p =0.2615 for ipsi, n=6 rats). Control data on the abscissa is plotted against inactivation data on the ordinate. Purple, movement toward the side ipsilateral to the inactivation site; blue, movement toward the side contralateral to the inactivation site; Error bars (s.e.m.) are not visible because they were obscured by the markers in all cases. (b) Mean wait times in the center port were not significantly different following muscimol inactivation of M2 (left; p = 0.7612 for contra, p =0.8896 for ipsi, n=5 rats) or pStr (right; p = 0.9128 for contra, p =0.9412 for ipsi, n=6 rats). All p-values were computed from paired t-tests. Error bars (s.e.m.) are not visible because they were obscured by the markers in all cases.

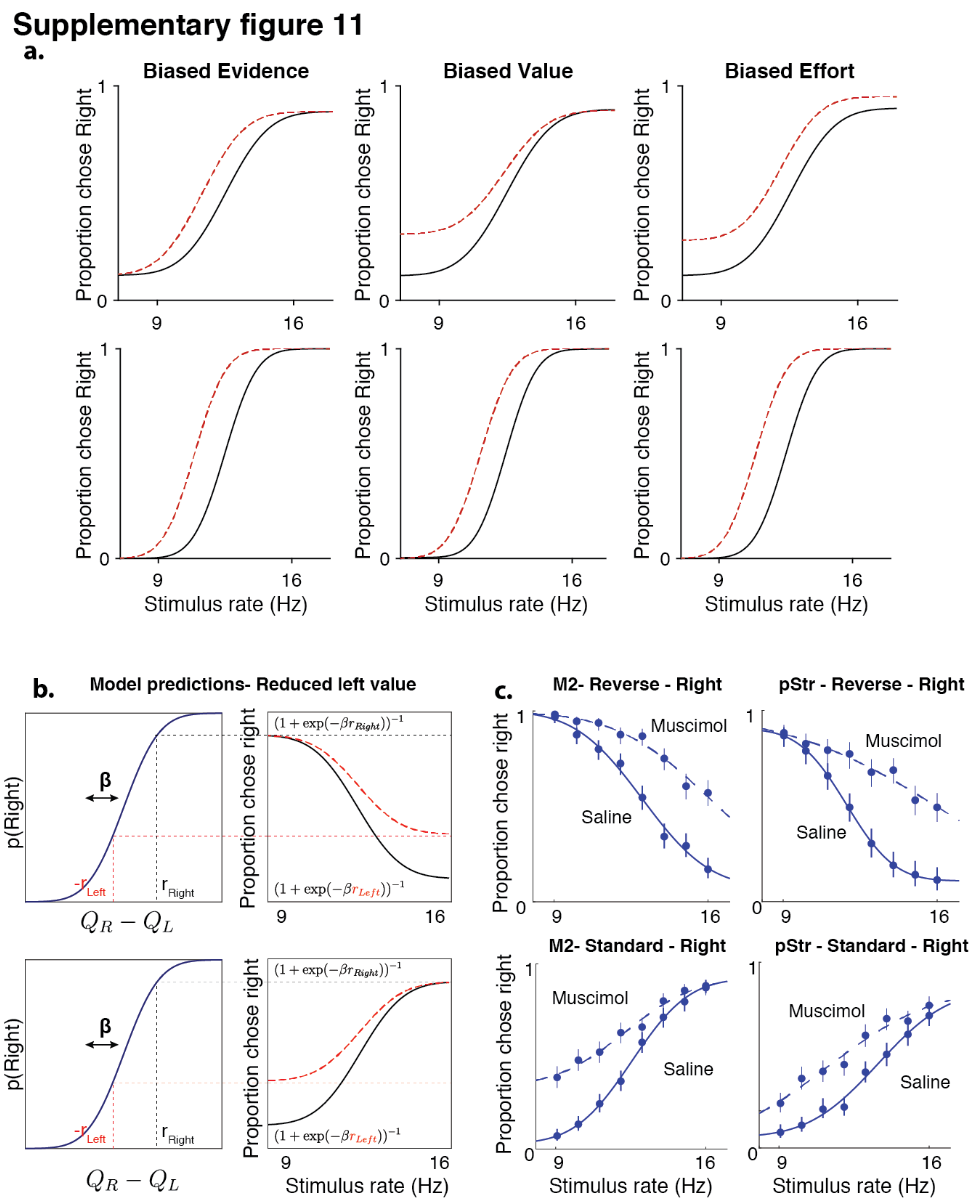

Supplementary Figure 11: Lapses differentiate perturbations to different stages of the decision-making

process (a) Model predictions for biased sensory evidence (left), decreased contralateral action value (center) and increased effort in performing contralateral movements (right). The three kinds of perturbations affect decisions at the sensory, value or motor stages and predict different effects on lapses (top), but reduce to the same effect (horizontal shift) in the absence of lapses (bottom). (b) Model predictions for rightward inactivations on standard (top) and reversed (bottom) contingencies - in both cases, the model predicts that reduced leftward action values should only affect lapses on the side associated with leftward movements. (c) Inactivation data on visual trials from M2 (left) or pStr (Right) shows a pattern of effects consistent with action value deficits, irrespective of the stimulus-response contingency.