



# Opponent control of behavior by dorsomedial striatal pathways depends on task demands and internal state

Scott S. Bolkan<sup>1,3</sup>, Iris R. Stone<sup>1,3</sup>, Lucas Pinto<sup>1</sup>, Zoe C. Ashwood<sup>1</sup>, Jorge M. Iravedra Garcia<sup>1</sup>, Alison L. Herman<sup>1</sup>, Priyanka Singh<sup>1</sup>, Akhil Bandi<sup>1</sup>, Julia Cox<sup>1</sup>, Christopher A. Zimmerman<sup>1</sup>, Jounhong Ryan Cho<sup>1</sup>, Ben Engelhard<sup>1</sup>, Jonathan W. Pillow<sup>1,2</sup> ✉ and Ilana B. Witten<sup>1,2</sup> ✉

**A classic view of the striatum holds that activity in direct and indirect pathways oppositely modulates motor output. Whether this involves direct control of movement, or reflects a cognitive process underlying movement, remains unresolved. Here we find that strong, opponent control of behavior by the two pathways of the dorsomedial striatum depends on the cognitive requirements of a task. Furthermore, a latent state model (a hidden Markov model with generalized linear model observations) reveals that—even within a single task—the contribution of the two pathways to behavior is state dependent. Specifically, the two pathways have large contributions in one of two states associated with a strategy of evidence accumulation, compared to a state associated with a strategy of repeating previous choices. Thus, both the demands imposed by a task, as well as the internal state of mice when performing a task, determine whether dorsomedial striatum pathways provide strong and opponent control of behavior.**

The striatum is composed of two principal outputs, the direct and indirect pathways, which are thought to exert opposing effects on behavior<sup>1</sup>. In support of this view, many influential studies have shown that pathway-specific activation of the striatum produces opposing behavioral biases<sup>2–14</sup>. For example, direct or indirect pathway activation oppositely influences locomotion<sup>2–4,14</sup>, licking<sup>5,11,15</sup>, left/right rotations<sup>2,3,11,16</sup>, repetition/cessation of activation-paired behaviors<sup>6–8</sup> and left/right movements to report value-based decisions<sup>9,13</sup>.

Despite this pioneering work, it remains unresolved whether the endogenous activity of the two pathways provides opposing control over the generation of movements, or instead contributes to the cognitive process of deciding which movement to perform. This is in part because pathway-specific manipulations have disproportionately relied on artificial and synchronous activation, rather than inhibition of endogenous activity<sup>2–11,13</sup>. The imbalance toward reports of activation suggests a wealth of negative results from inhibition, raising questions about the function of the endogenous activity, and whether it contributes to cognition. In fact, most previous pathway-specific activation studies have not used cognitively demanding tasks, making it difficult to dissociate a role in the decision toward a movement versus the generation of the movement itself<sup>2–4,6,11,14,16</sup>. In contrast, studies of the striatum that were not pathway specific have instead focused on cognitively demanding behaviors<sup>17–24</sup>. Taken together, this raises the possibility that striatal pathways exert opposing control of movement in the context of decision-making, rather than directly controlling motor output irrespective of cognition.

Thus, to determine if the contribution of endogenous activity in striatal pathways depends on cognition, we examined the effects of pathway-specific inhibition across a set of virtual reality tasks that had the same motor output and similar sensory features, but different

cognitive requirements. This allowed us to ask if a task's demands determined the effect of pathway-specific inhibition on behavior. Second, we used a latent state model to identify time-varying states within the same task. This allowed us to determine if the contribution of each pathway to behavior changed across time, even within the same task.

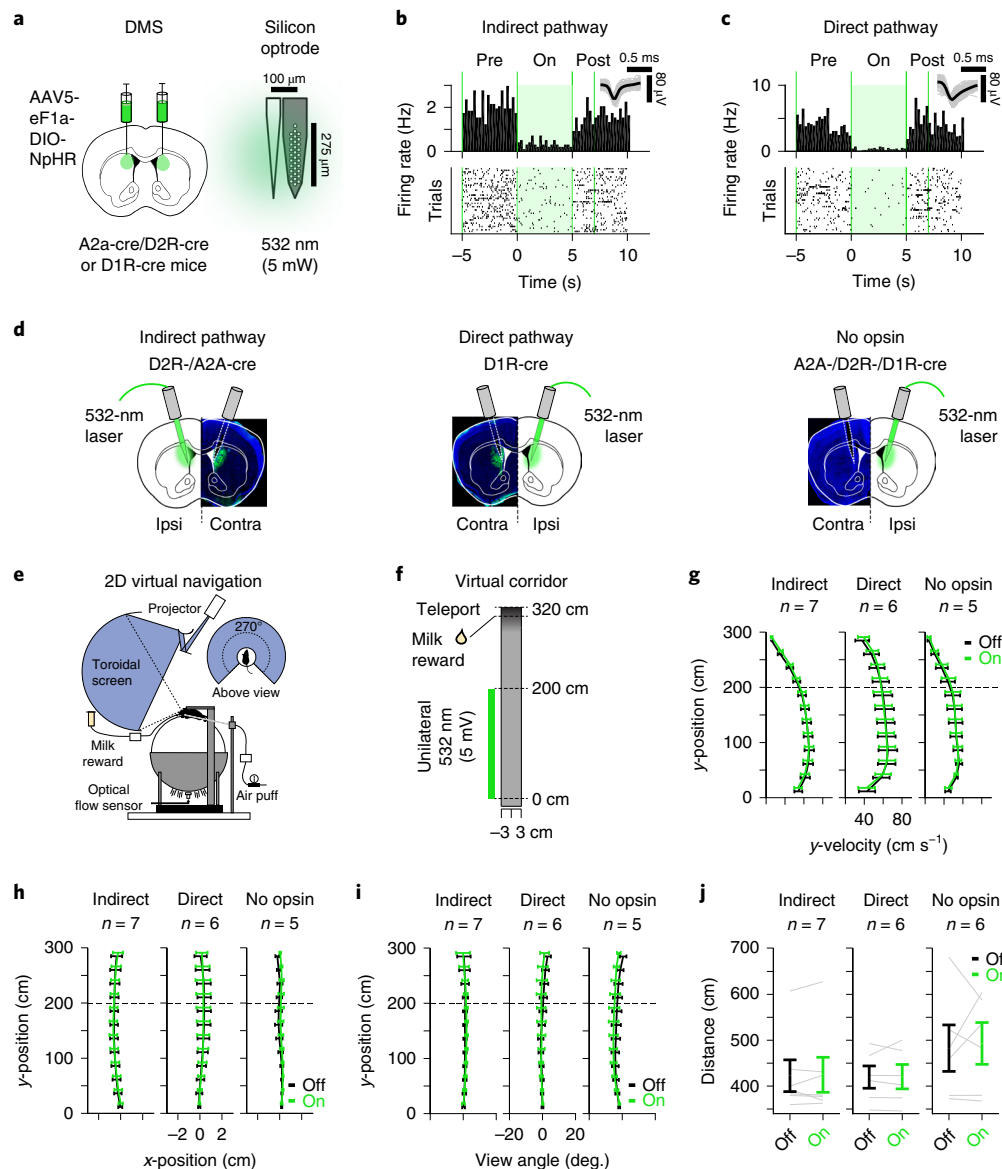
We found that inhibition of neither pathway produced a detectable influence on behavior as mice navigated a virtual corridor in the absence of a decision-making requirement. In contrast, pathway-specific inhibition produced strong and opposing biases on decisions based on the accumulation of evidence in a virtual T-maze<sup>25</sup>, and had weaker effects on choice during less demanding task variants. Our latent state model further revealed that even within the evidence accumulation task, mice occupy different states across time that differ in the weighting of sensory evidence and trial history, as well as the extent that pathway-specific inhibition impacts choice. Thus, by comparing the effects of pathway-specific inhibition across behavioral tasks, and across time within a task, we conclude that both demands of the task and internal state of the mice determine whether striatal pathways exert strong and opposing control over behavior.

## Results

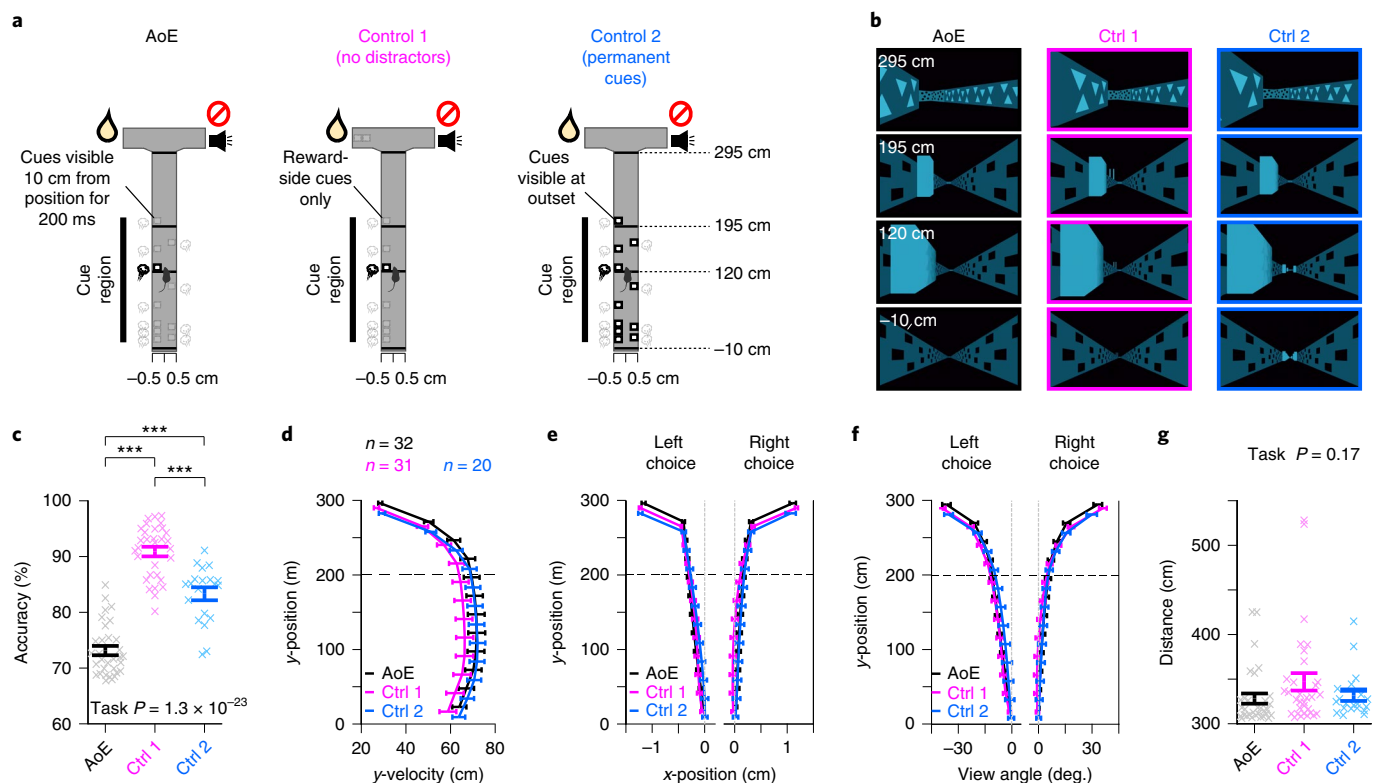
**Inhibition of pathway-specific dorsomedial striatum activity is effective.** We first sought to validate the effectiveness of halorhodsopsin<sup>18</sup> (NpHR)-mediated inhibition of indirect and direct striatal pathway activity in awake, head-fixed mice (Fig. 1a and Extended Data Fig. 1a,b). Toward this end, we bilaterally delivered virus carrying Cre-dependent NpHR to the dorsomedial striatum (DMS) in transgenic mouse lines (A2a-Cre/D2R-Cre/D1R-Cre), which we verified to have high degrees of specificity and penetrance for each pathway (Supplementary Fig. 1). We confirmed that 532-nm (5 mW)

<sup>1</sup>Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA. <sup>2</sup>Department of Psychology, Princeton University, Princeton, NJ, USA.

<sup>3</sup>These authors contributed equally: Scott S. Bolkan, Iris R. Stone. ✉e-mail: [pillow@princeton.edu](mailto:pillow@princeton.edu); [iwitten@princeton.edu](mailto:iwitten@princeton.edu)



**Fig. 1 | Pathway-specific dorsomedial striatum inhibition has no detectable impact on movement in mice navigating a virtual corridor.** **a**, Schematic of viral delivery of Cre-dependent NpHR to the DMS of A2a-Cre, D2R-Cre or D1R-Cre mice. Schematic of optrode, that is, a 32-channel silicon probe coupled with tapered optical fiber, which delivered 532-nm (5-mW) light to the DMS of awake, ambulating mice. **b**, Example peristimulus time histograms (PSTH; top) and rasters of trial-by-trial spike times (bottom) from a DMS single-unit recorded in an ambulating A2a-Cre mouse expressing Cre-dependent NpHR (indirect pathway). Inset shows an average spike waveform (black) and 100 randomly sampled spike waveforms (gray). A trial consisted of 5 s without laser (pre, -5 to 0 s), a 5-s laser sweep (on, 0 to 5 s) and a 10-s intertrial interval (ITI; 40 total trials). **c**, As in **b** but for DMS single-unit recorded in a D1R-Cre mouse expressing Cre-dependent NpHR (direct pathway). **d**, Schematic of bilateral fiber-optic implantation in behaving mice, with example histology from a mouse expressing NpHR in the indirect (D2R-/A2a-Cre) or direct (D1R-Cre) pathways, or control mouse without opsin (no opsin, A2a-/D2R-Cre or D1R-Cre). Light (532 nm; 5 mW) was delivered unilaterally to the left or right hemisphere on alternate testing sessions and lateralized behavior was defined as ipsilateral or contralateral relative to the laser hemisphere. **e**, Schematic of head fixation of mice in a VR apparatus allowing 2D navigation. Displacements of an air-suspended spherical ball in the anteroposterior (and mediolateral) axes of the mouse-controlled y- (and x-) position movements in a visual VR environment. **f**, Schematic of the virtual corridor (6-cm wide, 330-cm long), consisting of a start region (-10 to 0 cm), an inhibition region (0-200 cm) in which mice received unilateral 532 nm illumination on a random subset of trials (30%), a reward location (310 cm) where mice received reward, and a teleportation location (320 cm) where mice were transported to the start region following a variable ITI with mean of 2 s. **g**, Average y-velocity ( $\text{cm s}^{-1}$ ) across mice as a function of y-position (0-300 cm in 25-cm bins) while navigating the virtual corridor on laser-off (black) or laser-on (green) trials in groups receiving DMS indirect ( $n = 7$  mice,  $n = 1,288$  laser-off and  $n = 1,712$  laser-on trials) or direct ( $n = 6$  mice,  $n = 1,088$  laser-off and  $n = 757$  laser-on trials) pathway inhibition, or illumination of the DMS in the absence of NpHR expression (no opsin,  $n = 5$  mice,  $n = 1,178$  laser-off and  $n = 827$  laser-on trials). **h**, Same as **g** but for average x-position (cm) contralateral to the unilaterally coupled laser hemisphere. **i**, Same as **g** but for view angle (degrees, contralateral to laser hemisphere). **j**, Average across-mouse distance traveled (cm) to traverse the virtual corridor during laser-off (black) or laser-on (green) trials for mice receiving DMS indirect ( $n = 7$  mice,  $n = 2,109$  laser-off and  $n = 1,574$  laser-on trials) or direct ( $n = 6$  mice,  $n = 1,330$  laser-off and  $n = 930$  laser-on trials) pathway inhibition, or DMS illumination in the absence of NpHR ( $n = 6$  mice,  $n = 1,688$  laser-off and  $n = 1,199$  laser-on trials). Solid bars depict the mean  $\pm$  s.e.m. across mice; gray lines indicate individual mouse mean values.



**Fig. 2 | A set of virtual reality T-mazes has similar sensory features and identical motor requirements but different cognitive demands.** **a**, Schematic of three VR-based T-maze tasks. **b**, Example mouse perspective at the same maze position (–10 cm, 120 cm, 195 cm and 295 cm) from the example trial depicted in **a** of the accumulation of evidence (AoE; black), no distractors (ctrl 1) or permanent cues (ctrl 2) tasks. **c**, Average choice accuracy (percentage correct) across mice performing the accumulation of evidence (black,  $n = 32$  mice,  $n = 52,381$  trials), no distractors (magenta, ctrl 1:  $n = 31$  mice,  $n = 56,783$  trials) or permanent cues (cyan, ctrl 2:  $n = 20$  mice,  $n = 27,870$  trials) tasks.  $P$  value denotes one-way analysis of variance (ANOVA) of task on accuracy ( $P = 1.3 \times 10^{-23}$ ,  $F_{2,80} = 109.4$ ). Asterisks indicate statistical significance of post hoc, unpaired, two-tailed rank-sum comparisons of accuracy between groups (top to bottom:  $***P = 3.9 \times 10^{-7}$ ,  $z = -5.1$ ;  $***P = 2.1 \times 10^{-11}$ ,  $z = -6.7$ ;  $***P = 2.1 \times 10^{-5}$ ,  $z = 4.3$ ). **d**, Average y-velocity (cm s<sup>-1</sup>) across mice as a function of y-position (0–300 cm in 25-cm bins) during performance of each task (colors and  $n$  as in **c**). **e**, Same as **d** but for average x-position (cm) on left/right choice trials. **f**, Same as **d** but for average view angle (degrees) on left/right choice trials. **g**, Average distance (cm) traveled per trial across mice (evidence accumulation,  $n = 32$  mice,  $n = 53,833$  trials; no distractors (ctrl 1):  $n = 32$  mice,  $n = 60,074$  trials; permanent cues (ctrl 2):  $n = 20$  mice,  $n = 29,192$  trials).  $P$  value reflects one-way ANOVA of task on distance ( $P = 0.16$ ,  $F_{2,81} = 1.8$ ). Solid bars denote the across-mouse mean  $\pm$  s.e.m. and the transparent 'x' indicates the mean for individual mice.

of light delivery to the DMS through a tapered optical fiber produced rapid, sustained and reversible inhibition of spiking in mice expressing NpHR in the indirect pathway (Fig. 1b and Extended Data Fig. 1c–e,  $n = 18/60$ , 30% of neurons significantly inhibited) or the direct pathway (Fig. 1c and Extended Data Fig. 1f–h,  $n = 21/50$ , 42% of neurons significantly inhibited). Moreover, we observed: (1) minimal excitation during illumination<sup>26,27</sup> (Extended Data Fig. 1d,g), (2) minimal effects on spiking upon laser offset (Extended Data Fig. 1d,g), indicating limited post-inhibitory rebound, and (3) stability in the efficacy of inhibition across time (Supplementary Fig. 2). Together, our findings indicate that NpHR-mediated inhibition of DMS pathways is effective.

**Dorsomedial striatum pathway inhibition does not impact virtual corridor navigation.** To determine if endogenous activity in DMS pathways provides bidirectional control of motor output in the absence of a decision, we carried out unilateral inhibition of indirect and direct pathways in head-fixed mice running on an air-supported ball to traverse a two-dimensional (2D) linear corridor in virtual reality (VR; Figs. 1d–f, 6-cm  $\times$  330-cm corridor). Illumination of the DMS was restricted to 0–200 cm (laser on for 30% of trials; hemisphere of illumination alternated across days).

The parameters of the virtual corridor and inhibition period were selected to closely match the stem of the VR-based T-maze decision-making tasks that are the focus of subsequent experiments.

We found no detectable impact of pathway-specific DMS inhibition, nor DMS illumination alone, on multiple indicators of motor output during virtual corridor navigation. This included measures of velocity, x-position or view angle relative to the laser hemisphere, and distance traveled (Fig. 1g–j; see Extended Data Fig. 2 for additional measures). Similarly, we obtained null effects of pathway-specific inhibition on velocity (and spatial preference) in freely behaving mice in a conditioned place preference (CPP) assay (Supplementary Fig. 3).

These negative findings argue against a major involvement of endogenous activity in DMS pathways in the execution of movement in the absence of a decision. This is consistent with the dearth of reports demonstrating strong and opposing modulation of behavior by striatal pathways using pathway-specific optogenetic inhibition.

**Three virtual reality T-mazes with varying cognitive demands.** We next considered the possibility that, rather than contributing directly to a motor output, endogenous activity in DMS pathways

may instead have opposing influence over decisions in a manner that is dependent on cognitive demand. To test this idea, we trained mice to perform a set of VR-based, decision-making tasks<sup>25</sup> that shared identical motor readouts (left or right choice), had highly similar sensory environments, and yet differed in their cognitive requirements (Fig. 2a,b).

The first task was an ‘evidence accumulation’ task, in which visuo-tactile cues were transiently presented on each side of the central stem of a virtual T-maze according to a Poisson distribution (‘cue region’, 0–200 cm), and mice were rewarded for turning to the maze side with the greater number of cues (Fig. 2a,b; black). Thus, mice were required to continually accumulate sensory cues over several seconds into a memory (or motor plan) that guided their left/right decision.

In two additional control tasks, we made modifications intended to weaken the cognitive demands of each task. In the first control task (‘no distractors’), cues were presented on the rewarded maze side during the same maze region (0–200 cm) according to the same Poisson distribution, but distractor cues on the side of the non-rewarded arm were omitted (Fig. 2a,b; magenta). The absence of distractors on the non-rewarded side meant that each cue signaled reward with 100% probability, and thus gradual evidence accumulation was not required. Further ensuring that evidence accumulation was not required, an additional cue at the end of the maze was present only during the cue period (0–200 cm) to signal the rewarded side.

In the second control task (‘permanent cues’), the sensory statistics of the cues were identical to those in the evidence accumulation task, but rather than transient visual cue presentation, visual cues were permanently visible from trial onset (Fig. 2a,b; cyan). This maintained the same conceptual task structure of the evidence accumulation task while decreasing the memory demands, as the sensory cues (or the motor plan) did not need to be remembered until the cues were passed.

We assessed how task demands impacted choice accuracy in each task. Consistent with the greatest cognitive and mnemonic

demand in the evidence accumulation task, we found that overall choice accuracy was significantly lower compared to both control tasks (Fig. 2c, AoE: 73.1  $\pm$  0.8%. Ctrl 1: 90.6  $\pm$  0.9%. Ctrl 2: 83.3  $\pm$  1.2% mean  $\pm$  s.e.m.).

While the motor requirements of a decision were the same across tasks (crossing an  $x$ -position threshold at the end of the central stem; Methods), we examined the possibility that cross-task differences in cognitive requirements altered movement within the stem of the maze (0–300 cm). We observed no consistent cross-task differences in velocity,  $x$ -position or view angle on left or right choice trials, nor distance traveled (Fig. 2d–g; see Extended Data Fig. 3a–f for additional measures). We further compared the relationship between behavior in the stem of the maze and choice across tasks by using a decoder to predict choice based on the trial-by-trial  $x$ -position or view angle (Extended Data Fig. 3g–j) at successive maze positions (0–300 cm in 25-cm bins). While we were able to predict choice from either measure above chance levels in all three tasks (consistent with previous studies<sup>25</sup>), choice prediction accuracy was statistically indistinguishable across tasks (Extended Data Fig. 3g–j). Together, this indicated that cross-task differences in task demands did not prompt mice to systematically adopt distinct motor strategies.

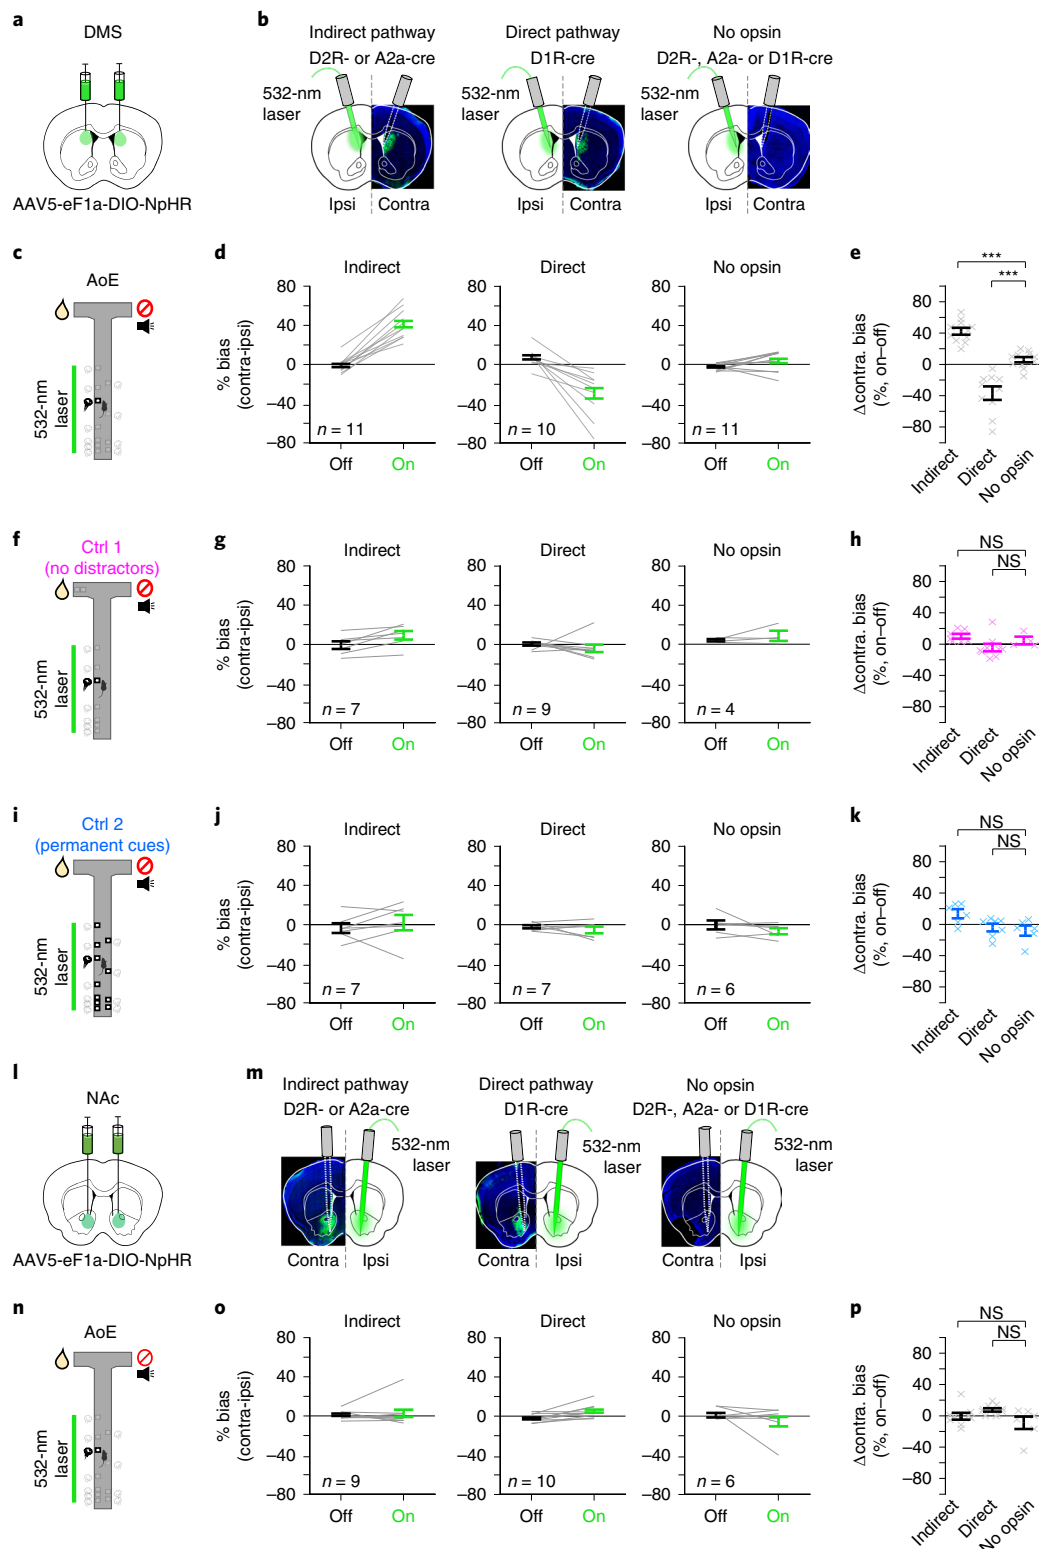
**Behavioral effects of dorsomedial striatum pathway inhibition depend on task demand.** We performed unilateral inhibition of DMS indirect and direct pathways restricted to the cue region (0–200 cm) of each task (Fig. 3a,b; laser on 10–20% of trials; hemisphere of illumination alternated across days). We found that inhibition of the indirect pathway produced a large bias toward contralateral choices during the accumulation of evidence task (Fig. 3c,d), which was significantly greater than that observed in control animals that did not express opsin (Fig. 3e, average contralateral bias: DMS indirect, 42.3%  $\pm$  4.4%, versus no opsin, 5.9%  $\pm$  3.6%). Similarly, inhibition of the direct pathway also produced a large choice bias during the accumulation of evidence task

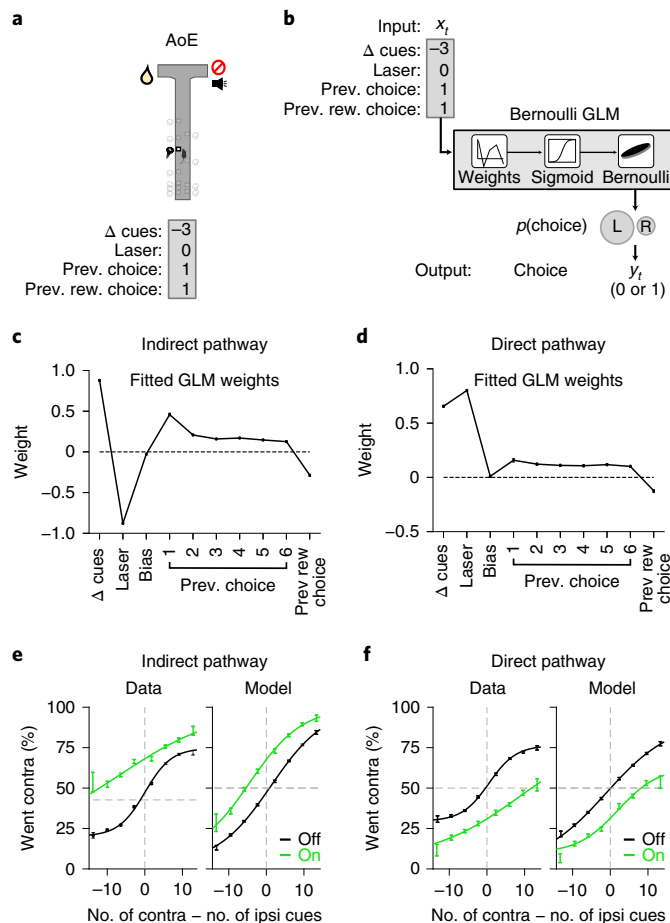
**Fig. 3 | Inhibition of dorsomedial striatum but not nucleus accumbens pathways has strong and opposing influence on choice during an evidence accumulation task, while having weaker effects during task variants with diminished cognitive demands.** **a**, Schematic of bilateral viral delivery of Cre-dependent NpHR to the DMS. **b**, Schematic of bilateral fiber-optic implantation of the DMS and unilateral inhibition in behaving mice, with example histology from a mouse expressing NpHR in the indirect (D2R-/A2a-Cre) or direct (D1R-Cre) pathways, or DMS illumination in the absence of NpHR (no opsin, A2a-/D2R- or D1R-Cre). Light (532 nm, 5 mW) was delivered unilaterally to the left or right hemisphere on alternate testing sessions and choice bias contralateral or ipsilateral to the hemisphere of inhibition was quantified. **c**, Schematic of the evidence accumulation task with delivery of 532-nm light restricted to the cue region (0–200 cm) on a random subset of trials (10–20%). **d**, Average across-mouse choice bias during the evidence accumulation task. Choice bias was defined as the difference between the percentage of correct performance on trials when the correct choice was contralateral or ipsilateral to the inhibited hemisphere (percentage correct, contralateral–ipsilateral, positive values indicate a contralateral bias). Bias was calculated separately on laser-off (black) and laser-on (green) trials for mice receiving unilateral indirect pathway inhibition ( $n = 11$  mice,  $n = 16,935$  laser-off and  $n = 3,390$  laser-on trials), unilateral direct pathway inhibition ( $n = 10$  mice;  $n = 14,030$  laser-off and  $n = 3,103$  laser-on trials), or unilateral illumination of the DMS in the absence of NpHR ( $n = 11$  mice,  $n = 21,422$  laser-off and  $n = 5,113$  laser-on trials). **e**, Difference in contralateral choice bias (percentage correct) between laser-off and laser-on trials (percentage bias, on–off) in mice performing the evidence accumulation task and receiving indirect pathway inhibition, direct pathway inhibition or DMS illumination in the absence of NpHR. Asterisks indicate the significance of an unpaired, two-tailed Wilcoxon rank-sum comparison of indirect to no opsin: \*\*\* $P = 1.1 \times 10^{-4}$ ,  $z = 3.9$ ; direct to no opsin: \*\*\* $P = 2.2 \times 10^{-4}$ ,  $z = -3.7$ . **f–h**, Same as **c–e** but for the no-distractors (ctrl 1) task. Indirect:  $n = 7$  mice,  $n = 13,706$  laser-off and  $n = 3,288$  laser-on trials; direct:  $n = 9$  mice,  $n = 14,647$  laser-off and  $n = 3,682$  laser-on trials; no opsin:  $n = 4$  mice,  $n = 3,654$  laser-off and  $n = 901$  laser-on trials. Asterisks indicate the significance of an unpaired, two-tailed Wilcoxon rank-sum comparison of indirect to no opsin: not significant (NS),  $P = 0.22$ ,  $z = 1.2$ . Direct to no opsin: NS,  $P = 0.08$ ,  $z = -1.8$ . **i–k**, As in **c–e** but for the permanent cues (ctrl 2) task. Indirect:  $n = 7$  mice,  $n = 4,033$  laser-off and  $n = 929$  laser-on trials; direct:  $n = 7$  mice,  $n = 6,061$  laser-off and  $n = 1,494$  laser-on trials; no opsin:  $n = 6$  mice,  $n = 3,975$  laser-off and  $n = 923$  laser-on trials. Asterisks indicate the significance of an unpaired, two-tailed Wilcoxon rank-sum comparison of indirect to no opsin: NS,  $P = 0.13$ ,  $z = 1.5$ . Direct to no opsin: NS,  $P = 0.62$ ,  $z = 0.5$ . **l**, As in **a** but for bilateral viral delivery of Cre-dependent NpHR to the NAc. **m**, Same as **b** but for bilateral fiber-optic implantation of the NAc and unilateral inhibition in behaving mice, with example histology from a mouse expressing NpHR in the indirect (D2R-/A2a-Cre) or direct (D1R-Cre) pathways, or NAc illumination in the absence of NpHR (no opsin, A2a-/D2R-Cre or D1R-Cre). **n–p**, As in **c** but for pathway-specific NAc inhibition during the accumulation of evidence task. Indirect:  $n = 9$  mice,  $n = 11,978$  laser-off and  $n = 2,604$  laser-on trials; direct:  $n = 10$  mice,  $n = 15,430$  laser-off and  $n = 3,348$  laser-on trials; no opsin:  $n = 7$  mice,  $n = 9,819$  laser-off and  $n = 1,488$  laser-on trials. Asterisks indicate the significance of an unpaired, two-tailed Wilcoxon rank-sum comparison of indirect to no opsin: NS,  $P = 0.86$ ,  $z = 0.18$ ; direct to no opsin: NS,  $P = 0.04$ ,  $z = 2.0$ . Solid bars denote across-mouse mean value  $\pm$  s.e.m. and the transparent ‘x’ indicates the mean for individual mice. To account for multiple-group comparisons, we considered  $P$  values significant after Bonferroni correction (two comparisons).



(Fig. 3d; average contralateral bias: DMS direct,  $-36.8\% \pm 8.6\%$ ), which was also significantly greater than that observed in control animals (Fig. 3e). However, in this case, the direction of the choice bias was in the opposite (ipsilateral) direction to that observed with indirect pathway inhibition (also see Extended Data Fig. 4a–i for psychometric curves).

Providing a stark contrast to the large effects of pathway-specific DMS inhibition on choice during the evidence accumulation task, inhibition of either pathway had significantly less impact on choice during the ‘no distractors’ and ‘permanent cues’ control tasks (Fig. 3f–k and Extended Data Fig. 5a–c; unpaired, two-tailed Wilcoxon rank-sum test of indirect pathway inhibition evidence accumulation





**Fig. 4 | A GLM reveals that sensory evidence, dorsomedial striatum pathway inhibition and trial history predict choice during the evidence accumulation task, but does not precisely recapitulate the shape of the psychometric curve. a**, Schematic of the evidence accumulation task and the coding of the external covariates for an example trial. **b**, Schematic of the Bernoulli GLM for an example trial, showing the relationship between external covariates (inputs) and choice on each trial. On each trial, a set of GLM weights maps each input ( $\Delta$  cues, laser, bias, previous choice and existence of a previous rewarded choice) to the probability of each outcome through a sigmoid function, which gives the probability of a 'rightward' choice on the current trial. **c**, Fitted GLM weights using aggregated data from all mice in the indirect pathway DMS inhibition group. The magnitude of each weight indicates the relative importance of that covariate in predicting choice, whereas the sign of the weight indicates the direction of the effect (for example, a negative laser weight indicates that if inhibition is in the right hemisphere, the mice will be more likely to turn left, while a positive weight on the previous choice indicates that if the previous choice was to the right, in the current trial this will bias the mice to turn right again). Error bars denote  $\pm 1$  posterior s.d. credible intervals. **d**, Same as **c** but for mice receiving DMS direct pathway inhibition. **e**, Fraction of contralateral choice trials as a function of the difference in contralateral versus ipsilateral cues for laser-off (black) and laser-on (green) trials, for mice receiving indirect pathway DMS inhibition for the data and for simulations of the model. Error bars denote 95% confidence intervals around the fraction of choices in each bin of the data; solid curves denote logistic fits ( $n=13$  mice,  $n=46,313$  laser-off and  $n=8,570$  laser-on trials). **f**, Same as **e** but for the mice receiving direct pathway inhibition of the DMS ( $n=13$  mice,  $n=41,250$  laser-off and  $n=7,927$  laser-on trials).

versus no distractors,  $P=8.0 \times 10^{-4}$ ,  $z=3.4$ , or evidence accumulation versus permanent cues,  $P=0.001$ ,  $z=3.3$ ; and direct pathway inhibition evidence accumulation versus no distractors,

$P=0.002$ ,  $z=-3.1$ , or evidence accumulation versus permanent cues,  $P=0.005$ ,  $z=-2.8$ ). In fact, the effects of pathway-specific DMS inhibition on choice bias in either control task did not significantly differ from those observed in control animals (Fig. 3h, for 'no distractors'; Fig. 3k, for 'permanent cues'; see also Extended Data Fig. 4a–i for psychometric curves).

Thus, inhibition of DMS pathways elicited strong and opposing effects on choice in the task with the greatest cognitive demand, which required the accumulation of sensory evidence across multiple seconds to arrive at a decision and had a far limited impact on choice in task variants with reduced cognitive demand.

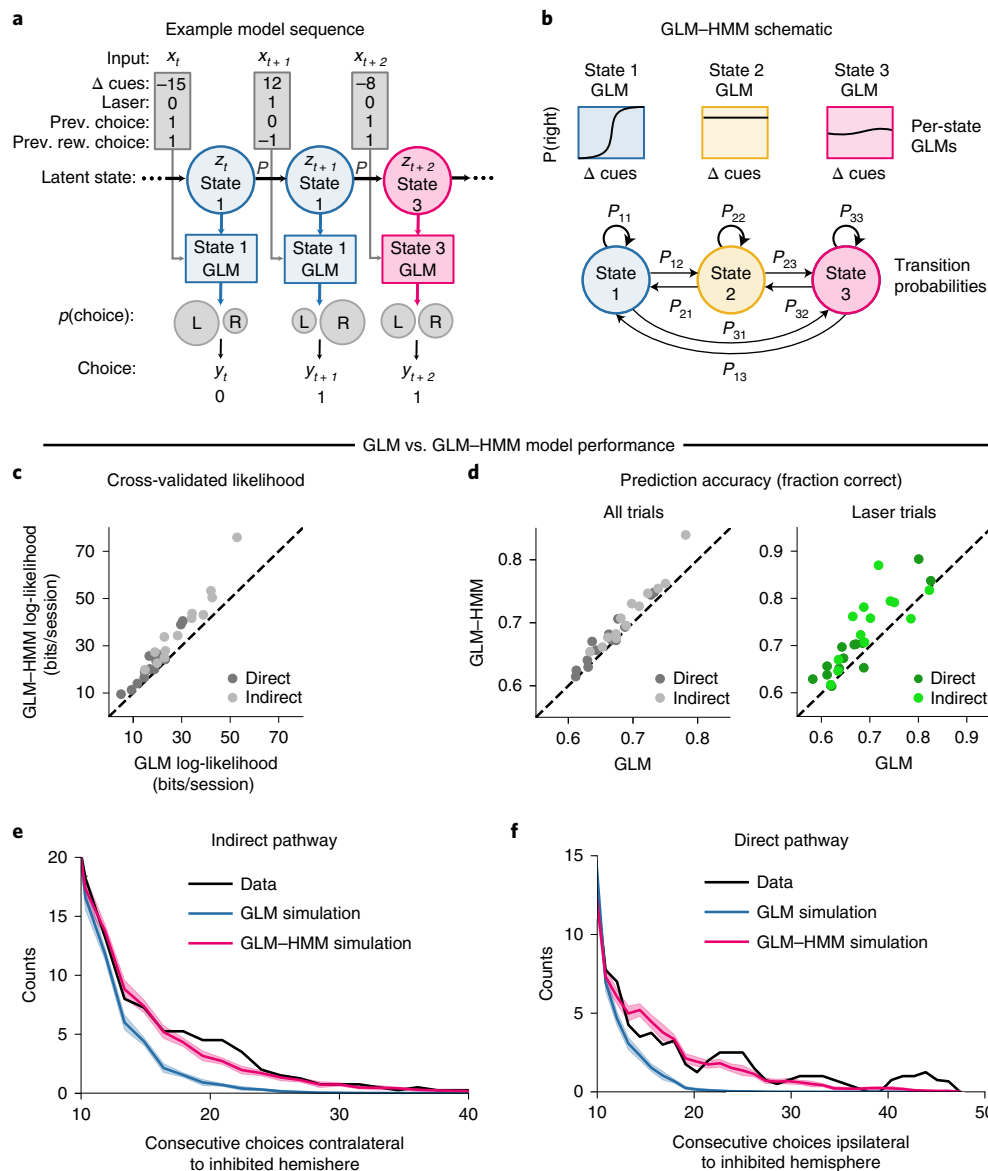
While DMS pathway inhibition had minimal impact on movement in a virtual corridor (Fig. 1 and Extended Data Fig. 2), we considered the possibility that pathway-specific DMS inhibition altered motor performance in the T-mazes. We found no cross-task differences in the effects of pathway-specific inhibition on measures of velocity, distance traveled or per-trial standard deviation in view angle (Extended Data Fig. 6a–i). However, we found subtle but opposing effects of pathway-specific inhibition on average  $x$ -position and view angle (Extended Data Fig. 6j–k) in the evidence accumulation task. The direction of these biases was similar in the control tasks, but consistently smaller than in the evidence accumulation task. Thus, in line with the close relationship between  $x$ -position/view angle and choice in the absence of inhibition in each task (Extended Data Fig. 3g–j), pathway-specific DMS inhibition produced the same general pattern of cross-task effects on choice bias (Extended Data Fig. 5b–d) and  $x$ -position/view angle (Extended Data Fig. 6j,k). As the quantitative relationship between  $x$ -position or view angle and choice is indistinguishable across tasks in the absence of neural inhibition (Extended Data Fig. 3g–j), cross-task differences in motor strategy do not provide a trivial explanation for these effects. Rather, taken together with the absence of an effect of pathway-specific DMS inhibition on motor output in the virtual corridor (Fig. 1h,i), these data imply that the effects of inhibition on behavior depend on cognitive demands.

#### Little effect of nucleus accumbens pathway inhibition on choice.

We next sought to determine whether opponent control of choice by striatal pathways during the evidence accumulation task was specific to the DMS, or if it extended to the ventral striatum. To this end, we delivered unilateral laser illumination to the nucleus accumbens (NAc) of mice expressing NpHR in the indirect or direct pathways (or non-opsin control mice), which was restricted to the cue region (0–200 cm) of the evidence accumulation task (Fig. 3l–p and Extended Data Fig. 4j–l).

Providing a clear functional dissociation between DMS and NAc, effects of pathway-specific NAc inhibition on choice bias were significantly smaller than those observed with inhibition of DMS pathways (Extended Data Fig. 5e,f; unpaired, two-tailed Wilcoxon rank-sum test of DMS versus NAc indirect pathway inhibition,  $P=2.6 \times 10^{-4}$ ,  $z=3.6$ ; of DMS versus NAc direct pathway inhibition,  $P=1.8 \times 10^{-4}$ ,  $z=-3.7$ ), and were also not significantly different from NAc control animals (Fig. 3o,p). It is unlikely that this dissociation between DMS and NAc can be explained by greater coexpression of pathway-specific markers in the ventral versus dorsal striatum<sup>28</sup>, as both subregions exhibited equally low colocalization of D1R and D2R receptors (Supplementary Fig. 1j–l).

**Bernoulli generalized linear model does not fully capture psychometric curves.** Our inactivation experiments suggest that DMS pathways make strong contributions to behavior during a cognitively demanding evidence accumulation task, but do not contribute strongly to similar tasks with weaker cognitive demands. However, even during the evidence accumulation task, it is possible that the animals' level of cognitive engagement varies over time. This raises



**Fig. 5 | A GLM-HMM better explains choice during the evidence accumulation task than the GLM, particularly on trials with dorsomedial striatum pathway inhibition.** **a**, Example sequence of three trials of the evidence accumulation task, showing the relationship between external covariates (inputs), latent state and choice on each trial. On each trial, the latent state defines which GLM weights map inputs ( $\Delta$  cues, laser, previous choice and previous rewarded choice) to the probability of choosing right or left. The transition probability  $P$  governs the probability of changing states between trials. See Methods for information on how the inputs were coded. **b**, Schematic of GLM-HMM. The model has three latent states with fixed probabilities of transitioning between them. Each state is associated with a distinct decision-making strategy, defined by a mapping from external covariates or inputs, such as  $\Delta$  cues, to choice probability. **c**, Cross-validated log-likelihood demonstrating the increased performance of the GLM-HMM over a standard Bernoulli GLM on held-out sessions. Dots represent model performance for individual mice ( $n = 13$  for each group). **d**, Same as **c** but showing prediction accuracy as a fraction of the choices correctly predicted by each model across all trials or on the subset of trials when the laser was on. **e**, Histograms showing the number of consecutive laser trials for which the animal's choice was in the same direction as the expected biasing effect of the laser (that is, a choice contralateral to the laser hemisphere during DMS indirect pathway inhibition). Data are shown in black, the GLM simulation is shown in blue and the GLM-HMM simulation is shown in pink. For the simulations, data of the same length as the real data were generated 100 times and the resulting histograms averaged. Curves denote smoothed counts using a sliding window average (window size of three bins). Shaded regions around the GLM and GLM-HMM curves indicate 95% confidence intervals. **f**, Same as **e** but for mice receiving direct pathway inhibition of the DMS; therefore, laser-biased choices were defined as those ipsilateral to the hemisphere of inhibition.

the possibility that the contributions of the two pathways to behavior could change over time, even within the same task.

To address this possibility, we sought to understand the factors that contribute to decisions in the evidence accumulation task. As a first step, we used a Bernoulli generalized linear model (GLM)

to predict choice based on a set of external covariates (Fig. 4a,b). These covariates included the sensory evidence (difference between the number of right and left cues, or ' $\Delta$  cues'), the recent choice and reward history, the delivery of optical inhibition ('laser'), as well as a bias. Note that we set the value of the laser covariate to +1

(or -1) on trials with right (or left) hemisphere inhibition, and zero otherwise. A positive (or negative) GLM weight on this covariate thus captured an ipsilateral (or contralateral) 'laser'-induced bias in choices relative to the hemisphere of inhibition. For the choice history covariates, a positive weight indicates a tendency toward repeating past choices (Methods).

We fit the GLM to aggregated behavioral data from mice inhibited in each DMS pathway and found that sensory evidence, trial history and optical inhibition all contributed to predicting choice (Fig. 4c,d). As expected, the effect of inhibition of each pathway was large and opposite in sign. However, the GLM did not accurately capture the animals' psychometric curve, describing the probability of a rightward choice as a function of the sensory evidence (Fig. 4e,f). This led us to consider variants of the standard GLM that might better account for choice behavior.

**GLM-HMM better explains choice data with dorsomedial striatum inhibition.** The standard GLM describes choice as depending on a fixed linear combination of sensory evidence, trial history and the presence or absence of optical inhibition. However, an alternative possibility is that mice use a weighting function that varies over time. To test this idea, we adopted a latent state model with different GLM weights for different states (Fig. 4). The model consists of a hidden Markov model (HMM) with Bernoulli GLM observations, or GLM-HMM<sup>29–32</sup> (Fig. 5a,b). Each hidden state is associated with a unique set of GLM weights governing choice behavior in that state. Probabilistic transitions between states occur after every trial, governed by a fixed matrix of transition probabilities (Methods).

The GLM-HMM explained the choice data in the evidence accumulation task better than the GLM across multiple measures. To compare models, we computed the test log-likelihood of each animal's data using cross-validation with held-out sessions (three-state GLM-HMM in Fig. 5; see Extended Data Fig. 7a–e and Methods ('GLM-HMM') for more information on model selection). The three-state GLM-HMM achieved an average of a 6.2 bits per session (bps) increase in log-likelihood, making an average session ~76 times more likely under the GLM-HMM (Fig. 5c). Furthermore, the GLM-HMM correctly predicted choice on held-out data more often than the GLM, especially on laser trials (Fig. 5d; average improvement across mice of 1.6% on all trials, 3.5% on trials with optical inhibition and 4.1% on trials with optical inhibition when considering only mice with at least 100 inhibition trials).

Most interestingly, the GLM-HMM was better able to capture the temporal structure in the effect of inhibition on choice. Specifically, the choice data contained long runs in which choice was consistent with the bias direction predicted by pathway-specific inhibition ('laser'), a feature which GLM-HMM simulations recapitulated, but GLM simulations did not (Fig. 5e,f). Thus, taken

together, the GLM-HMM provided a better model of the choice data than a standard GLM, particularly on trials with pathway-specific DMS inhibition.

**GLM-HMM identifies states with varying dorsomedial striatum dependence.** We examined the state-dependent weights of the GLM-HMM and found substantial differences across states in the weighting of sensory evidence, previous choice and, most intriguingly, inhibition of DMS pathways (Fig. 6a,b). In particular, two of the three states (states 1 and 2) displayed a large weighting of sensory evidence on choice, while the 'laser' weight was large only in state 2. In contrast, in state 3, choice history had a larger weight than in the other states, and neither sensory evidence nor 'laser' had much influence on choice.

To characterize state-dependent psychometric performance, we used the fitted model to compute the posterior probability of each state given the choice data and assigned each trial to its most probable state (Fig. 6c,d). We then examined the psychometric curves for trials assigned to each state. In state 3, performance was low (Fig. 6g) and DMS inhibition had little effect on behavior (Fig. 6c,d). This is consistent with the high GLM weight on choice history in this state and low weights on sensory evidence and laser (Fig. 6a,b). This implies relatively little contribution of DMS pathways during a task-disengaged state when mice pursued a strategy of repeating previous choices rather than accumulating sensory evidence. When considered together with comparisons of the effect of pathway-specific DMS inhibition in control T-maze tasks where performance is high (Fig. 2c) but effects of inhibition are limited (Fig. 3f–k and Extended Data Fig. 5b,c), this implies a dissociation between task performance and the contributions of DMS pathways to behavior.

Compared to state 3, sensory evidence heavily modulated behavior in both states 1 and 2, and performance was accordingly high (Fig. 6c,d,g). Interestingly, the effect of DMS pathway inhibition was much larger in state 2. These results were again consistent with the GLM weights: both state 1 and 2 had high weighting of sensory evidence and low weighting of choice history but greatly differed in their weighting of the 'laser' (Fig. 6a,b). The discovery of state 2 implies that DMS pathways contribute most heavily to choices in a state in which mice are pursuing a strategy of evidence accumulation, consistent with cross-task comparisons of the effects of inhibition (Fig. 3). The discovery of state 1, which differed most noticeably from state 2 in the extent that the laser affected choice, may suggest the existence of another neural mechanism for evidence accumulation with minimal DMS dependence.

We found that GLM-HMM simulations closely recapitulated these state-dependent psychometric curves (Fig. 6e,f). This not only validated our fitting procedure but provided additional evidence

**Fig. 6 | A GLM-HMM uncovers states during the evidence accumulation task with different weighting on sensory evidence, choice history and dorsomedial striatum pathway inhibition.** **a**, Fitted GLM weights for the three-state model from mice in the indirect pathway DMS inhibition group. Error bars denote  $\pm 1$  posterior s.d. for each weight. The magnitude of the weight represents the relative importance of that covariate in predicting choice, and the sign of the weight indicates whether it biases choice to the left or to the right. Thus, for example, a negative 'laser' weight indicates that inhibition of the right hemisphere makes the mouse more likely to turn left, while a positive 'previous choice' weight makes the mouse more likely to repeat its previous choice. **b**, Per-state GLM weights for the direct pathway group. **c**, Fraction of contralateral choices as a function of the difference in contralateral versus ipsilateral cues in each trial for mice in the indirect pathway inhibition group. To compute psychometric functions for each state, we assigned each trial to its most probable state under the fitted model. Error bars denote  $\pm 1$  s.e.m. for light-off (solid) and light-on (dashed) trials. Solid curves denote logistic fits to the concatenated data across mice for light-off (solid) and light-on (dashed) trials. **d**, Same as **c** but for the mice receiving direct pathway inhibition of the DMS. **e**, Same as **c** but for data simulated from the model fit to mice receiving indirect pathway inhibition of the DMS (Methods). **f**, Same as **e** but for mice receiving direct pathway inhibition of the DMS. **g**, Performance in each state for mice receiving DMS inhibition in the indirect and direct pathways, shown as the percentage of total trials assigned to that state in which the mice made the correct choice. Colored bars denote the average performance across all mice. Black dots show averages for individual mice ( $n=13$  mice for both groups). **h**, The percentage of 'laser-on' trials that the model assigned to each state for mice receiving DMS inhibition in the indirect and direct pathways. Colored bars denote the average performance across all mice. Black dots show averages for individual mice ( $n=13$  mice for both groups). **i**, The posterior probability of each state for the five trials before and after a laser-on trial, averaged across all such periods ( $n=8,570$ , indirect;  $n=7,927$ , direct).



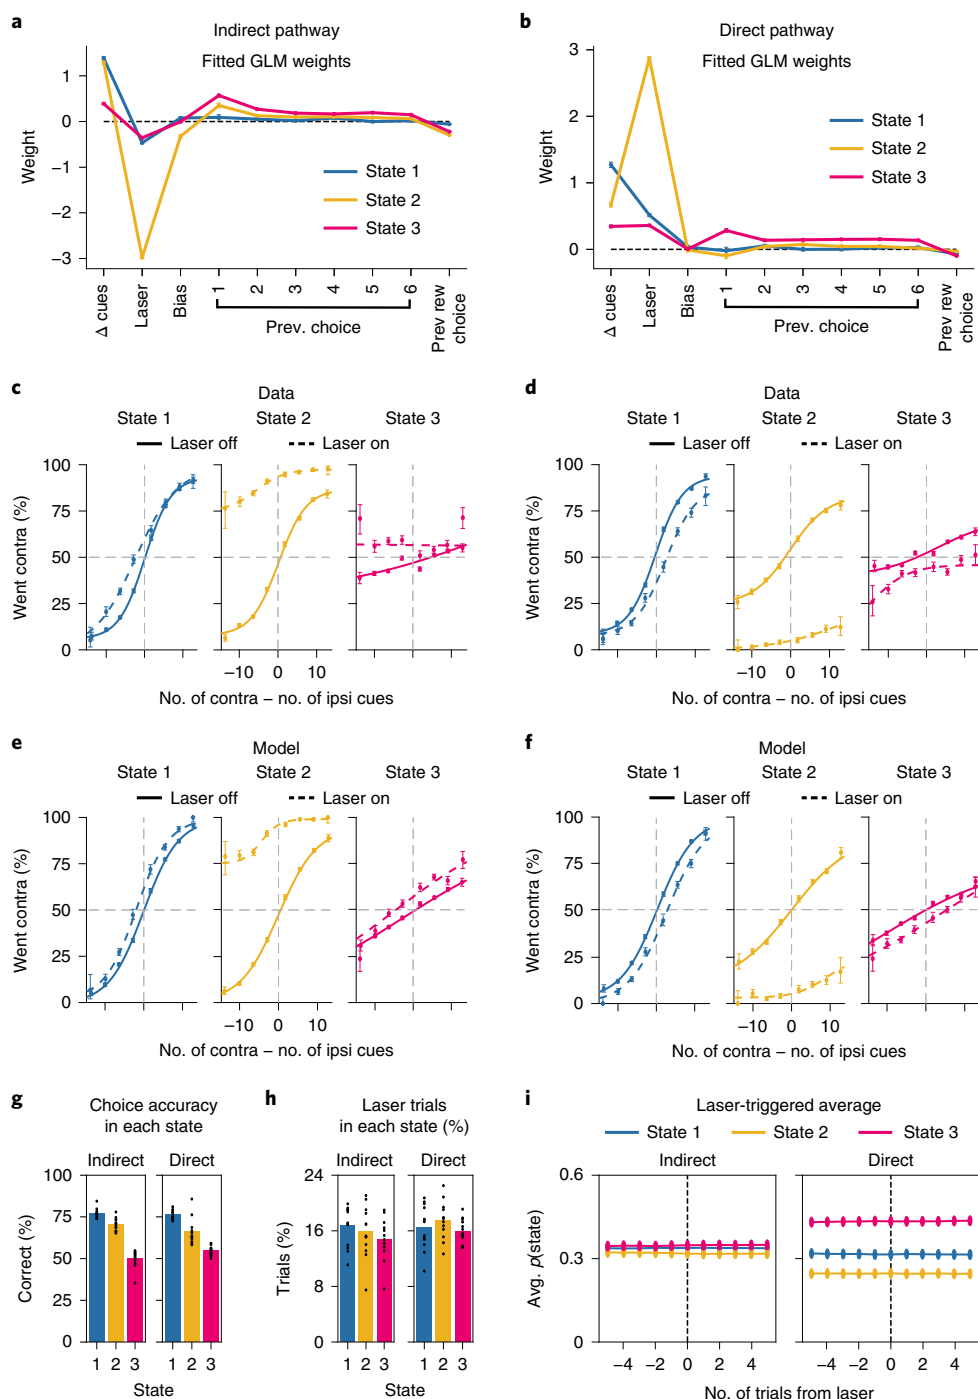
that a multistate model provides a good account of the animals' decision-making behavior during the evidence accumulation task.

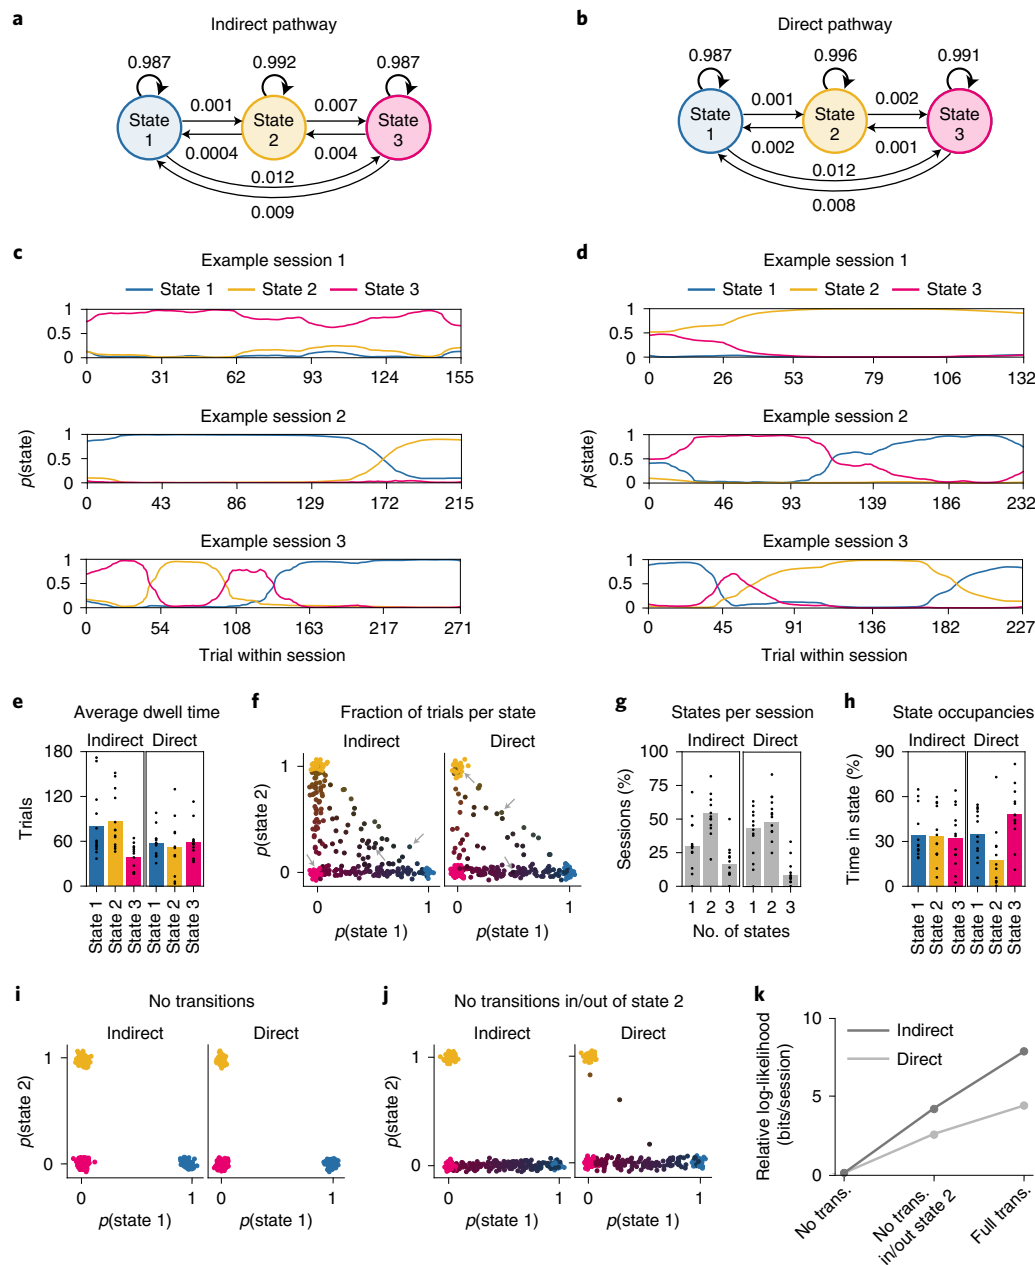
While the effect of the laser differed across states, the probability of being in a particular state did not change on or after trials with optical inhibition (Fig. 6i), implying that DMS pathway inhibition itself did not generate transitions between states. In addition, the fraction of trials with optical inhibition was equivalent across states (~15% of all trials in each state; Fig. 6h). This implies that the model did not identify states simply based on the presence of laser trials.

We obtained similar states when fitting the model to a combined dataset including all groups of mice (those receiving DMS indirect and direct pathway inhibition, as well as control mice receiving DMS illumination in the absence of NpHR; Extended Data Fig. 7f).

As when fitting each group separately, the combined fit revealed that both inhibition groups contained a single state with large weights on sensory evidence and the laser. In contrast, the control mice had small laser weights across all three states.

We also examined the results of fitting the four-state GLM-HMM (Extended Data Fig. 7d,e), given it had a slightly higher cross-validated log-likelihood than the three-state model (Extended Data Fig. 7a). In this case, the weights for states 1 and 2 were very similar to those in the three-state model; the key difference was that the choice history state (state 3 of the three-state model) was further subdivided into two states that differed in having a slight rightward versus a slight leftward bias. This suggests that while the model may uncover finer-grained structure in the data beyond three states





**Fig. 7 | Diversity across sessions in the timing and number of GLM-HMM state transitions.** **a**, Transition probabilities for the indirect pathway group. **b**, Same as **a** but for the direct pathway group. **c**, The posterior probability of being in each state for each trial for three example sessions from a mouse in the indirect pathway group. **d**, Same as **c** but for two mice from the direct pathway group. **e**, Dwell times showing the average consecutive number of trials that the mice spent in each state for mice with indirect (range 39–86 trials, average session length 202 trials) and direct (range 52–59 trials, average session length 185 trials) pathway inhibition. Black dots show averages for individual mice ( $n=13$  mice for both groups). **f**, The fraction of trials that the mice spent in each state in each session. Each dot represents an individual session ( $n=271$ , indirect pathway;  $n=266$ , direct pathway). Color coding reinforces the state composition of each session (for example, blue indicates that the mouse spent 100% of the session in state 1). A small amount of Gaussian noise was added to the position of each dot for visualization purposes. Gray arrows identify the example sessions shown in **c** and **d**. **g**, The fraction of sessions in which the mice entered one, two or all three states. Gray bars denote the average fraction of sessions for all mice. Black dots show averages for individual mice ( $n=13$  mice for both groups). **h**, Time spent in each state represented as a percentage of total trials for mice inhibited in the indirect and direct pathways. Colored bars denote the average state occupancies across all mice. Black dots show averages for individual mice ( $n=13$  mice both groups). **i**, Same as **f** except state assignments were obtained from a model in which the transition probabilities were restricted to disallow transitions between states (that is, all off-diagonal transition probabilities equal zero; Methods). **j**, Same as **f** except state assignments were obtained from a model in which transitions were disallowed between state 2 and the other states. **k**, Comparison of the cross-validated log-likelihood of the data when fitting GLM-HMMs with the reduced models from **i** and **j**, relative to the log-likelihood of the full model, in bps.

(Extended Data Fig. 7d,e), these states yield diminishing interpretive insight on the weighting of sensory evidence, choice history and DMS pathway inhibition across time.

**Diversity in timing and number of GLM-HMM state transitions.** The fitted transition matrix revealed a high probability of remaining in the same state across trials (Fig. 7a,b). These transition

probabilities produced a diversity in the timing and number of state transitions across sessions, which we visualized by calculating the maximum posterior probability of each state on each trial (Fig. 7c,d and 'GLM-HMM' in Methods). In some sessions, mice persisted in the same state, while in many sessions, mice visited two or even all three states (see example sessions in Fig. 7c,d, summaries of state occupancies across sessions in Fig. 7e–h and a summary of all individual mice in Supplementary Fig. 4). Average single-state dwell times ranged from 39 to 86 trials (Fig. 7g). This was shorter than the average session length of 194 trials, consistent with visits to multiple states per session.

While individual sessions were heterogeneous in terms of their state occupancies, averaged across sessions, the posterior probability of being in each state tended to be stable across trials (Fig. 7e and Extended Data Fig. 8a,b). Model simulations recapitulated these state transition characteristics, including dwell times and state occupancies (Extended Data Fig. 9), further indicating our model captures latent structure in our data.

One notable exception in the stability of posterior probabilities of each state across time was an increase in state 3 probability toward the end of a session (Extended Data Fig. 8a), potentially reflecting a decrease in task engagement related to reward satiety. Consistent with a relationship to satiety, within-session transitions into state 3 were associated with higher amounts of previously accumulated reward and higher preceding rates of reward (Extended Data Fig. 8c–f). In addition, while the posterior probability of each state showed minimal modulation surrounding a rewarded trial, the probability of state 3 was much more likely surrounding trials with excess travel (Extended Data Fig. 8g–j), an indicator of non-goal-directed movement and task disengagement. Indeed, the probability of state 3 gradually increased and decreased approximately 25 trials before and following excess travel trials, consistent with the average dwell time for state 3 (Fig. 7e).

Given the presence of sessions in which mice occupied a single state, we considered model variants that disallowed within-session state transitions. Our goal was to determine if these variant models could provide a better explanation of the data, or alternatively, if within-session state transitions are in fact an important structural feature for explaining the data. In one model variant, we disallowed transitions between states entirely (Fig. 7k). In the other, we tested the possibility that state 2, which is unique in the strength of its laser weight, captured a session-specific feature of inhibition by disallowing transitions in and out of that state (Fig. 7l). Using cross-validation, we found that neither alternative model explained the data as well as a model with unrestricted transitions (Fig. 7m), indicating that within-session transitions between states was an important feature of the model.

**Motor performance across GLM-HMM states.** Given the close relationship between excess travel and the posterior probability of state 3, we considered the possibility that other measures of motor behavior varied across states. We found that on trials without DMS pathway inhibition (Extended Data Fig. 10a–g,o–u), mice exhibited no obvious differences across states in velocity, *x*-position or view angle (Extended Data Fig. 10a–d,o–r). However, during state 3 relative to state 1 and 2, we observed an increased tendency in measures of non-goal-directed movements (Extended Data Fig. 10e–g,s–u). This is consistent with the higher probability of state 3 around trials with excess travel (Extended Data Fig. 8h,j), and the interpretation of state 3 as a task-disengaged state.

We also considered the possibility that DMS pathway inhibition had state-dependent effects on motor output (Extended Data Fig. 10h–n,v–bb). We observed limited effects of inhibition on velocity, per-trial standard deviation in view angle and distance traveled across all three states. However, similar to our cross-task comparisons (Extended Data Fig. 6j,k), DMS pathway inhibition produced

a small but opposing bias in average *x*-position (Extended Data Fig. 10j,x) and view angle (Extended Data Fig. 10k,y), which was greatest in the state with the largest laser weight (state 2; Fig. 6). This is consistent with our conclusions that the effects of DMS inhibition on behavior are state dependent, and that *x*-position and view angle are closely linked indicators of choice in the context of VR-based T-maze tasks (Extended Data Fig. 3g–j).

## Discussion

Our findings indicate that the opposing contributions of DMS pathways to movement are minimal in the absence of a decision (Fig. 1), while the pathways provide large and opponent contributions to decision-making. Moreover, this contribution depends on the demands of a task (Fig. 2), as the effect of inhibition is much larger during decisions that require evidence accumulation relative to control tasks with weaker cognitive requirements yet similar sensory features and motor requirements (Fig. 3). The GLM-HMM further revealed that even within the evidence accumulation task, the contribution of DMS pathways to choice is not fixed. For example, DMS pathways have little contribution when mice pursue a strategy of repeating previous choices during the evidence accumulation task (Fig. 6). Thus, together our findings imply that opposing contributions of DMS pathways to behavior are dependent on task demands and internal state.

**Cross-task differences in effects of dorsomedial striatum pathway inhibition.** Our finding that DMS activity contributes to the evidence accumulation task, but not to task variants with weaker cognitive demands, is broadly consistent with previous work based on lesions, pharmacology, and recordings implicating DMS in short-term memory and the dynamic comparison of the value of competing options<sup>18–21,24,33–36</sup>. But then why have most previous optogenetic pathway-specific manipulations emphasized an opposing role for DMS pathways in the direct control of motor output<sup>2–4,8,11,16</sup>, rather than on decision-making? Prior work has overwhelmingly relied on the synchronous activation of striatal pathways, as opposed to the inhibition used here. While DMS pathway activation is sufficient to bias movements such as spontaneous rotations, we observed relatively little impact of inhibition on behavior in the absence of a decision (Fig. 1). Taken together with previous work, our findings may thus imply limits in the use of artificial activations in assessing striatal pathway function. Our results may also imply that DMS pathways would not necessarily display opposing correlates of movements<sup>14,17,37–40</sup>, but rather opposing correlates of a decision process<sup>12,13,19,21,33,36</sup>.

While our observations are consistent with the classic view of opposing contributions of striatal pathways to behavior<sup>1</sup>, several prominent studies have instead challenged this view by reporting non-opposing behavioral effects of activating each pathway<sup>41–47</sup>. This may be because the pathways of a specific striatal subregion only exert opposing control on behavior in a specific context, for example, during cognitively demanding decision-making as shown here for the DMS, or during an interval timing task that requires the proactive suppression of actions, as shown for the dorsolateral striatum<sup>27</sup>. Along these lines, our comparison to NAc pathways, where inhibition produced weak effects on behavior in similar directions (Fig. 3p), may imply that we have not discovered the context in which NAc pathways have opposing contributions to behavior.

We designed our T-maze tasks to have very similar sensory features and identical motor requirements, and yet very different cognitive demands, as assessed by task accuracy. That being said, the sensory features were not identical. Therefore, while unlikely, we cannot rule out that the subtle sensory differences contributed to the cross-task differences in the effect of pathway inhibition. A future direction would be to maintain an identical sensory

environment across tasks and instead change the decision-making rule to be more cognitively demanding.

**Within-task changes in effects of dorsomedial striatum pathway inhibition.** Complimenting our cross-task comparison, we reveal the new insight that mice occupy time-varying latent states within a single task and that the contribution of DMS pathways to choice depends on the internal state of mice. The application of a GLM–HMM was critical in uncovering this feature of behavior, allowing the unsupervised discovery of latent states that differ in how external covariates were weighted to influence a choice<sup>31,48,49</sup>. This provided two insights on the contributions of DMS pathways to behavior.

First, the impact of DMS inhibition was diminished when mice occupied a task-disengaged state in which choice history heavily predicted decisions, while conversely, the impact of DMS inhibition was accentuated when mice occupied a task-engaged state in which sensory evidence strongly influenced choice (Fig. 6). This strengthens our conclusion from the cross-task comparison, which is that DMS pathways have a greater contribution to behavior when actively accumulating evidence toward a decision.

Second, mice occupied two similar task-engaged states that were modestly distinguished in overall accuracy (Fig. 6g) and prominently distinguished by the influence of DMS inhibition on choice (Fig. 6a–d). While transitions between these two states were relatively rare on the same day, there were days that included both states (Fig. 7). The discovery of these two states leads to the intriguing suggestion that mice are capable of accumulating evidence toward a decision in at least two neurally distinct manners: one that depends on each DMS pathway (state 2), and another that does not (state 1). This may relate to demonstrations that neural circuits have substantial capacity for compensation to perturbations<sup>30,51</sup>, and our modeling approach may provide a new avenue for the identification of such compensatory mechanisms on relatively short timescales.

While our work focused on the three-state GLM–HMM, our conclusions do not depend on assuming exactly three states. In fact, the cross-validated log-likelihood of our data is higher for four states than three. Yet the conclusions from the four-state model were similar to those from the three-state model (compare Fig. 6a,b to Extended Data Fig. 7d,e), and additional gains in log-likelihood decrease for larger numbers of states. Nevertheless, it will be important for future work to compare the GLM–HMM framework used here, which assumes discrete states, to models that assume continuously varying states<sup>32,52</sup>.

Altogether, our studies provide new perspectives on the neural mechanisms by which DMS pathways exert opponent control over behavior, with particular emphasis on the importance of accounting for task demands, internal state and associated behavioral strategies when assessing neural mechanisms. To this end, we expect our behavioral and computational frameworks to be of broad utility in uncovering the neural substrates of decision-making in a wide range of settings.

## Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41593-022-01021-9>.

Received: 12 June 2020; Accepted: 21 January 2022;

Published online: 7 March 2022

## References

- Alexander, G. E. & Crutcher, M. D. Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends Neurosci.* **13**, 266–271 (1990).
- Kravitz, A. V. et al. Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature* **466**, 622–626 (2010).
- Roseberry, T. K. et al. Cell-type-specific control of brainstem locomotor circuits by basal ganglia. *Cell* **164**, 526–537 (2016).
- Bartholomew, R. A. et al. Striatonigral control of movement velocity in mice. *Eur. J. Neurosci.* **43**, 1097–1110 (2016).
- Bakshurin, K. I. et al. Opponent regulation of action performance and timing by striatonigral and striatopallidal pathways. *eLife* **9**, e54831 (2020).
- Lobo, M. K. et al. Cell-type-specific loss of BDNF signaling mimics optogenetic control of cocaine reward. *Science* **330**, 385–390 (2010).
- Kravitz, A. V., Tye, L. D. & Kreitzer, A. C. Distinct roles for direct and indirect pathway striatal neurons in reinforcement. *Nat. Neurosci.* **15**, 816–818 (2012).
- Yttri, E. A. & Dudman, J. T. Opponent and bidirectional control of movement velocity in the basal ganglia. *Nature* **533**, 402–406 (2016).
- Tai, L.-H., Lee, A. M., Benavidez, N., Bonci, A. & Wilbrecht, L. Transient stimulation of distinct subpopulations of striatal neurons mimics changes in action value. *Nat. Neurosci.* **15**, 1281–1289 (2012).
- Nonomura, S. et al. Monitoring and updating of action selection for goal-directed behavior through the striatal direct and indirect pathways. *Neuron* **99**, 1302–1314 (2018).
- Lee, J., Wang, W. & Sabatini, B. L. Anatomically segregated basal ganglia pathways allow parallel behavioral modulation. *Nat. Neurosci.* **23**, 1388–1398 (2020).
- Cui, L. et al. Asymmetrical choice-related ensemble activity in direct and indirect-pathway striatal neurons drives perceptual decisions. Preprint at *bioRxiv* <https://doi.org/10.1101/2021.11.16.468594> (2021).
- Tang, Y. et al. Opposing regulation of short-term memory by basal ganglia direct and indirect pathways that are coactive during behavior. Preprint at *bioRxiv* <https://doi.org/10.1101/2021.12.15.472735> (2021).
- Parker, J. G. et al. Diametric neural ensemble dynamics in parkinsonian and dyskinetic states. *Nature* **557**, 177–182 (2018).
- Chen, Z. et al. Direct and indirect pathway neurons in ventrolateral striatum differentially regulate licking movement and nigral responses. *Cell Rep.* **37**, 109847 (2021).
- Lee, H. J. et al. Activation of direct and indirect pathway medium spiny neurons drives distinct brain-wide responses. *Neuron* **91**, 412–424 (2016).
- London, T. D. et al. Coordinated ramping of dorsal striatal pathways preceding food approach and consumption. *J. Neurosci.* **38**, 3547–3558 (2018).
- Balleine, B. W., Delgado, M. R. & Hikosaka, O. The role of the dorsal striatum in reward and decision-making. *J. Neurosci.* **27**, 8161–8165 (2007).
- Yartsev, M. M., Hanks, T. D., Yoon, A. M. & Brody, C. D. Causal contribution and dynamical encoding in the striatum during evidence accumulation. *eLife* **7**, e34929 (2018).
- Lau, B. & Glimcher, P. W. Value representations in the primate striatum during matching behavior. *Neuron* **58**, 451–463 (2008).
- Ding, L. & Gold, J. I. Separate, causal roles of the caudate in saccadic choice and execution in a perceptual decision task. *Neuron* **75**, 865–874 (2012).
- Barnes, T. D., Kubota, Y., Hu, D., Jin, D. Z. & Graybiel, A. M. Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* **437**, 1158–1161 (2005).
- Yin, H. H. et al. Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nat. Neurosci.* **12**, 333–341 (2009).
- Akhlaghpour, H. et al. Dissociated sequential activity and stimulus encoding in the dorsomedial striatum during spatial working memory. *eLife* **5**, e19507 (2016).
- Pinto, L. et al. An Accumulation-of-evidence task using visual pulses for mice navigating in virtual reality. *Front. Behav. Neurosci.* **12**, 36 (2018).
- Owen, S. F., Liu, M. H. & Kreitzer, A. C. Thermal constraints on in vivo optogenetic manipulations. *Nat. Neurosci.* **22**, 1061–1065 (2019).
- Cruz, B. F., Soares, S. & Paton, J. J. Striatal circuits support broadly opponent aspects of action suppression and production. Preprint at *bioRxiv* <https://doi.org/10.1101/2020.06.30.180539> (2020).
- Kupchik, Y. M. et al. Coding the direct/indirect pathways by D1 and D2 receptors is not valid for accumbens projections. *Nat. Neurosci.* **18**, 1230–1232 (2015).
- Bengio, Y. & Frasconi, P. An input output HMM architecture. *Adv. Neural Inf. Process. Syst.* **7**, 427–234 (1994).
- Escola, S., Fontanini, A., Katz, D. & Paninski, L. Hidden Markov models for the stimulus-response relationships of multistate neural systems. *Neural Comput.* **23**, 1071–1132 (2011).
- Calhoun, A. J., Pillow, J. W. & Murthy, M. Unsupervised identification of the internal states that shape natural behavior. *Nat. Neurosci.* **22**, 2040–2049 (2019).
- Ashwood, Z. C. et al. Mice alternate between discrete strategies during perceptual decision-making. *Nat. Neurosci.* **25**, 201–212 (2022).



33. Donahue, C. H., Liu, M. & Kreitzer, A. C. Distinct value encoding in striatal direct and indirect pathways during adaptive learning. Preprint at *bioRxiv* <https://doi.org/10.1101/277855> (2018).
34. Shin, J. H., Kim, D. & Jung, M. W. Differential coding of reward and movement information in the dorsomedial striatal direct and indirect pathways. *Nat. Commun.* **9**, 404 (2018).
35. Delevich, K., Hoshal, B., Collins, A. G. & Wilbrecht, L. Choice suppression is achieved through opponent but not independent function of the striatal indirect pathway in mice. Preprint at *bioRxiv* <https://doi.org/10.1101/675850> (2020).
36. Frank, M. J. & Badre, D. Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: computational analysis. *Cereb. Cortex* **22**, 509–526 (2012).
37. Cui, G. et al. Concurrent activation of striatal direct and indirect pathways during action initiation. *Nature* **494**, 238–242 (2013).
38. Barbera, G. et al. Spatially compact neural clusters in the dorsal striatum encode locomotion relevant information. *Neuron* **92**, 202–213 (2016).
39. Sippy, T., Lapray, D., Crochet, S. & Petersen, C. C. H. Cell-type-specific sensorimotor processing in striatal projection neurons during goal-directed behavior. *Neuron* **88**, 298–305 (2015).
40. Jin, X., Tecuapetla, F. & Costa, R. M. Basal ganglia subcircuits distinctively encode the parsing and concatenation of action sequences. *Nat. Neurosci.* **17**, 423–430 (2014).
41. Soares-Cunha, C. et al. Activation of D2 dopamine receptor-expressing neurons in the nucleus accumbens increases motivation. *Nat. Commun.* **7**, 1–11 (2016).
42. Cole, S. L., Robinson, M. J. F. & Berridge, K. C. Optogenetic self-stimulation in the nucleus accumbens: D1 reward versus D2 ambivalence. *PLoS ONE* **13**, e0207694 (2018).
43. Vicente, A. M., Galvão-Ferreira, P., Tecuapetla, F. & Costa, R. M. Direct and indirect dorsolateral striatum pathways reinforce different action strategies. *Curr. Biol.* **26**, R267–R269 (2016).
44. Tecuapetla, F., Jin, X., Lima, S. Q. & Costa, R. M. Complementary contributions of striatal projection pathways to action initiation and execution. *Cell* **166**, 703–715 (2016).
45. Geddes, C. E., Li, H. & Jin, X. Optogenetic editing reveals the hierarchical organization of learned action sequences. *Cell* **174**, 32–43 (2018).
46. Wang, L., Rangarajan, K. V., Gerfen, C. R. & Krauzlis, R. J. Activation of striatal neurons causes a perceptual decision bias during visual change detection in mice. *Neuron* **98**, 669 (2018).
47. Peak, J., Chieng, B., Hart, G. & Balleine, B. W. Striatal direct and indirect pathway neurons differentially control the encoding and updating of goal-directed learning. *eLife* **9**, e58544 (2020).
48. Eldar, E., Morris, G. & Niv, Y. The effects of motivation on response rate: a hidden semi-Markov model analysis of behavioral dynamics. *J. Neurosci. Methods* **201**, 251–261 (2011).
49. Ahilan, S. et al. Learning to use past evidence in a sophisticated world. *PLoS Comput. Biol.* **15**, e1007093 (2019).
50. Goshen, I. et al. Dynamics of retrieval strategies for remote memories. *Cell* **147**, 678–689 (2011).
51. Fetsch, C. R. et al. Focal optogenetic suppression in macaque area MT biases direction discrimination and decision confidence, but only transiently. *eLife* **7**, e36523 (2018).
52. Roy, N. A. et al. Extracting the dynamics of behavior in sensory decision-making experiments. *Neuron* **109**, 597–610 (2021).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2022

## Methods

**Animals.** For optogenetic experiments we used both male and female transgenic mice on heterozygous backgrounds, aged 2–6 months of age, from the following three strains backcrossed to a C57BL/6J background (Jackson Laboratory, 000664) and maintained in-house: Drd1-Cre ( $n=45$ , EY262Gsat, MMRRC-UCD), Drd2-Cre ( $n=24$ , ER44Gsat, MMRRC-UCD) and A2a-Cre ( $n=18$ , KG139Gsat, MMRRC-UCD). An additional 35 mice were excluded from all optogenetic analyses due to failed task acquisition ( $n=11$  mice) or failed viral/fiber-optic targeting of DMS ( $n=8$ ) or NAc ( $n=16$ ). An additional four Drd1-Cre mice, three A2a-Cre mice and two Drd2-Cre mice were used for electrophysiological characterization of NpHR-mediated inhibition, or fluorescence in situ hybridization (FISH) characterization of Cre expression profiles. FISH experiments also utilized 2 Drd1a-tdTomato mice (Jax, 016204). Mice were co-housed with same-sex littermates and maintained on a 12-h light–12-h dark cycle. All surgical procedures and behavioral training occurred in the dark cycle. All procedures were conducted in accordance with National Institute of Health guidelines and were reviewed and approved by the Institutional Animal Care and Use Committee at Princeton University.

**Surgical procedures.** All mice underwent sterile stereotaxic surgery to implant ferrule-coupled optical fibers (Newport, 200  $\mu\text{m}$  core, 0.37 NA) and a custom titanium headplate for head fixation under isoflurane anesthesia (5% induction, 1.5% maintenance). Mice received a preoperative antibiotic injection of Baytril (enrofloxacin; 5 mg per kg body weight) intramuscularly, as well as analgesia pre-operatively and 24 h later in the form of subcutaneous meloxicam injections (2 mg per kg body weight). A microsyringe pump controlling a 10- $\mu\text{l}$  glass syringe (Nanofil) was used to bilaterally deliver virus targeted to either the DMS (0.74 mm anterior, 1.5 mm lateral, –3.0 mm ventral) or the NAc (1.3 mm anterior, 1.2 mm lateral, –4.7 mm ventral). For optogenetic inhibition, the following viruses were used: AAV2/5-eF1a-DIO-eNpHR3.0-EYFP-WPRE-hGH (UPenn,  $1.3 \times 10^{13}$  parts per ml) or AAV2/5-eF1a-DIO-eNpHR3.0-EYFP-WPRE-hGH (PNI Viral Core,  $2.2 \times 10^{14}$  parts per ml; 1:5 dilution). For FISH experiments, AAV2/5-eF1a-DIO-EYFP-hGHpA (PNI Viral Core,  $6.0 \times 10^{13}$  parts per ml) was used to label D1R<sup>+</sup> and D2R<sup>+</sup> neurons in D1R-Cre and A2a-Cre transgenic lines. In all experiments, virus was delivered at a rate of 0.2  $\mu\text{l min}^{-1}$  for a total volume of 0.3–0.7  $\mu\text{l}$  in the DMS, or 0.3–0.4  $\mu\text{l}$  in the NAc. To accommodate patch fiber coupling, optical fibers were implanted at angles (DMS: 15°, NAc 10°). The following coordinates (in unrotated space) were targeted: DMS, 0.74 mm anterior, 1.1 mm lateral, –3.6 mm ventral; NAc: 1.3 mm anterior, 0.55 mm lateral, –5.0 mm ventral. Fibers were fixed to the skull using dental cement. Mice were allowed to recover and closely monitored for 5 d before beginning water restriction and behavioral training.

**Optrode recording for NpHR validation.** Following the surgical procedures described above, Cre-dependent NpHR was virally delivered bilaterally to the DMS of mice ( $n=3$  A2a-Cre;  $n=2$  D1R-Cre) via small (~300  $\mu\text{m}$ ) craniotomies made using a carbide drill (Extended Data Fig. 1a). The craniotomies were filled with a small amount of silicon adhesive (Kwik-Sil, World Precision Instruments) and then covered with ultraviolet-curing optical adhesive (Norland Optical Adhesive 61), while a custom-designed headplate for head fixation was cemented to the skull. Following a recovery period of ~4 weeks, awake mice were head-fixed on a plastic running wheel attached to a breadboard via Thorlabs posts and holders, which was fixed immediately adjacent to a stereotaxic setup (Kopf) enclosed within a Faraday cage (Extended Data Fig. 1b). Silicon and optical adhesive was removed from the craniotomies and a 32-channel, single-shank silicon probe (A1x32-Poly3, NeuroNexus) coupled to a tapered optical fiber (65  $\mu\text{m}$ , 0.22 NA) was stereotactically inserted under visual guidance of a stereoscope and allowed to stabilize for ~30 min. Signals were acquired at 20 kHz using a digital headstage amplifier (RHD2132, Intan) connected to an RHD USB data acquisition board (C3100, Intan). A screw implanted over the cerebellum served as ground. Continuous signal was imported into MATLAB for referencing to a local probe channel and high-pass filtering at 200 Hz, and then imported into Offline Sorter v3 (Plexon) for spike thresholding and single-unit sorting. During recording, the optical fiber was connected via a patch cable to a 532-nm laser, which was triggered by a transistor–transistor logic (TTL) pulse sent by a pulse generator controlled by a computer running Spike2 software. TTL pulse times were copied directly to the RHD USB data acquisition board. Laser sweeps consisted of 40 deliveries of 5 s of light (5 mW, measured from fiber tip), separated by 15-s ITIs. One to three recordings were made at different depths within a single probe penetration (minimum separation of 300  $\mu\text{m}$ ), with each hemisphere receiving one to three penetrations at different mediolateral or anteroposterior coordinates. For recordings in mice carried out over multiple days, craniotomies were filled with Kwik-Sil and covered with silicone elastomer between recordings (Kwik-Cast, World Precision Instruments).

**Virtual reality behavior.** *Virtual reality setup.* Mice were head-fixed over an 8-inch Styrofoam ball suspended by compressed air (~60 p.s.i.) facing a custom-built Styrofoam toroidal screen spanning a visual field of 270° horizontally and 80° vertically. The setup was enclosed within a custom-designed cabinet built

from optical rails (Thorlabs) and lined with sound-attenuating foam sheeting (McMaster-Carr). A DLP projector (Optoma HD141X) with a refresh rate of 120 Hz projected the VR environment onto the toroidal screen (Fig. 1e).

An optical flow sensor (ADNS-3080 APM2.6), located beneath the ball and connected to an Arduino Due, ran custom code to transform real-world ball rotations into virtual-world movements (<https://github.com/sakoay/AccumTowersTools/tree/master/OpticalSensorPackage/>) within the MATLAB-based ViRMEn<sup>53</sup> software engine (<http://pni.princeton.edu/pni-software-tools/virmen/>). The ball and sensor of each VR rig were calibrated such that ball displacements ( $dX$  and  $dY$ , where  $X$  (and  $Y$ ) are parallel to the anteroposterior (and mediolateral) axes of the mouse) produced translational displacements proportional to ball circumference in the virtual environment of equal distance in corresponding  $X$  and  $Y$  axes. The  $y$ -velocity of the mouse is given by  $\sqrt{dY^2/dt}$ , where  $dt$  is the elapsed time from the previous sampling of the sensor. The virtual view angle of mice was obtained by first calculating the current displacement angle as:  $\omega = \text{atan2}(-dX \times \sin(dY), |dY|)$ . Then the rate of change of view angle ( $\theta$ ) for each sampling of the sensor is given by equation (1):

$$\frac{d\theta}{dt} = \sin(\omega) \times \min\left(e^{1.4|\omega|^{1.2}} - 1, \frac{\pi}{2}\right) - \theta \quad (1)$$

This exponential function was tuned to (1) minimize the influence of small ball displacements and thus stabilize virtual-world trajectories, and (2) increase the influence of large ball displacements to allow sharp turns into the maze arms<sup>25</sup>.

Reward and whisker air puffs were delivered by sending a TTL pulse to solenoid valves (NResearch), which were generated according to behavioral events on the ViRMEn computer. Each TTL pulse resulted in either the release of a drop of reward (~4–8  $\mu\text{l}$  of 10% sweetened condensed milk in water, vol/vol) to a lick tube, or the release of air flow (40 ms, 15 p.s.i.) to an air puff cannula (Ziggy's Tubes and Wires, 16 gauge) directed to the left and right whisker pads from the rear position. The ViRMEn computer also controlled TTL pulses sent directly to a 532-nm DPSS laser (Shanghai, 200 mW).

**Behavioral shaping.** Following post-surgical recovery, over the course of 4–7 d, mice were extensively handled while gradually restricting water intake to an allotted volume of 1–2 ml per day. Throughout water restriction, mice were closely monitored to ensure no signs of dehydration were present and that body mass was at least 80% of the pre-restriction value. Mice were then introduced to the VR setup where behavior was shaped to perform the accumulation of evidence task as previously described in detail<sup>25</sup> (Supplementary Fig. 5a) or the permanent cues (control 2) task (Supplementary Fig. 5f). We tested a total of 32, 34 and 20 mice in the accumulation of evidence, no distractors (control 1) and permanent cues (control 2) tasks, respectively. No mice received optogenetic testing in all three tasks, but 7 mice received optogenetic testing in the accumulation of evidence and no distractors task, and 19 mice received optogenetic testing in the no distractors and permanent cues tasks (Supplementary Table 1).

Shaping followed a similar progression in both tasks. In the first four shaping mazes of both procedures, a visual guide located in the rewarded arm was continuously visible, and the maze stem was gradually extended to a final length of 300 cm (Supplementary Fig. 5a,f). In mazes 5–7 of the evidence accumulation shaping procedure (Supplementary Fig. 5a), the visual guide was removed and the cue region was gradually decreased to 200 cm, thus introducing the full 100-cm delay region of the testing mazes. The same shift to a 200-cm cue region and 100-cm delay region occurred in mazes 5–6 of the permanent cues shaping procedure, but without removing the visibility of the visual guide (Supplementary Fig. 5f). In mazes 8–9 of evidence accumulation shaping, distractor cues were introduced to the non-rewarded maze side with increasing frequency (mean side ratio (s.d.) of rewarded:non-rewarded side cues of 8.3:0.7 to 8.0:1.6 m<sup>-1</sup>). Distractor cues were similarly introduced with increasing frequency in mazes 6–8 of the permanent cues shaping procedure, while the visual guide was removed in mazes 7 and 8. In all evidence accumulation shaping mazes (mazes 1–9), cues were only made visible when mice were 10-cm from the cue location and remained visible until trial completion. In the final evidence accumulation testing mazes (maze 10 and 11), cues were made transiently visible (200 ms) after first presentation (10 cm from cue location), while the mean side ratio of rewarded:non-rewarded side cues changed from 8.0:1.6 (Supplementary Fig. 5a; maze 10) to 7.7:2.3 m<sup>-1</sup> (Supplementary Fig. 5a; maze 11). In contrast, throughout all shaping mazes (1–6) and testing mazes (7–8) of the permanent cues task, cues were visible from the onset of a trial.

The median number of sessions to reach the first evidence accumulation testing maze (maze 10) was 22 sessions, while the mean number of sessions was  $23.0 \pm 0.8$  (Supplementary Fig. 5b,c). Mice typically spent between 2 and 5 sessions on each shaping maze before progressing to the next, with performance increasing or remaining stable throughout (Supplementary Fig. 5d,e; maze 9: 74.1%  $\pm$  9.8% correct). The median number of sessions to reach the first permanent cues (control 2) testing maze (maze 7) was 17 sessions, while the mean number of sessions was  $18.0 \pm 1.5$  (Supplementary Fig. 5g–j). Mice typically spent between 2 and 4 sessions on each shaping maze before progressing to the next, with performance increasing or remaining largely stable throughout (Supplementary Fig. 5g–j; maze 6: 87%  $\pm$  4.3% correct).

**Optogenetic testing mazes.** The evidence accumulation task took place in a 330-cm-long virtual T-maze with a 30-cm start region (−30 to 0 cm), followed by a 200-cm cue region and finally a 100-cm delay region (Fig. 2a, black, left). While navigating the cue region of the maze, mice were transiently presented with high-contrast visual cues (wall-sized ‘towers’) on either maze side, which were also paired with a mild air puff (15 p.s.i., 40 ms) to the corresponding whisker pad. The side containing the greater number of cues indicated the future rewarded arm. A left or right choice was determined when mice crossed an  $x$ -position threshold > 15 cm, which was only possible within one of the maze arms (the width of choice arms was  $\pm 25$  cm relative to the center of the maze stem). Mice received reward (−4–8  $\mu$ l of 10% vol/vol sweetened condensed milk in drinking water) followed by a 3-s ITI after turning to the correct arm at the end of the maze, while incorrect choices were indicated by a tone followed by a 12-s ITI. In each trial, the position of cues was drawn randomly from a spatial Poisson process with a rate of  $8.0\text{ m}^{-1}$  for the rewarded side and  $1.6\text{ m}^{-1}$  for the non-rewarded side (Supplementary Fig. 5a; maze 10) or  $7.3:2.3\text{ m}^{-1}$  (Supplementary Fig. 5a; maze 11). Note that only maze 10 data were used for cross-task comparisons of optogenetic effects with permanent cues and no distractors control tasks to precisely match cue presentation statistics (Figs. 2 and 3 and Extended Data Figs. 3 and 6). Visual cues (and air puffs) were presented when mice were 10 cm away from their drawn location and ended 200 ms (or 40 ms) later. Cue positions on the same side were also constrained by a 12-cm refractory period. Each session began with warm-up trials of a visually guided maze (Supplementary Fig. 5a; maze 4), with mice progressing to the evidence accumulation testing maze after ten trials (or until accuracy reached 85% correct). During performance of the testing maze, if accuracy fell below 55% over a 40-trial running window, mice were transitioned to an easier maze in which cues were presented only on the rewarded side and did not disappear following presentation (Supplementary Fig. 5a; maze 7). These ‘easy blocks’ were limited to ten trials, after which mice returned to the main testing maze regardless of performance. Behavioral sessions lasted for ~1 h and typically consisted of ~150–200 trials.

All features of the ‘no distractors’ (control 1) task (Fig. 2b, magenta; Supplementary Fig. 5a, maze 12) were identical to the evidence accumulation task (Supplementary Fig. 5a, maze 10) except that (1) distractor cues were removed from the non-rewarded side, and (2) a distal visual guide located in the rewarded arm was transiently visible during the cue region (0–200 cm).

All features of the ‘permanent cues’ (control 2) task (Fig. 2b, cyan; Supplementary Fig. 5g, maze 8) were identical to the evidence accumulation task except that reward and non-reward side visual cues were made permanently visible from trial onset. As in the evidence accumulation task, whisker air puffs were only delivered when mouse position was 10 cm from visual cue location. Note that mice underwent optogenetic testing on two permanent cues mazes (mazes 7 and 8). Maze 8 shared identical reward to non-reward side cue statistics ( $8.0:1.6\text{ m}^{-1}$ ) as maze 10 of the evidence accumulation task. Therefore, for all cross-task comparisons of optogenetic inhibition, only data from these mazes were analyzed (Figs. 2 and 3 and Extended Data Figs. 3–6).

To discourage side biases, in all tasks we used a previously implemented debiasing algorithm<sup>25</sup>. This was achieved by changing the underlying probability of drawing a left or a right trial according to a balanced method described in detail elsewhere<sup>25</sup>. In brief, the probability of drawing a right trial,  $p_R$ , is given by equation (2):

$$p_R = \frac{\sqrt{e_R}}{(\sqrt{e_R} + \sqrt{e_L})} \quad (2)$$

where  $e_R$  and  $e_L$  are the weighted average of the fraction of errors the mouse has made in the past 40 right (and left) trials. The weighting for this average is based on a half-Gaussian with  $\sigma = 20$  trials in the past, which ensures that most recent trials have larger weight on the debiasing algorithm. To discourage the generation of sequences of all-right (or all-left) trials, we capped  $\sqrt{e_R}$  and  $\sqrt{e_L}$  to be within the range of (0.15, 0.85). Because the empirical fraction of drawn right trials could deviate from  $p_R$ , particularly when the number of trials is small, we applied an additional pseudorandom drawing prescription to  $p_R$ . Specifically, if the empirical fraction of right trials (calculated using a  $\sigma = 60$  trials half-Gaussian weighting window) was above  $p_R$ , right trials were drawn with probability  $0.5p_R$ , whereas if this fraction was below  $p_R$ , right trials were drawn with probability  $0.5(1 + p_R)$ .

**Virtual corridor.** Following shaping in the behavioral tasks, mice were transitioned to free navigation in a virtual corridor arena in the same VR apparatus described above. The virtual corridor was 6 cm in diameter and 330 cm in effective length (Fig. 1e,f). This included a start region (−10 to 0 cm), a reward location (310 cm) in which mice received 4  $\mu$ l of 10% vol/vol sweetened condensed milk in drinking water, and a teleportation region (320 cm) in which mice were transported back to the start region following a variable ITI with mean of 2 s. Mice were otherwise allowed to freely navigate the virtual corridor during ~70-min sessions. The virtual environment was controlled by the ViRMEn software engine, with real-to-virtual-world movement transformations as described above.

**Optogenetics during virtual reality behavior.** According to a previously published protocol<sup>54</sup>, optical fibers (200  $\mu$ m, 0.37 NA) were chemically etched using 48%

hydrofluoric acid to achieve tapered tips 1.5–2 mm (DMS-targeted) or 1–1.5 mm (NAC-targeted) in length. Following behavioral shaping in VR (and >6 weeks of viral expression), mice underwent optogenetic testing. On alternate daily sessions, optical fibers in the left or right hemisphere were unilaterally coupled to a 532-nm DPSS laser (Shanghai, 200 mW) via a multimode fiber-optic patch cable (PFP, 62.5  $\mu$ m). On a random subset of trials (10–30%), mice received unilateral laser illumination (5 mW, measured from patch cable) that was restricted to the first passage through 0–200 cm of the virtual corridor (Fig. 1 and Extended Data Fig. 2), or the cue region (0–200 cm) of each T-maze decision-making task (Fig. 3). The laser was controlled by TTL pulses generated using a National Instruments DAQ card on a computer running the ViRMEn-based virtual environment.

**Conditioned place preference test.** Mice underwent a real-time conditioned place preference (CPP) test with bilateral optogenetic inhibition paired to one side of a two-chamber apparatus (Supplementary Fig. 3). The CPP apparatus consisted of a rectangular Plexiglass box with two chambers (29 cm  $\times$  25 cm) separated by a clear portal in the center. The same gray, plastic flooring was used for both chambers, but each chamber was distinguished by vertical or horizontal black and white bars on the chamber walls. During a baseline test, mice were placed in the central portal while connected to patch cables coupled to an optical commutator (Doric) and were allowed to freely move between both sides for 5 min. In a subsequent 20-min test, mice received continuous, bilateral optogenetic inhibition (532 nm, 5 mW) when located in one of the two chamber sides (balanced across groups). Video tracking, TTL triggering and data analysis were carried out using EthoVision software (Noldus). Mice who displayed a bias for one chamber side greater than 45 s during the baseline test were excluded from the analysis.

**Behavior analyses. Data selection.** See Supplementary Table 1 for a list of all mice with optogenetic testing data in the virtual corridor, the accumulation of evidence, no distractors or permanent cues tasks. The following describes the trial selection criteria for inclusion in analyses throughout.

For cross-task comparisons (Figs. 2 and 3 and Extended Data Figs. 3–6), we analyzed only trials from evidence accumulation maze 10 (Supplementary Fig. 5a), ‘no distractors’ maze 12 (Supplementary Fig. 5a,g) and ‘permanent cues’ maze 8 (Supplementary Fig. 5g), which each followed matching rewarded and non-rewarded side cue probability statistics (saved for the by-design absence of non-rewarded cues in the ‘no distractors’ control task). For model-based analyses of the evidence accumulation task (Figs. 4–7 and Extended Data Figs. 7–10), both maze 10 and maze 11 data were included, which differed only modestly in the side ratio of reward to non-reward side cues (Supplementary Fig. 5a; ~50% of trials were maze 10 or 11). In all tasks and all analyses throughout, we removed initial warm-up blocks (Supplementary Fig. 5a; maze 4, approximately 5–15% of total trials). For model-based analyses of the evidence accumulation task (Figs. 4–7 and Extended Data Figs. 7–10), we included interspersed ‘easy blocks’ capped at ten trials in length (Supplementary Fig. 5a; maze 7, see above). These trial blocks comprised approximately ~5% of total trials, were included to avoid gaps in trial history, and were treated identically to the main evidence accumulation mazes by the models. These trials were removed from cross-task comparisons of optogenetic inhibition (Figs. 2 and 3 and Extended Data Fig. 3–6).

For analysis of optogenetic inhibition during virtual corridor navigation (Fig. 1 and Extended Data Fig. 2), we removed trials with excess travel of >10% of maze stem (or >330 cm) and mice with <150 total trials from measures of  $y$ -velocity,  $x$ -position and average view angle. Trials with excess travel had similar proportions across laser-off and laser-on trials and pathway-specific inhibition and control groups (indirect pathway: 8.1% of laser-off and 8.2% of laser-on trials; direct pathway: 8.2% of laser-off and 8.1% of laser-on trials; no opsin control: 7.0% of laser-off and 6.9% laser-on trials). Excess travel trials reflected the minority of trials in which mice made multiple traversals of the virtual corridor, thus skewing measures of average  $y$ -velocity,  $x$ -position and view angle during the majority of ‘clean’ corridor traversals. Importantly, we excluded no trials in direct measurements of distance, per-trial view angle standard deviation, and trials with excess travel to detect potential effects of pathway-specific DMS inhibition (or DMS illumination alone) on these measures (Fig. 1 and Extended Data Fig. 2).

Similarly, for all cross-task comparisons (Figs. 2 and 3 and Extended Data Figs. 3 and 6), we removed trials with excess travel for all analyses comparing choice,  $y$ -velocity,  $x$ -position and average view angle. To better capture task-engaged behavior, we also only considered trial blocks in which choice accuracy was greater than >60% for these measures. Excess travel trials were not excluded for cross-task comparisons of effects on measures of distance, per-trial view angle standard deviation and trials with excess travel. Exact trial and mouse number are reported in figure legends.

For cross-task comparisons of motor output measures (Extended Data Fig. 10), we did not exclude trial blocks with choice accuracy of <60%, given that different GLM–HMM states were associated with performance levels. However, for the reasons outlined above, we removed excess travel trials from measures of  $y$ -velocity,  $x$ -position and average view angle, and additionally only considered mice who occupied all three GLM–HMM states after trial selection. For measures of per-trial standard deviation in view angle, distance and excess travel, we applied no trial selection criteria, but only mice who occupied all three GLM–HMM states were included for analysis.



**General performance indicators.** Accuracy was defined as the percentage of trials in which mice chose the maze arm corresponding to the side having the greater number of cues (Fig. 2c). For measures of choice bias, sensory evidence and choice were defined as either ipsilateral or contralateral relative to the unilaterally coupled laser hemisphere. Choice bias was calculated separately for laser-off and laser-on trials as the difference in choice accuracy between trials where sensory evidence indicated a contralateral reward versus when sensory evidence indicated an ipsilateral reward (percentage correct, 'contralateral' – 'ipsilateral', positive values indicate greater contralateral choice bias; Fig. 3d,g,j,o). Delta choice bias was calculated as the difference in contralateral choice bias between laser-off and laser-on trials ('on' – 'off', positive values indicate laser-induced contralateral choice bias; Fig. 3e,g,k,p). In Extended Data Fig. 8c–f, reward at GLM–HMM state transition reflects the total amount of reward (ml) received from the start of the session up to the trial before a state transition (mice typically received 1–1.5 ml per session). Reward rate at GLM–HMM state transition was calculated as the sum of reward received from the start of the session up to the trial before a state transition, divided by the sum of all trial durations from the start of the session up to the trial before a state transition. GLM–HMM transitions were defined as a within-session change in the most likely state based on the posterior probability of each state (see 'GLM–HMM' below for more details).

**Psychometric curve fitting.** Psychometric performance was assessed based on the percentage of contralateral choices as a function of the difference in the number of contralateral and ipsilateral cues ( $n$  contra –  $n$  ipsi; Fig. 6c–f and Extended Data Fig. 4). In Extended Data Fig. 4, transparent lines reflect the mean performance of individual mice in bins (–16:4:16,  $n$  contra –  $n$  ipsi) of sensory evidence during laser-off (black) and laser-on (green) trials, while bold lines reflect the corresponding mean and s.e.m. across mice. In Fig. 6, psychometric curves, indicating the probability of a rightward choice as a function of the sensory evidence, were fit to the following four-parameter sigmoid using maximum likelihood fitting<sup>65</sup>:

$$p(\text{choice} = R|\Delta) = \lambda + \frac{1 - \lambda - \gamma}{1 + \exp(-(\Delta - \sigma)\mu)} \quad (3)$$

where  $\lambda$  and  $\gamma$  are the right and left lapse rates, respectively,  $\sigma$  is the offset,  $\mu$  is the slope, and  $\Delta$  is the difference in the number of contralateral and ipsilateral cues on a given trial. Each point in Fig. 6c–f represents the binned difference in cues in increments of 4 from –16 to 16 (as in Extended Data Fig. 4), from which we calculated the percentage of contralateral choice trials for each bin.

**Motor performance indicators.** Y-velocity (cm s<sup>–1</sup>) was calculated on every sampling iteration (120 Hz, or every ~8 ms) of the ball motion sensor as  $dY/dt$  where  $dY$  was the change in Y-position displacement in VR, and  $dt$  was the elapsed time from the previous sampling of the sensor. The y-velocity for all iterations in which a mouse occupied y-positions 0–300 cm in 25-cm bins were then averaged across iterations in each bin to obtain per-trial y-velocity as a function of y-position. Binned y-velocity as a function of y-position was then averaged across trials for individual mice, and the average and s.e.m. across mice reported throughout (Figs. 1g and 2d and Extended Data Figs. 6a–c and 10b,p,i,w; averaged across y-position 0–200 cm in Extended Data Figs. 2b and 3b).

X-position trajectory (cm) as a function of y-position was calculated per trial by first taking the x-position at y-positions 0–300 cm in 1-cm steps, which was defined as the x-position at the last sampling time  $t$  in which  $y(t) \geq Y$ , and then averaging the y-position bins of 25 cm from 0 to 300 cm. Binned x-position as a function of y-position was then averaged across left/right (or ipsilateral/contralateral) choice trials for individual mice, and the average and s.e.m. across mice was reported throughout (Figs. 1h and Fig. 2e and Extended Data Fig. 10c,q; averaged across y-positions 0–200 cm in Extended Data Figs. 2c, 3c, 6j,l and 10j,x). Average view angle trajectory (degrees) was calculated in the same manner as x-position (Figs. 1i and 2f and Extended Data Fig. 10d,r; average across y-positions 0–200 cm in Extended Data Figs. 2d, 3d, 6j–l and 10k,y). View angle standard deviation was calculated by first sampling the per-trial view angle from 0 to 300 cm of the maze in 5-cm steps. The standard deviation in view angle was then calculated for each trial, and then averaged across trials for individual mice. The average and s.e.m. across mice are reported throughout (Extended Data Figs. 2e, 3f, 6g–i and 10e,s,l,z). This measure sought to capture unusually large deviations in single-trial view angles, which would be indicative of excessive turning or rotations.

Distance traveled was measured per trial as the sum of the total  $x$  and  $y$  displacement calculated at each sampling iteration  $t$ , as  $\sqrt{dX^2 + dY^2}$ . Distance traveled per trial was then averaged across trials for individual mice and the average and s.e.m. across mice was reported throughout (Figs. 1j and 2g and Extended Data Figs. 2f, 6d–f and 10f,t,m,aa). Excess travel was defined as the fraction of trials with total distance traveled per trial (calculated as above) greater than 10% of maze length ( $0 > 330$  cm). The average and s.e.m. across mice was reported throughout (Extended Data Figs. 2g, 3e, 6d–f and 10g,u,n,bb).

Decoding of choice based on the trial-by-trial  $x$ -position (Extended Data Fig. 3g,i) or view angle (Extended Data Fig. 3h,j) of mice was carried out by performing a binomial logistic regression using the MATLAB function 'glmfit'.

In Extended Data Fig. 3g,h, the logistic regression was fit separately for individual mice at successive y-positions in each T-maze stem (0–300 cm in 25-cm bins), where the trial-by-trial average  $x$ -position (or view angle) at each y-position bin (calculated as above) was used to generate weights predicting the probability of a left or right choice given a particular  $x$ -position (or view angle) value. Individual mouse fits were weighted according to the proportion of left and right choice trials. Fivefold cross-validation (resampled for new folds ten times) was used to evaluate prediction accuracy on held-out trials. A choice probability greater than or equal to 0.5 was decoded as a right choice, and prediction accuracy for individual mice was calculated as the fraction of predicted choices matching actual mouse choice, averaged across cross-validation sets. The same approach was used in Extended Data Fig. 3i,j, except that the training data were randomly sampled across all mice from a single task (50% of total trials, resampled 50 times; training data from evidence accumulation, no distractors or permanent cues task). The learned weights were then used to predict choice based on held-out  $x$ -position (or view angle) data from all three tasks, with prediction accuracy calculated as the fraction of predicted choices matching actual choice, parsed by individual mice, and averaged across cross-validation sets. A package of code for behavioral analysis in VR-based T-maze settings is available at <https://github.com/BrainCOGS/behavioralAnalysis/>. In addition, all analyses described here can be replicated at <https://github.com/ssbolkan/BolkanStoneEtAl/>.

**General statistics and reproducibility.** We performed one-way ANOVAs to assess effects of the factor task (three levels: evidence accumulation, no distractors or permanent cue) on choice accuracy (Fig. 2c), distance traveled (Fig. 1g), average y-velocity (0–200 cm; Extended Data Fig. 3b), average  $x$ -position (0–200 cm) on left or right choice trials (Extended Data Fig. 6c), average view angle (0–200 cm) on left or right choice trials (Extended Data Fig. 3d), percentage of trials with excess travel (Extended Data Fig. 3e), per-trial standard deviation in view angle (Extended Data Fig. 3f), delta (laser on–off) choice bias (Extended Data Fig. 5b–d), delta (laser on–off) distance traveled (Extended Data Fig. 6d–f, left), delta (laser on–off) percentage of trials with excess travel (Extended Data Fig. 6d–f, right), delta (laser on–off) per-trial standard deviation in view angle (Extended Data Fig. 6g–i), delta (laser on–off) average  $x$ -position (0–200 cm; Extended Data Fig. 6j–l, left), and delta (laser on–off) average view angle (0–200 cm; Extended Data Fig. 6j–l, right). Post hoc comparisons between tasks were made when a main effect of the factor task had a  $P$  value less than  $\alpha < 0.05/2$  to account for multiple-group comparisons (Extended Data Fig. 5b–d). One exception to this is Extended Data Fig. 6j–l, where all post hoc comparisons were made for laser effects on delta (on–off)  $x$ -position and view angle (and displayed with corresponding exact  $P$  values) to provide greater clarity around trend-level effects. We did not assume normality of the data in all post hoc comparisons, which used the nonparametric, unpaired, two-tailed Wilcoxon rank-sum test. A  $P$  value below 0.025 was considered significant to correct for multiple comparisons ( $P = 0.05/2$  comparisons per group). Exact  $P$  values, degrees of freedom, and  $z$ -statistics are reported in the text and/or legends.

We performed one-way ANOVAs to assess effects of the factor group (three levels: indirect pathway inhibition, direct pathway inhibition or no opsin illumination) on delta (laser on–off) y-velocity (Extended Data Fig. 2b), delta (laser on–off)  $x$ -position (Extended Data Fig. 2c), delta (laser on–off) view angle (Extended Data Fig. 2d), delta (laser on–off) per-trial standard deviation in view angle (Extended Data Fig. 2e), delta (laser on–off) distance (Extended Data Fig. 2f), delta (laser on–off), and delta (laser on–off) percentage of trials with excess travel (Extended Data Fig. 2g) and delta (laser on–off) preference (that is, time) and speed during the real-time CPP test (Supplementary Fig. 3).

We performed a repeated-measure one-way ANOVA to assess effects of the within-subject factor state (three levels: GLM–HMM state 1, 2 or 3) on within-session accumulated reward or reward rate before GLM–HMM transition (Extended Data Fig. 8c–f), per-trial standard deviation in view angle (Extended Data Fig. 10e,s), distance traveled (Extended Data Fig. 10f,t), percentage of trials with excess travel (Extended Data Fig. 10g,u), delta (laser on–off) average  $x$ -position 0–200 cm (Extended Data Fig. 10j,x), delta (laser on–off) average view angle 0–200 cm (Extended Data Fig. 10k,y), delta (laser on–off) per-trial standard deviation in view angle (Extended Data Fig. 10l,z), delta (laser on–off) distance traveled (Extended Data Fig. 10m,aa) and delta (laser on–off) percentage of trials with excess travel (Extended Data Fig. 10n,bb). Post hoc comparisons between groups were made when a main effect of the factor task had a  $P$  value  $< 0.05$ . One exception is in Extended Data Fig. 6j,k,x,y, where all post hoc comparisons were made for laser effects on delta (on–off)  $x$ -position and view angle (and displayed with corresponding exact  $P$  values) to provide greater clarity around trend-level effects. We did not assume normality of the data in all post hoc comparisons, which used the nonparametric, unpaired, two-tailed Wilcoxon rank-sum test. A  $P$  value below 0.025 was considered significant to correct for multiple comparisons ( $P = 0.05/2$  comparisons per group). Exact  $P$  values, degrees of freedom and  $z$ -statistics are reported in the text and/or legends.

In Fig. 3e,h,k,p, we used the nonparametric, unpaired, two-tailed Wilcoxon rank-sum test to assess effects of indirect or direct pathway inhibition versus no opsin illumination of brain tissue on delta (laser on–off) choice bias. A  $P$  value below 0.025 was considered significant to correct for multiple comparisons ( $P = 0.05/2$  comparisons per group). Exact  $P$  values, degrees of freedom and  $z$ -statistics are reported in the text and/or legends.



Related to Fig. 1a–c, Extended Data Fig. 1 and Supplementary Fig. 2, we used a paired, two-tailed Wilcoxon signed-rank test on cross-trial average firing rates (baseline-pre versus laser-on or laser-on versus laser-post) to determine the significance of laser modulation of single-neuron activity. A Bonferroni-corrected  $P$  value was used to determine significance ( $P < 0.00083$  for 60 indirect pathway neurons or  $P < 0.001$  for 50 direct pathway neurons).

All experiments, but not analysis, were conducted blind to experimental conditions. All experiments were conducted over multiple cohorts (see Supplementary Table 1 for details of individual mouse testing across experiments). Individual cohorts consisted of a random selection of test groups (for example, DMS indirect and direct pathway inhibition, and DMS no-opsin mice). We did not account for cohort effects in our statistical analyses, but no obvious cohort-dependent effects were qualitatively observable. No statistical method was used to predetermine group sample sizes; rather, animal and trial numbers were targeted to match or exceed those from similar studies.

**Bernoulli GLM.** *Coding of covariates and choice output.* We coded the external covariates (referred to as inputs in Fig. 4b) and output (the mouse's choice) on each trial as follows:

$\Delta$  cues: an integer value from  $-16$  to  $16$ , divided by the standard deviation of the  $\Delta$  cues across all sessions in all mice, representing the standardized difference between the number of cues on the right and left sides of the maze.

Laser: a value of  $1$ ,  $-1$  or  $0$  depending on whether optogenetic inhibition was on the right hemisphere, left hemisphere or off, respectively.

Previous choice: a value of  $1$  or  $-1$  if the choice on a previous trial was to the right or left, respectively. We set the value to  $0$  at the start of each session when there was an absence of previous choices (for example, for the third trial of a session, previous choices 3–6 would be coded as  $0$ ).

Previous rewarded choice: a value of  $1$ ,  $-1$  or  $0$  depending on whether the previous choice was correct and to the right, correct and to the left or incorrect, respectively.

Choice output: a value of  $1$  or  $0$  depending on whether the mouse turned right or left.

**Fitting.** We used a Bernoulli GLM, also known as logistic regression, to model the binary (right/left) choices of mice as a function of task covariates. This also corresponds to a one-state GLM–HMM (Fig. 4 and Extended Data Fig. 7a). The model was parameterized by a weight vector (carrying weights for sensory evidence, choice and reward history, and DMS inhibition). On each trial  $t$ , the weights map the external covariates to the probability of each choice  $y_t$ . The model can be written according to equation (4) as:

$$p(y_t = 1 | x_t) = \frac{1}{1 + \exp(-w^T x_t)}, \quad (4)$$

where  $y=1$  indicates a rightward choice, and  $y=0$  indicates a leftward choice. We fit the model by penalized maximum likelihood, which involved minimizing the negative log-likelihood function plus a squared penalty term on the model weights. The log-likelihood function is given by the conditional probability of the choice data  $Y = y_1, \dots, y_T$  given all the external covariates  $X = x_1, \dots, x_T$ , considered as a function of the model parameters given by equation (5):

$$\begin{aligned} \log p(Y | X, w) &= \sum_{t=1}^T \log p(y_t | x_t) \\ &= \sum_{t=1}^T y_t \log p(y_t = 1 | x_t) + \sum_{t=1}^T (1 - y_t) \log (1 - p(y_t = 1 | x_t)) \end{aligned} \quad (5)$$

We then minimized the loss function, given by  $-\log p(Y | X, w) + \frac{1}{2} w^T w$ , using Python's 'scipy.optimize.minimize'. This can be interpreted as a log-posterior over the weights, with  $\frac{1}{2} w^T w$  representing the negative log of a Gaussian prior distribution with mean zero and variance 1, which regularizes by penalizing large weight values. A variance of 1 was chosen because the resulting penalty was sufficient to resolve cases where the weights would grow unusually large and lead to decreases in the log-likelihood during fitting. We computed the posterior standard deviation of the fitted GLM weights (shown as error bars in Fig. 4c,d) by taking the diagonal elements of the inverse negative Hessian (matrix of second derivatives) of the log-posterior at its maximum<sup>36,57</sup>.

**GLM–HMM.** *Model architecture.* To incorporate discrete internal states, we used a hidden Markov model (HMM) with a Bernoulli GLM governing the decision-making behavior in each state. The model is defined by a transition matrix and a vector of GLM weights for each state. The transition matrix contains a fixed set of probabilities that govern the probability of changing from a state  $z \in \{1, \dots, K\}$  on trial  $t$  to any other state on the next trial. We refer to these as transition probabilities, which can be abbreviated according to equation (6):

$$P_{ij} = p(z_{t+1} = j | z_t = i) \quad (6)$$

Each GLM has a unique set of weights  $w_k$  that maps the external covariates  $x_t$  (coded as described in 'Bernoulli GLM') to the probability of the choice  $y_t$  for each

of the  $k$  states. These probabilities can be expressed as a modified version of Eq. (4), as given by equation (7):

$$p(y_t = 1 | x_t, z_t = j) = \frac{1}{1 + \exp(-w_j^T x_t)} \quad (7)$$

where  $w_j$  is the vector of GLM weights for state  $j$ . Note that in this expression, the choice probability depends on both the external covariates (inputs) and the state via the state-dependent GLM weights<sup>29–31</sup>. We refer to these as 'observation probabilities', which can be abbreviated according to equation (8):

$$\phi_{ij} = p(y_t = j | x_t = i) \quad (8)$$

**Fitting.** We fit the GLM–HMM to the data using the expectation–maximization (EM) algorithm<sup>30</sup>. The EM algorithm computes the maximum likelihood estimate of the model parameters using an iterative procedure that involves an E-step (expectation), in which the posterior distribution of the latent variables is calculated, followed by an M-step (maximization), in which the values of the model parameters are updated given the posterior distribution of the latents. These steps are repeated until the log-likelihood of the model converges on a local optimum<sup>57</sup>.

The log-likelihood (also referred to as the log marginal likelihood) is obtained from the joint probability distribution over the latent states  $Z = z_1, \dots, z_T$  and the observations  $Y = y_1, \dots, y_T$  on each trial given the model parameters  $\theta$ . Marginalizing over the latents, the log-likelihood is computed as the log of the sum over states of the marginal probabilities and is written as shown in equation (9):

$$\log p(Y | X, \theta) = \log \left( \sum_z p(Y, Z | X, \theta) \right) \quad (9)$$

The set of parameters  $\theta$  governing the model consists of a transition matrix and the state-dependent GLM weights, which we described above. We initialized the transition matrix by sampling each row from a Dirichlet distribution, with a larger concentration parameter over the entries along the diagonal ( $\alpha_{ii} = 5$ ,  $\alpha_{ij} = 1$ ), reflecting a slight bias toward self-transitions. For the GLM weights, we reasoned that the true values for each state would likely be in approximately the same range as the true values for the one-state (GLM) case. Therefore, we initialized the per-state GLM weights  $w_k$  with  $k \in \{1, \dots, K\}$  by first fitting a basic GLM ('Bernoulli GLM') to find  $w_0$ . Then, as we didn't want the initial weights to be the same in each state, we initialized  $w_k = w_0 + e_k$  where  $e \sim \mathcal{N}(0, 0.2)$ .

The goal of the E-step of the EM algorithm is to compute  $p(Z | X, Y, \theta)$ , the posterior probability of the latent states given the observations and the model parameters. This can be obtained using a two-stage message passing algorithm known as the forward–backward algorithm<sup>30</sup>. The forward pass, sometimes called 'filtering', finds the normalized conditional probability  $\hat{\alpha}(z_t)$  for each state  $z$  at trial  $t$  by iteratively computing the following according to equation (10):

$$\begin{aligned} \hat{\alpha}(z_t) &= p(z_t | y_{1:t}, x_{1:t-1}) \\ &= \frac{p(y_t | z_t, x_t)}{p(y_t | y_{1:t-1})} \sum_{k=1}^K p(z_{t-1} = k, y_{1:t-1} | x_{1:t-1}) p(z_t | z_{t-1} = k) \\ &= c_t^{-1} \phi_{z_{t-1}, y_t} \sum_{k=1}^K \hat{\alpha}(z_{t-1} = k) P_{k, z_t} \end{aligned} \quad (10)$$

where  $c_t = p(y_t | y_{1:t-1})$  is a scale factor ensuring the probabilities over states sum to 1, which is computed by summing the unnormalized probabilities. We set the prior distribution over states before the first trial, denoted  $\hat{\alpha}(z_0)$ , to be the uniform distribution. Note that this is a normalized version of the forward–backward algorithm that avoids underflow errors<sup>57</sup>.

The backward step, also referred to as 'smoothing', takes the information from the forward pass and works in the reverse direction, carrying the information about future states backwards in time to further refine the latent state probabilities. Here we find the normalized conditional probability  $\hat{\beta}(z_t)$  for each state  $z$  at trial  $t$  by iteratively computing the following according to equation (11):

$$\begin{aligned} \hat{\beta}(z_t) &= \frac{p(y_{t+1:T} | x_{t+1:T}, z_t)}{p(y_{t+1:T} | y_{1:t})} \\ &= \frac{1}{p(y_{t+1} | y_{1:t})} \sum_{k=1}^K \frac{p(y_{t+2:T} | x_{t+2:T}, z_{t+1} = k)}{p(y_{t+2:T} | y_{1:t+1})} p(y_{t+1} | z_{t+1} = k, x_{t+1}) p(z_{t+1} = k | z_t) \\ &= c_{t+1}^{-1} \sum_{k=1}^K \hat{\beta}(z_{t+1} = k) \phi_{k, y_{t+1}} P_{z_t, k} \end{aligned} \quad (11)$$

where  $\hat{\beta}(z_T) = 1$ .

From these two conditional probabilities, we can calculate the marginal posterior probabilities of the latent states given by equation (12):

$$\gamma(z_t) = p(z_t | X, Y, \theta) = \hat{\alpha}(z_t) \hat{\beta}(z_t) \quad (12)$$

which was the goal of the E-step. We can also compute the joint posterior distribution of two successive latents, given by equation (13):

$$\xi(z_t, z_{t+1}) = \hat{\alpha}(z_t) \hat{\beta}(z_{t+1}) P_{z_t, z_{t+1}} \phi_{z_{t+1}, y_{t+1}} \quad (13)$$

which will be important for computing updates in the M-step. Because the format of the data included sessions from several different mice over many days, we computed the forward-backward pass separately for each session. This ensured that the learned transition probabilities would not take into account the effect of the last trial of one session on the first trial of the next session.

The M-step of the EM algorithm takes the newly computed posterior probabilities over the latents and uses them to update the values of the model parameters (Eqs. (6)–(8)) by maximizing the expression for  $P$  and  $w$ . Because the transition probabilities are fixed, we can compute their updates using the closed-form solution given by equation (14):

$$P_{jk} = \frac{\sum_{t=2}^T \xi(z_{tj}, z_{t+1,k})}{\sum_{l=1}^K \sum_{t=2}^T \xi(z_{tj}, z_{t+1,l})} \quad (14)$$

This closed-form update can be derived by applying the appropriate Lagrange multipliers to the complete-data log-likelihood function<sup>57</sup>.

Maximization for  $w$  involves minimizing the negative of the log-likelihood function, weighted by the marginal posterior probabilities of the latent states, plus a squared penalty term on the model weights. This penalty can be interpreted as the negative log of a Gaussian prior with mean zero and variance 1, which regularizes by penalizing large weight values. The resulting loss function is given by equation (15):

$$L(w) = \sum_{t=1}^T \sum_{k=1}^K \gamma(z_t = k) \log(p(y_t | x_t, z_t = k)) - \frac{1}{2} w_k^T w_k \quad (15)$$

which we optimized using numerical optimization and the L-BFGS-B algorithm as previously described (see ‘Bernoulli GLM’).

Both E- and M-steps of the EM algorithm are guaranteed to increase the log-posterior. We alternated E- and M-steps until the difference between the log-posteriors over ten iterations was smaller than a given tolerance ( $\text{tol} = 1 \times 10^{-3}$ ). Because the EM algorithm only guarantees that the log-likelihood will converge upon a local optimum<sup>57</sup>, we fit the model 20 times using different initializations of the weights and transition matrix and verified that the top four or more fits all converged on the same solution (meaning that the weights for each fit were the same within a tolerance of  $\pm 0.05$ ) to confirm that the algorithm had indeed found the global optimum. After determining the best fit, we computed the posterior standard deviation of the fitted GLM weights (shown as error bars in Fig. 6a,b and Extended Data Fig. 7d–f) by computing the inverse Hessian of the optimized log-posterior using Python’s autograd package.

**Model selection.** In Extended Data Fig. 7a, we performed cross-validation on the data from both the indirect and direct pathway inhibition groups, which revealed that three to five latent states were sufficient to reach a plateau in likelihood. To obtain a test set, we selected ~20% of sessions from the data to hold out from model fitting. Test sessions were chosen by randomly selecting an approximately equal number of sessions from each of the 13 mice in either group. Constraining the held-out data in this way ensured that the cross-validation results were not affected by possible individual differences across mice. We then calculated the log-likelihood of the test data after fitting the model under parameterizations of 1–5 states to the remaining ~80% of sessions. We express the log-likelihood in bps, defined according to equation (16):

$$L_{\text{bps}} = l \cdot \frac{\hat{L} - \hat{L}_0}{T \cdot \log(2)} \quad (16)$$

where  $l$  is the average session length,  $T$  is the number of trials in the test set and  $\hat{L}_0$  is the log-likelihood of the test-set data under the bias-only Bernoulli GLM. To obtain the bias term  $b$  we computed the following given by equation (17):

$$b = \frac{T(\text{right})}{T(\text{right}) + T(\text{left})} \quad (17)$$

where  $T(\text{side})$  is the number of trials in the test set in which the mice turned in that direction. For all cross-validation results presented in the paper, we report the averaged  $L_{\text{bps}}$  from five different test sets. We followed the same procedure as above in Extended Data Fig. 7c, selecting the optimal number of previous choices using a three-state GLM–HMM under parameterizations of 1–8 previous choices while holding the number of all other external inputs ( $\Delta$  cues, laser, bias and previous rewarded choice) constant. In Extended Data Fig. 7b, we simulated data from the inferred parameters (see ‘Simulating data’) for a two-state GLM–HMM fit to the real data for the indirect pathway inhibition group. We then performed cross-validation as described above on both the full simulated dataset (‘all data’) and for a small subset (‘5% of data’). We chose to simulate from the two-state model

to differentiate the simulation from the real data and to demonstrate the results for an arbitrary choice in the true number of states.

**Testing.** In Fig. 5c, we compared the performance of the GLM–HMM to the GLM by calculating the log-likelihood of the test sets of individual mice. To do so, we held out data and fitted the model across all animals using the same procedure described above. However, we then split the test set by mouse (thus creating 13 different test sets) and calculated the log-likelihood for each individual animal, thus expressing the log-likelihood in units of mouse bits per session (mbps) given by equation (18):

$$L_{\text{mbps}} = l \cdot \frac{\hat{L}_m - \hat{L}_{0m}}{T_m \cdot \log(2)} \quad (18)$$

Here,  $\hat{L}_m$  is the optimized log-likelihood of the model in question (either the GLM or three-state GLM–HMM) for a single mouse. Similarly,  $\hat{L}_{0m}$  is the optimized log-likelihood under the bias-only Bernoulli GLM and  $T_m$  is the total number of trials for that mouse. We then repeated the procedure for five test sets and took the average of the results for each mouse.

In Fig. 5d, we evaluated the prediction accuracy of the GLM for each animal by taking the same training and test sets that we used to find the log-likelihoods and using Eq. (2) to calculate the probability of turning right on each trial. We then compared this probability to the mouse’s actual choice on that trial, labeling the trial as correct if the model predicted a 50% or greater probability of turning in the direction of the mouse’s true choice. We then calculated the prediction accuracy for each mouse as the number of correct trials divided by the total number of trials for that mouse. To evaluate the prediction accuracy of the GLM–HMM for each animal, we computed  $p(y_t | x_{1:t-1}, y_{1:t-1})$ , or the predictive distribution for trial  $t$  of the test set using the observations from trials 1 to  $t-1$ . This arises from averaging over the state probabilities given previous choice data to get a prediction for a particular trial. That is, we ran the forward pass (see ‘Fitting’) to obtain the state probabilities  $p(z_t | x_{1:t-1}, y_{1:t-1})$ , computed the initial choice probabilities  $p(y_t | x_t, z_t)$  using Eqs. (7) and (8), and then calculated the predictive distribution according to equation (19):

$$p(y_t | x_{1:t-1}, y_{1:t-1}) = \sum_{k=1}^K p(y_t | x_t, z_t = k) p(z_t = k | x_{1:t-1}, y_{1:t-1}) \quad (19)$$

We then ran this forward over all the trials in the test set for each mouse. Finally, we computed the prediction accuracy using the same method described for the GLM prediction accuracy.

**State assignments.** To determine the most likely state on each trial (Figs. 6c–i and 7g,j) and Extended Data Figs. 7g,h, 8 and 10 and Supplementary Fig. 4), we assigned each trial to the state with maximum marginal probability given the inputs and choice data, as computed by the forward-backward algorithm given by equation (20):

$$s_t = \arg \max_z (p(z_t | \mathbf{X}, \mathbf{Y}, \theta)) \quad (20)$$

**Simulating data.** For the analyses in Fig. 6e,f and Extended Data Figs. 7 and 9, we evaluated the ability of the three-state GLM–HMM to predict choices and state transitions that matched the animals’ actual behavior in each state. Regarding the covariates for the simulation, we kept the evidence ( $\Delta$  cues) and optogenetic inhibition from the real data but populated the trial history covariates using simulated previous choices. To simulate choices on each trial, we first computed the observation probabilities (Eqs. (7) and (8)) using  $x'_t$  (the external covariates) and  $w_k$  (the learned weights from the model fitted to real data). The state  $k$  on each trial was randomly chosen from a distribution given by the learned transition probabilities  $P_{z_{t-1}, z_t}$  from the model fitted to real data. We then randomly generated choices  $y'_t$  from the distribution of observation probabilities. Repeating this process for each trial to obtain  $x'_{1:T}$  and  $y'_{1:T}$ , we fit the model to the simulated data using the same procedure described previously (see ‘Fitting’) to obtain the posterior probability over states. For Fig. 6e,f and Extended Data Fig. 7, we computed the psychometric curves for each state using these posterior probabilities and the simulated choices (see ‘Psychometric curve fitting’).

**Model comparisons.** For the two alternative model comparisons with restricted transition probabilities (Fig. 7k–m), we fit the three-state GLM–HMM using the same general procedure as described above. However, in the case where we disallowed transitions during a session (Fig. 7k), we fixed the transition matrix to the identity matrix and only fit the state-dependent GLM weights. In the case where we disallowed transitions in and out of state 2 (Fig. 7l), we derived a constrained M-step that forced the transition probabilities for state 2 to 0. In detail, the constrained M-step involved zeroing out the transition probabilities associated with state 2 and then renormalizing so the rows of the transition matrix summed to 1. Note that the three sessions that appear to still allow transitions in and/or out of state 2 for mice inhibited in the direct pathway of the DMS (Fig. 7l, right) were due

to rare cases where the model had high uncertainty about the state, and the most probable state flipped between state 2 and another state at some point during the session. In Fig. 7m, solid curves denote the average log-likelihood for five different test sets. Held-out data for test sets were selected as a random 20% of sessions, using an approximately equal number of sessions for each mouse.

**Fluorescence in situ hybridization and stereological quantification.** In situ hybridization (Supplementary Fig. 1) was performed using the RNAscope Multiplex Fluorescent Assay (ACD, no. 323110) with the following probes: Mm-Drd1a (406491), Mm-Drd2-C2 (406501-C2, 1:50 dilution in C1 solution) and Cre-01-C3 (474001-C3, 1:50 dilution in C1 solution). Likely due to lower expression of Cre mRNA in D1R-Cre and A2a-Cre mice, we did not detect unambiguous Cre fluorescence signal in these lines. We therefore relied on Cre-dependent viral expression of AAV5-DIO-EYFP to report Cre<sup>+</sup> neurons alongside Drd1a and Drd2 probes in sections from two D1R-Cre and two A2R-Cre mice, but used all three probes in sections from two D2R-Cre mice. In D1R-Cre and A2R-Cre mice, the Drd1a and Drd2 probes were fluorescently linked to TSA Plus Cy-3 and TSA Plus Cy-5, respectively (Perkin Elmer). In D2R-Cre mice, Drd1a, Drd2 and Cre probes were linked to TSA Plus Cy-3, TSA Plus Fluorescein or TSA Plus Cy-5, respectively. All fluorophores were reconstituted in dimethylsulfoxide according to Perkin Elmer instructions and diluted 1:1,200 in TSA buffer included in the RNAscope kit. After in situ hybridization, slides were coverslipped using Fluoromount-G containing DAPI (SouthernBiotech).

We then obtained  $\times 20$  confocal z-stacks from the DMS, NAcCore and NAcShell in all lines and manually quantified specificity, penetrance and D1R<sup>+</sup>/D2R<sup>+</sup> overlap using LASX software (Leica). Specificity was determined as the percentage of the following: GFP<sup>+</sup> neurons coexpressing Drd1 in D1R-Cre mice (DMS,  $n=5$  sections, 193 cells; NAcCore,  $n=5$  sections, 298 cells; NAcShell,  $n=5$  sections, 363 cells), GFP<sup>+</sup> neurons coexpressing Drd2 in A2a-Cre mice (DMS,  $n=4$  sections, 144 cells; NAcCore,  $n=4$  sections, 326 cells; NAcShell,  $n=4$  sections, 312 cells) or Cre<sup>+</sup> neurons coexpressing Drd2 in D2R-Cre mice (DMS,  $n=5$  sections, 1,302 cells; NAcCore,  $n=5$  sections, 1,104 cells; NAcShell,  $n=5$  sections, 1,187 cells). Penetrance was determined as the percentage of Drd2<sup>+</sup> neurons coexpressing Cre in D2R-Cre mice (DMS,  $n=5$  sections, 1,269 cells; NAcCore,  $n=5$  sections, 1,055 cells; NAcShell,  $n=5$  sections, 1,144 cells). We did not assess penetrance in D1R-Cre or A2a-Cre lines because our Cre-dependent viral reporter did not fully penetrate all Cre<sup>+</sup> neurons. Quantification of D1R<sup>+</sup>/D2R<sup>+</sup> overlap in striatal regions was carried out on two D2R-Cre mice and/or two D1R-tdTomato mice and measured as both the percentage of D1R<sup>+</sup> neurons that were D2R<sup>+</sup> (DMS,  $n=10$  sections, 2,423 cells; NAcCore,  $n=10$  sections, 2,196 cells; NAcShell,  $n=10$  sections, 2,220 cells) and the percentage of D2R<sup>+</sup> neurons that were D1R<sup>+</sup> (DMS,  $n=5$  sections, 868 cells; NAcCore,  $n=5$  sections, 834 cells; NAcShell,  $n=5$  sections, 874 cells).

**Histology.** Mice were anesthetized with a 0.05-ml injection of Euthasol intraperitoneally and transcardially perfused with 1 $\times$  PBS, followed by fixation with 4% paraformaldehyde (PFA). Whole brains with intact fiber-optic implants were post-fixed in 4% PFA for 1–3 d, followed by brain dissection and another 24 h of post-fixation in PFA. For optogenetic experiments, brains were then transferred to PBS for coronal sectioning (50  $\mu$ m) on a vibratome. Viral expression and fiber-optic placements were assessed under slide-scanning (NanoZoomer, Hamamatsu) or single slide (Leica) epifluorescent microscopes (Supplementary Fig. 6). For FISH experiments, post-fixation dissected brains were transferred through a sucrose gradient: 10% sucrose in PBS for 6–8 h, 20% sucrose in PBS overnight and 30% sucrose in PBS overnight. Coronal sections (18  $\mu$ m) containing the DMS and NAc were made using a cryostat, mounted without coverslips on Superfrost plus slides (Fisher), and stored at  $-80^{\circ}$  before the FISH protocol. After the FISH protocol, tile-scanning and cellular resolution images of coverslipped slides were acquired using a confocal microscope (Leica TCS SP8).

**Reporting Summary.** Further information on research design is available in the Nature Research Reporting Summary linked to this article.

## Data availability

The data that support the findings of this study are publicly available on figshare at <https://doi.org/10.6084/m9.figshare.17299142.v1>.

## Code availability

Code for general use applications of GLM–HMM analyses developed in this study, including all applications to the present dataset, are available on GitHub at <https://github.com/irisstone/glmhmm/>. Code to analyze data and regenerate all other plots in this study is publicly available at <https://github.com/ssbolkan/BolkanStoneEtAl>.

## References

- Aronov, D. & Tank, D. W. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* **84**, 442–456 (2014).
- Hanks, T. D. et al. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**, 220–223 (2015).
- Wichmann, F. A. & Hill, N. J. The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept. Psychophys.* **63**, 1293–1313 (2001).
- Pillow, J. W., Ahmadian, Y. & Paninski, L. Model-based decoding, information estimation, and change-point detection techniques for multineuron spike trains. *Neural Comput.* **23**, 1–45 (2011).
- Bishop, C. M. Chapter 13: Sequential Data. in *Pattern Recognition and Machine Learning (Information Science and Statistics)* (eds. Jordan, M., Kleinberg, J. & Schölkopf, B.) 605–652 (Springer-Verlag, 2006).

## Acknowledgements

We thank the entire BRAINCoGs team as well as the laboratories of I.B.W. and J.W.P. for feedback on this work. We thank S. Stein and S. Baptista for technical support in animal training, and C. Kopecs for technical assistance. This work was supported by grants from the National Institutes of Health R01 DA047869 (to I.B.W.), F32MH118792 (to S.S.B.), F32NS101871 (to L.P.), K99MH120047 (to L.P.), U19 NS104648-01 (to J.W.P. and I.B.W.) and ARO W911NF1710554 (to I.B.W.), the Brain Research Foundation (to I.B.W.), Simons Collaboration on the Global Brain (to J.W.P. and I.B.W.), 1R01MH106689 (to I.B.W.) and the New York Stem Cell Foundation (to I.B.W.). I.B.W. is an NYSCF–Robertson Investigator.

## Author contributions

S.S.B. performed the experiments. I.R.S. developed the GLM–HMM. S.S.B. and I.R.S. analyzed the data with guidance from J.W.P. and I.B.W. L.P., Z.C.A. and B.E. provided technical and analysis support. J.M.I., A.L.H. and P.S. aided behavioral training. A.B. performed histology. J.C. provided support for electrophysiology. C.A.Z. and J.R.C. provided support for the virtual corridor. S.S.B. and I.B.W. conceived the experimental work. S.S.B., I.R.S., J.W.P. and I.B.W. interpreted the results and wrote the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

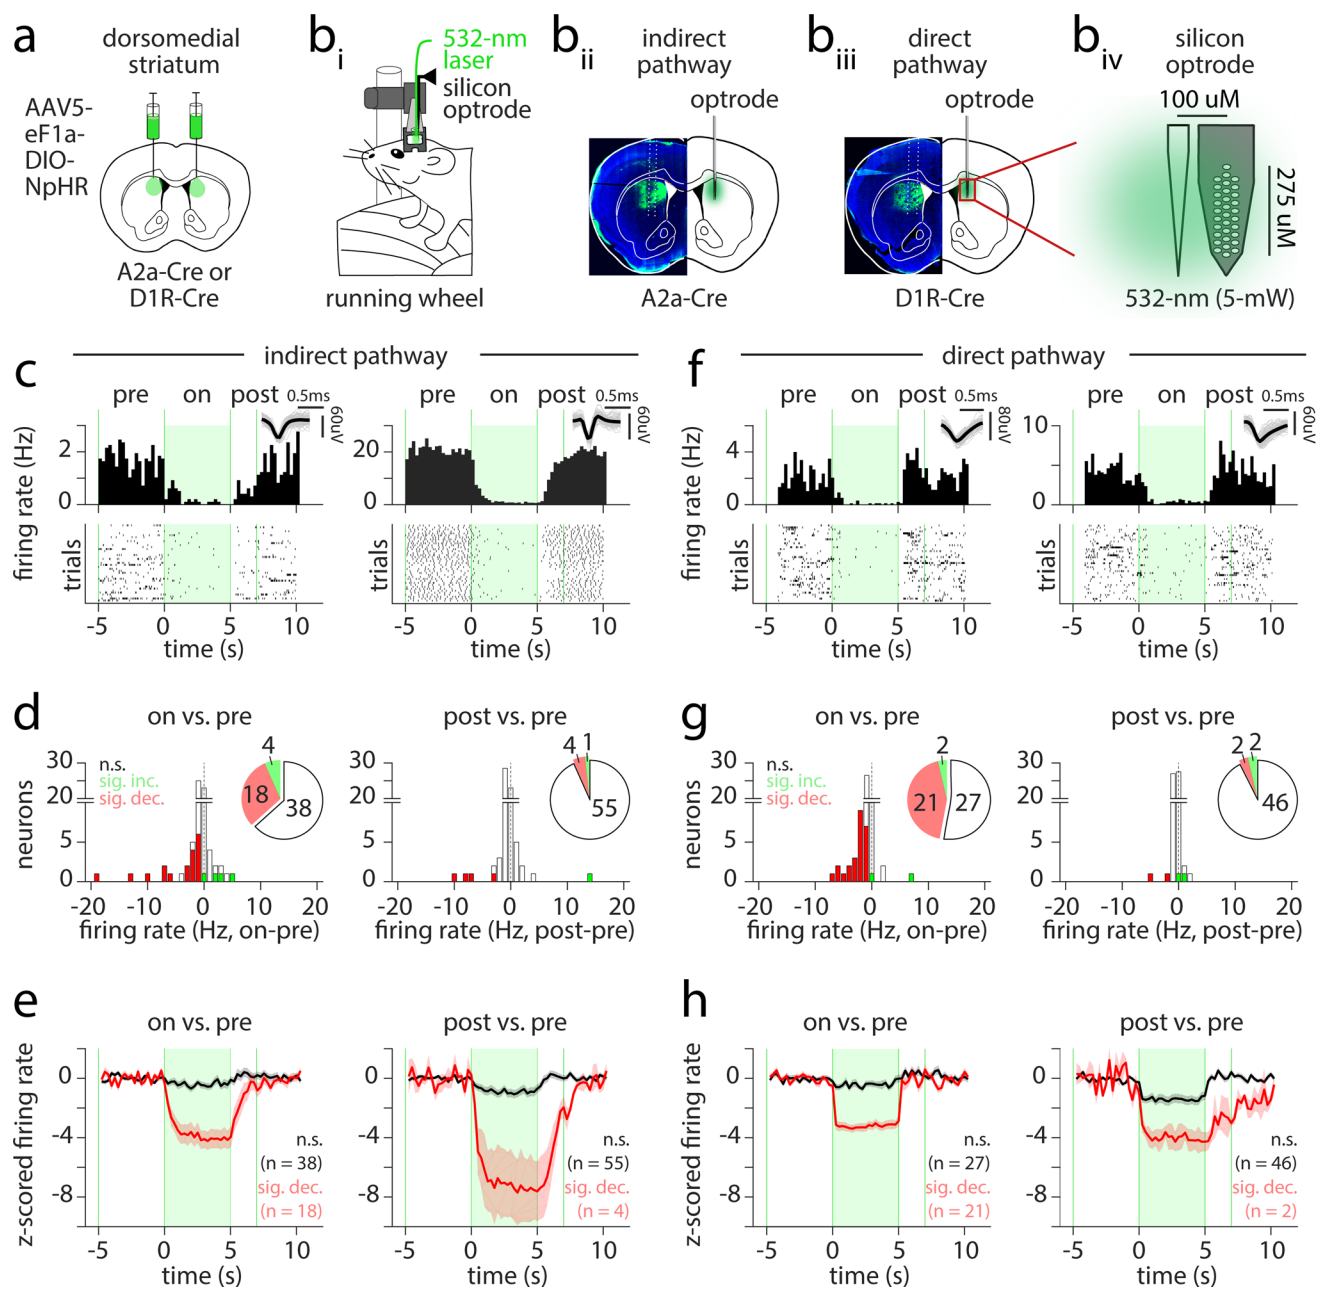
**Extended data** is available for this paper at <https://doi.org/10.1038/s41593-022-01021-9>.

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41593-022-01021-9>.

**Correspondence and requests for materials** should be addressed to Jonathan W. Pillow or Ilana B. Witten.

**Peer review information** *Nature Neuroscience* thanks Naoshige Uchida and the other, anonymous reviewer(s) for their contribution to the peer review of this work.

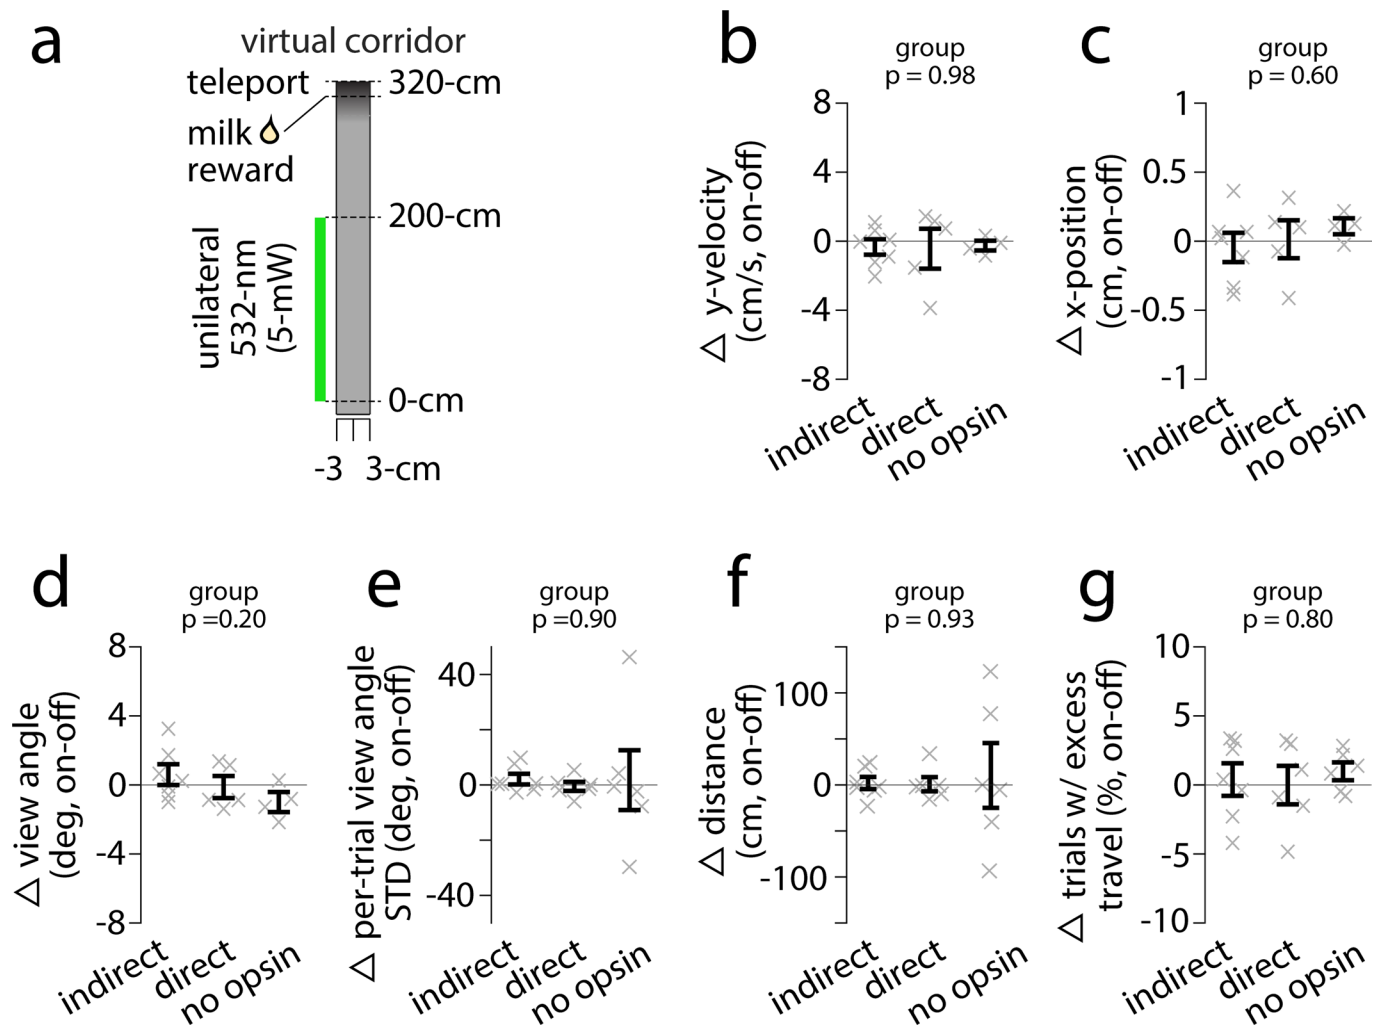
**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).



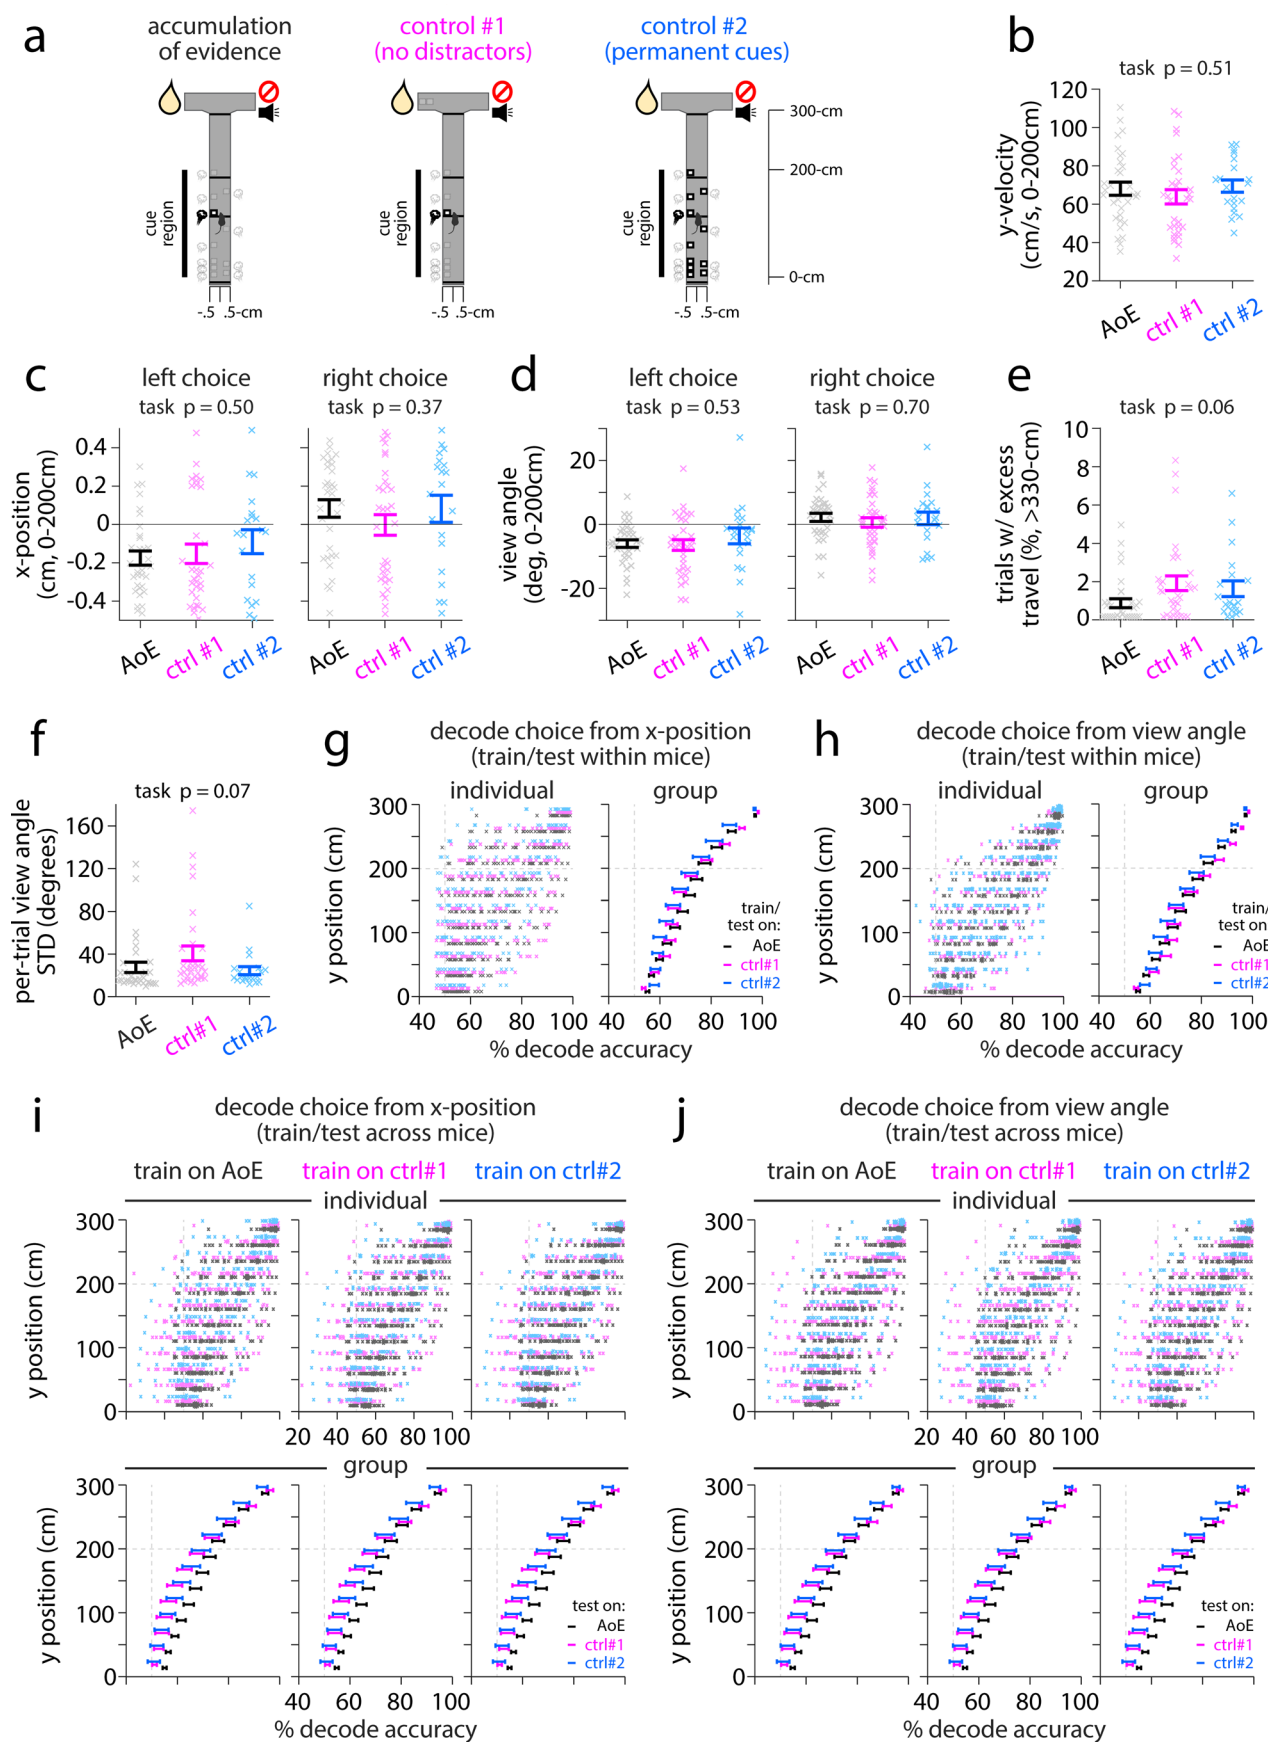
Extended Data Fig. 1 | See next page for caption.



**Extended Data Fig. 1 | Optogenetic inhibition of DMS pathways is effective, generating little post-inhibitory rebound, nor excitation during the inhibition period.** (a) Schematic of viral delivery of AAV5-eF1a-DIO-NpHR to the dorsomedial striatum (DMS) of A2a-Cre or D1R-Cre mice. (b,i) Schematic of electrophysiological recording and laser delivery (532-nm, 5-mW) to the DMS in awake, head-fixed mice ambulating on a running wheel. (b,ii) Example recording electrode tracks and cre-dependent NpHR expression in an A2a-Cre mouse targeting the indirect pathway of the DMS. (b,iii) As in b,ii but in a D1R-Cre mouse targeting the DMS direct pathway. (b,iv) Schematic of silicon optrode recording tip, including tapered optical fiber coupled to a 32-channel silicon probe. (c) Two example peristimulus time histograms (PSTH) (top) and raster plots of trial-by-trial spike times (bottom) from single neurons recorded from the DMS of an A2a-Cre mouse. Inset at top displays average spike waveform (black) and 100 randomly sampled spike waveforms (grey) for each neuron. A trial consisted of 5-s without laser (*pre*, -5 to 0-s), 5-s of 532-nm light (5-mW) delivery (*on*, 0 to 5-s), followed by a 10-s ITI (40 trials per recording site). The first 2-s following laser offset (*post*, 5-7-s) was used to assess post-inhibitory effects. (d) Left: Histogram of change in average firing rate (*on-pre*, Hz) for all neurons ( $n=60$ ) recorded from the DMS of A2a-Cre mice ( $n=3$ ). Colors indicate non-significant (black,  $n=38$  neurons), significantly decreased (red,  $n=18$  neurons) or increased (green,  $n=4$  neurons) changes in firing rate determined via paired, two-tailed signrank comparison of average across-trial baseline (*pre*) or laser (*on*) firing rates. A Bonferroni-corrected significance threshold was used to account for multiple neuron comparisons ( $p < 0.00083$ , or  $p = 0.05/60$  neuron comparisons). Right: same as left but for change in firing rate (*post-pre*, Hz): non-significant ( $n=55$  neurons), significantly decreased ( $n=4$ ) or increased ( $n=1$ ). Insets display pie-chart summaries of the proportion of non-significant (black unfilled), significantly decreased (red) or increased (green) neurons. (e) Left: Mean  $\pm 1$  s.e.m. z-scored firing rate across all non-significantly modulated *on* vs *pre* (black,  $n=38$ ) or significantly decreased *on* vs *pre* (red,  $n=18$ ) neurons recorded from A2a-Cre mice. Right: same as left but for all non-significantly modulated *post* vs *pre* (black,  $n=55$ ) or significantly decreased *post* vs *pre* (red,  $n=4$ ) neurons. (f) Same as c but for example neurons recorded from the DMS of D1R-Cre mice. (g) Same as d but for all neurons ( $n=50$ ) recorded from the DMS of D1R-Cre mice ( $n=2$ ). Left (*on-pre*): non-significant ( $n=27$ ), significantly decreased ( $n=21$ ), or increased ( $n=2$ ). Right (*post-pre*): non-significant ( $n=46$ ), significantly decreased ( $n=2$ ) or increased ( $n=2$ ). A Bonferroni-corrected significance threshold was used to account for multiple neuron comparisons ( $p < 0.001$ , or  $p = 0.05/50$  neuron comparisons). (h) same as e but for neurons recorded from the DMS of D1R-Cre mice.



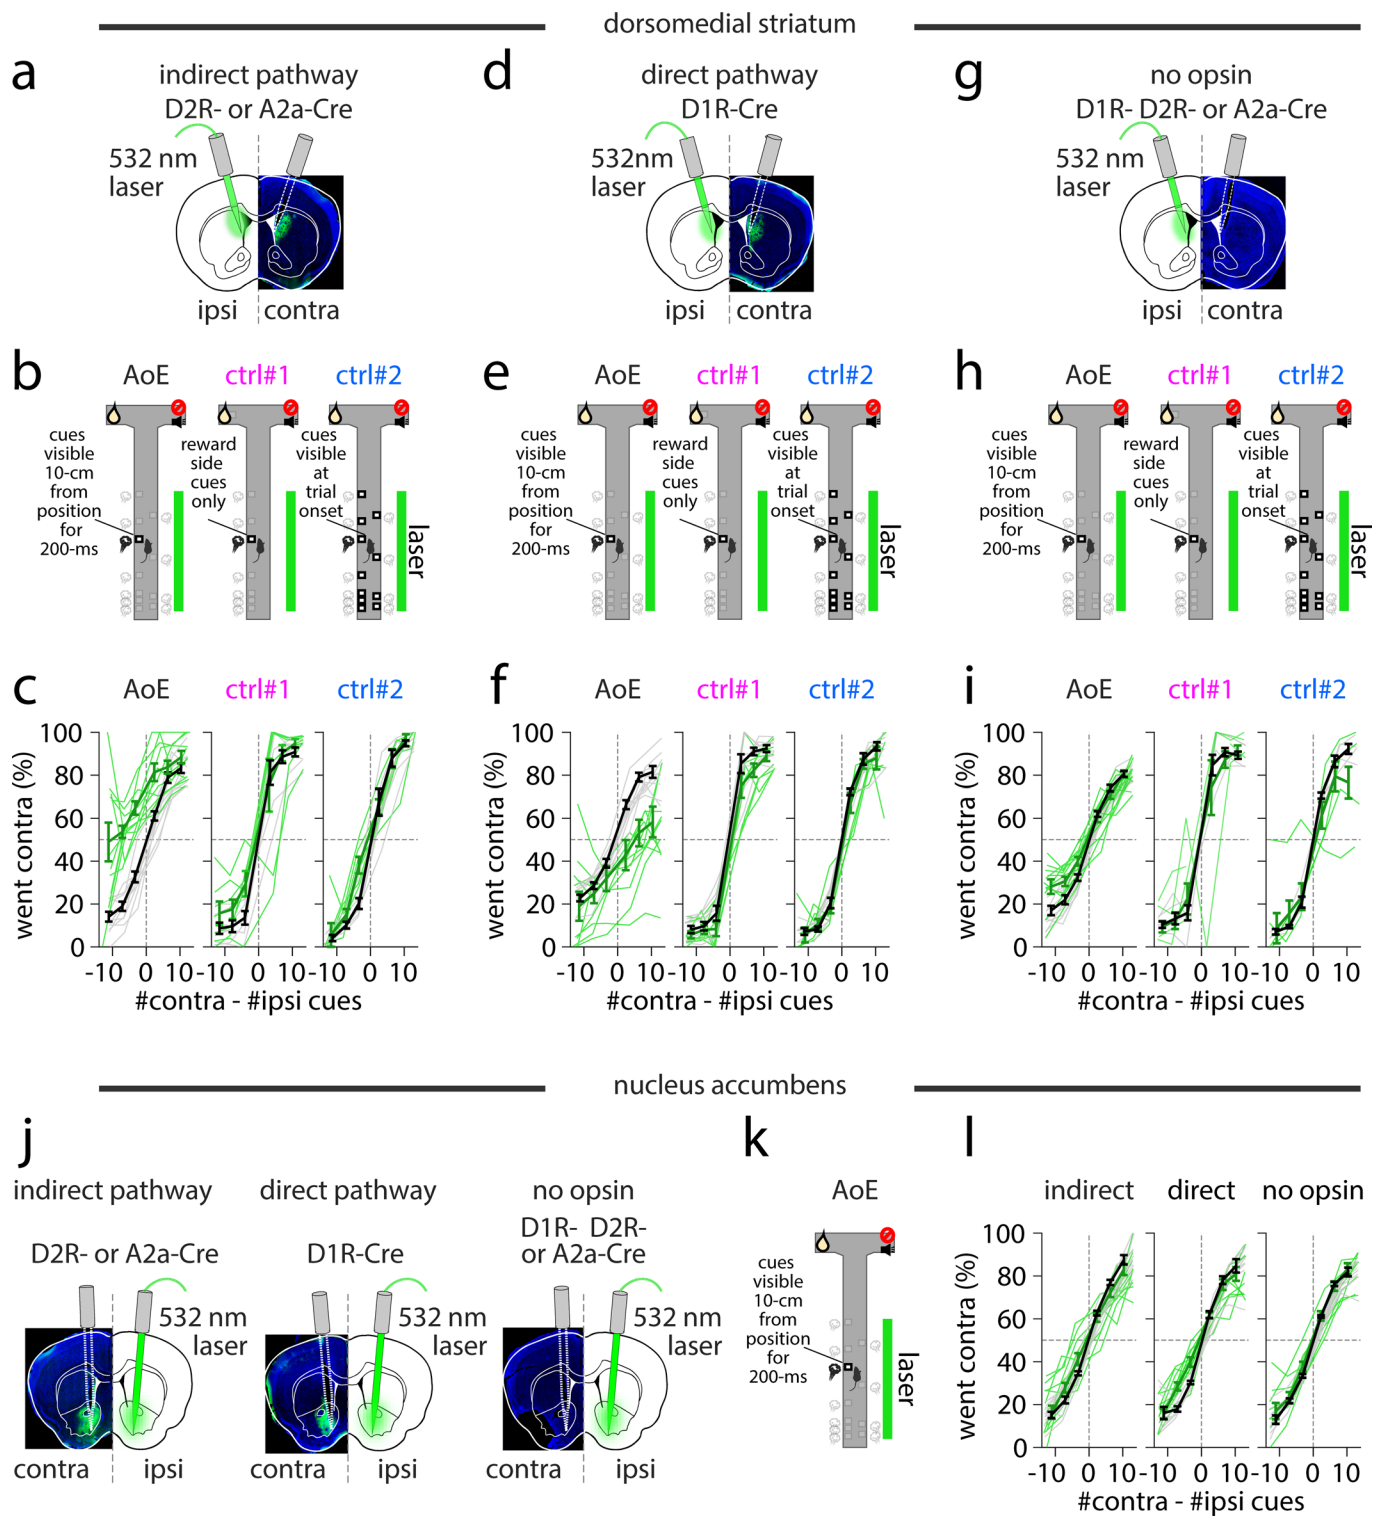
**Extended Data Fig. 2 | Non-significant motor effects of DMS pathway inhibition compared to non-opsin expressing controls during virtual corridor navigation.** (a) Schematic of virtual corridor and unilateral delivery of 532-nm light (5-mW) restricted to 0-200 cm. (b) Difference in average y-velocity (cm/s) during laser on and off trials (on-off) for mice receiving indirect ( $n = 7$  mice,  $n = 1,712$  laser off and  $n = 1,288$  laser on trials) or direct ( $n = 6$  mice,  $n = 1,088$  laser off and  $n = 757$  laser on trials) pathway inhibition of the DMS, or DMS illumination alone (no opsin;  $n = 5$  mice,  $n = 1,178$  laser off and  $n = 827$  laser on trials).  $p$ -value denotes significance of one-way ANOVA of group on delta y-velocity ( $p = 0.98$ ,  $F_{2,13} = 0.02$ ). (c) Same as b but for difference in x-position (cm, on-off) contralateral to the laser hemisphere ( $p = 0.60$ ,  $F_{2,13} = 0.53$ ). (d) Same as c but for difference in view angle (deg, on-off) contralateral to the laser hemisphere ( $p = 0.20$ ,  $F_{2,13} = 1.90$ ). (e) Same as c but for difference in mean standard deviation in view angle (deg, on-off). The mean of the standard deviation in view angles sampled in 5-cm steps from 0-300 cm were calculated per trial, and then averaged across all laser off (or on) trials for a mouse ( $p = 0.94$ ,  $F_{2,16} = 0.06$ ). Indirect:  $n = 7$  mice,  $n = 2,109$  laser off and  $n = 1,574$  laser on trials; direct:  $n = 6$  mice,  $n = 1,330$  laser off and  $n = 930$  laser on trials; no opsin:  $n = 6$  mice,  $n = 1,688$  laser off and  $n = 1,199$  laser on trials). (f) As in e but for difference in total distance travelled (cm, on-off) to complete a trial ( $p = 0.93$ ,  $F_{2,16} = 0.08$ ). (g) As in e but for the difference in percentage of trials with excess travel (defined as  $>10\%$  of corridor length to reward, or  $>330$  cm) ( $p = 0.76$ ,  $F_{2,18} = 0.28$ ). Throughout solid black lines indicate mean  $\pm 1$  s.e.m. across mice and transparent 'x' denote individual mouse mean throughout.



Extended Data Fig. 3 | See next page for caption.

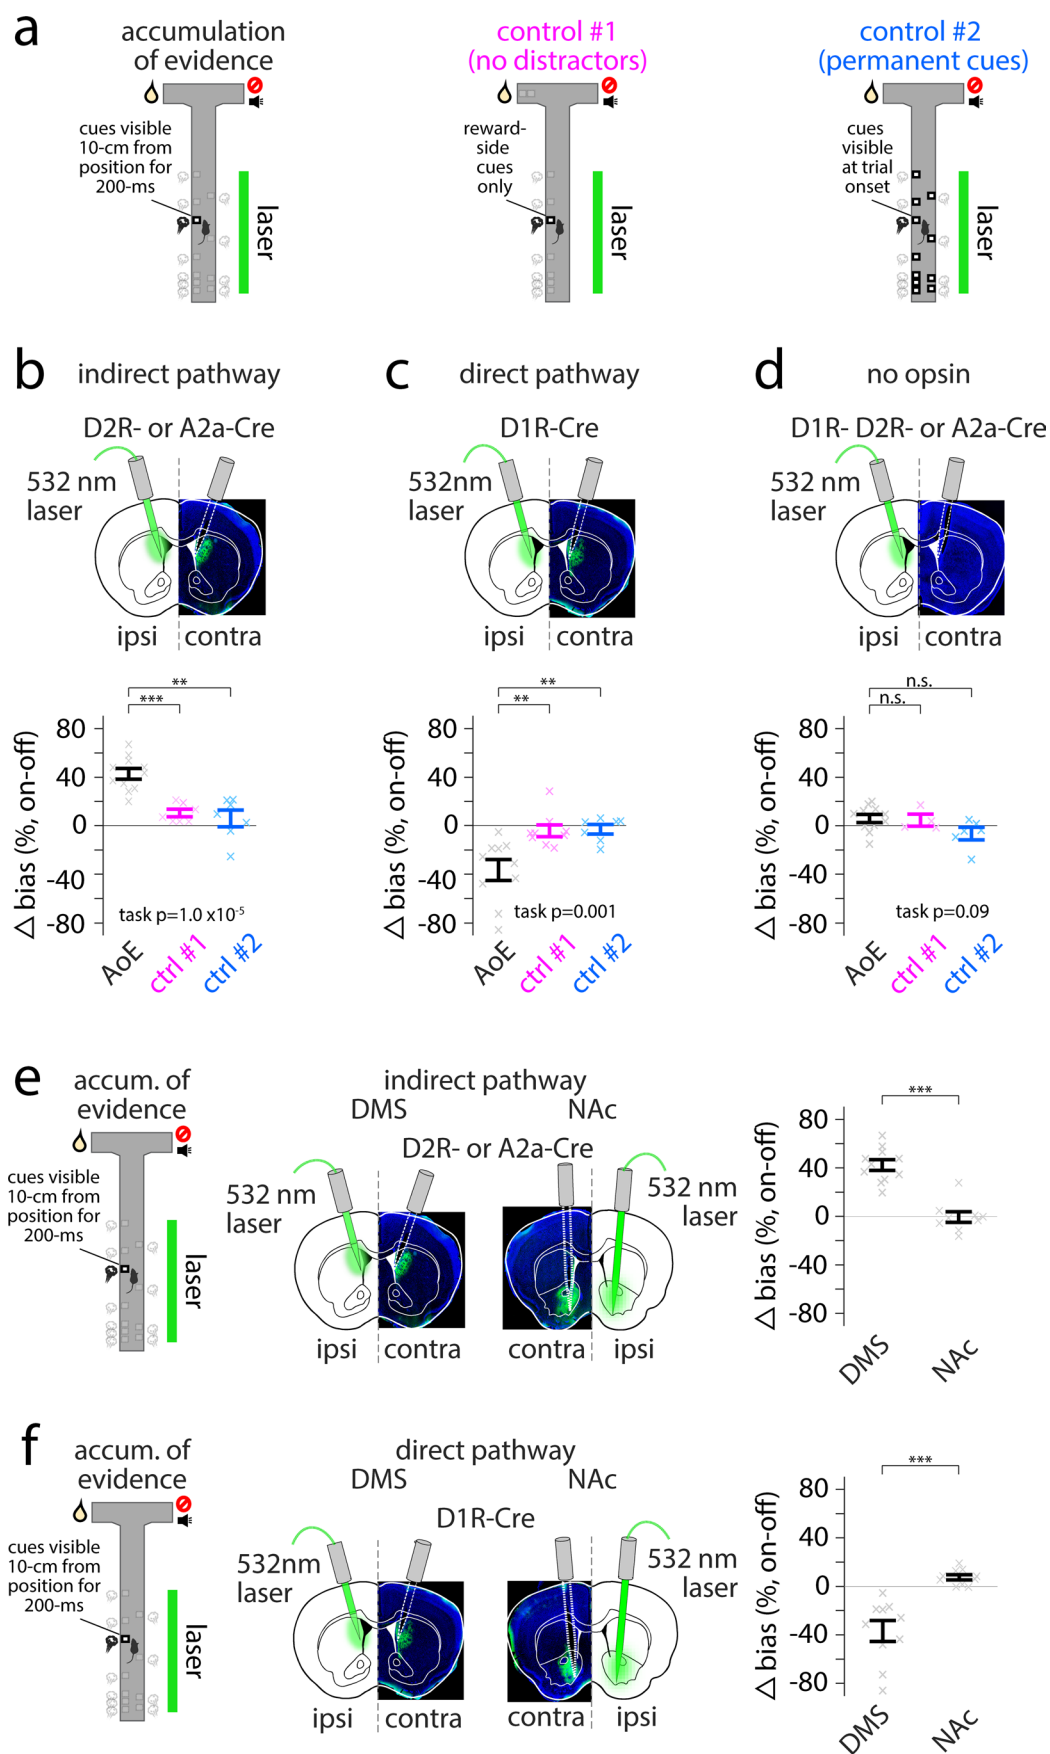
**Extended Data Fig. 3 | Similar motor performance across three virtual reality T-mazes.** Schematic of three virtual reality (VR)-based T-mazes that differ in task requirements. **(b)** Average y-velocity (cm/s) of mice during the cue region (0–200 cm) of the accumulation of evidence task (black,  $n = 32$  mice,  $n = 52,381$  trials), no distractors (ctrl #1) task (magenta:  $n = 32$  mice, 56,783 trials), or permanent cues (ctrl #2) task (cyan:  $n = 20$  mice,  $n = 27,870$  trials). Solid bars denote mean  $\pm 1$  s.e.m. across mice while transparent 'x' denotes individual mouse mean. p-value denotes one-way ANOVA of task on y-velocity ( $p = 0.51$ ,  $F_{2,80} = 0.67$ ). **(c)** Same as **b** but for average x-position (cm) during the cue region (0–200 cm) on left and right choice trials. p-value denotes one-way ANOVA of task on x-position (left choice:  $p = 0.50$ ,  $F_{2,80} = 0.70$ ; right choice:  $p = 0.37$ ,  $F_{2,80} = 1.0$ ). **(d)** Same as **b** but for average view angle (degrees) during the cue region (0–200 cm) on left and right choice trials (left choice:  $p = 0.53$ ,  $F_{2,80} = 0.64$ ; right choice:  $p = 0.70$ ,  $F_{2,80} = 0.37$ ). **(e)** As in **b** but for average percent of trials with excess travel (defined as travel  $> 10\%$  of maze stem, or  $> 330$  cm). Accumulation of evidence:  $n = 32$  mice,  $n = 53,833$  trials; control #2 (no distractors):  $n = 32$  mice,  $n = 60,074$  trials; control #2 (permanent cues):  $n = 20$  mice,  $n = 29,192$  trials. p-value denotes one-way ANOVA of task on excess travel ( $p = 0.06$ ,  $F_{2,81} = 2.9$ ). **(f)** As in **b** but for mean standard deviation in view angle (degrees) per trial ( $n$  as in **e**). p-value denotes one-way ANOVA of task on view angle deviation ( $p = 0.07$ ,  $F_{2,81} = 2.8$ ). **(g)** Average accuracy of decoding left/right choice based on the trial-by-trial x-position (cm) of mice as a function of y-position in the maze (0–300 cm in 25-cm bins). Training and test trial sets were selected within individual mice (80% train, 5-fold cross-validation, re-sampled 10 times). *Left:* Each 'x' depicts decoding accuracy at each y-position bin for individual mice performing the evidence accumulation (black), no distractors (ctrl #1, magenta), or permanent cues (ctrl #2, cyan) tasks. *Right:* Group mean and  $\pm 1$  s.e.m. across mice for each task ( $n$  as in **b**). **(h)** Same as **f** but for average accuracy of decoding left/right choice based on the trial-by-trial view angle (degrees) of mice ( $n$  as in **b**). **(i)** Average accuracy of decoding left/right choice based on the trial-by-trial x-position (cm) of mice as a function of y-position in the maze (0–300 cm in 25-cm bins). Training trial sets were randomly selected across all mice (50% total trials, re-sampled 50 times) performing either the accumulation of evidence (*left*, AoE, black), no distractors (*middle*, ctrl#1, magenta), or permanent cues (*right*, ctrl#2, cyan) tasks. Testing trial sets were the 50% of held-out trials in the task used for training, or all trials in the alternate tasks. *Top:* Each 'x' depicts average decoding accuracy across all training/test sets at each y-position bin for individual mice performing the evidence accumulation (black), no distractors (ctrl #1, magenta), or permanent cues (ctrl #2, cyan) tasks. *Right:* Group mean and  $\pm 1$  s.e.m. across mice for each task ( $n$  as in **a**). **(j)** Same as **i** but for average accuracy of decoding left/right choice based on the trial-by-trial view angle (degrees) of mice ( $n$  as in **b**).





Extended Data Fig. 4 | See next page for caption.

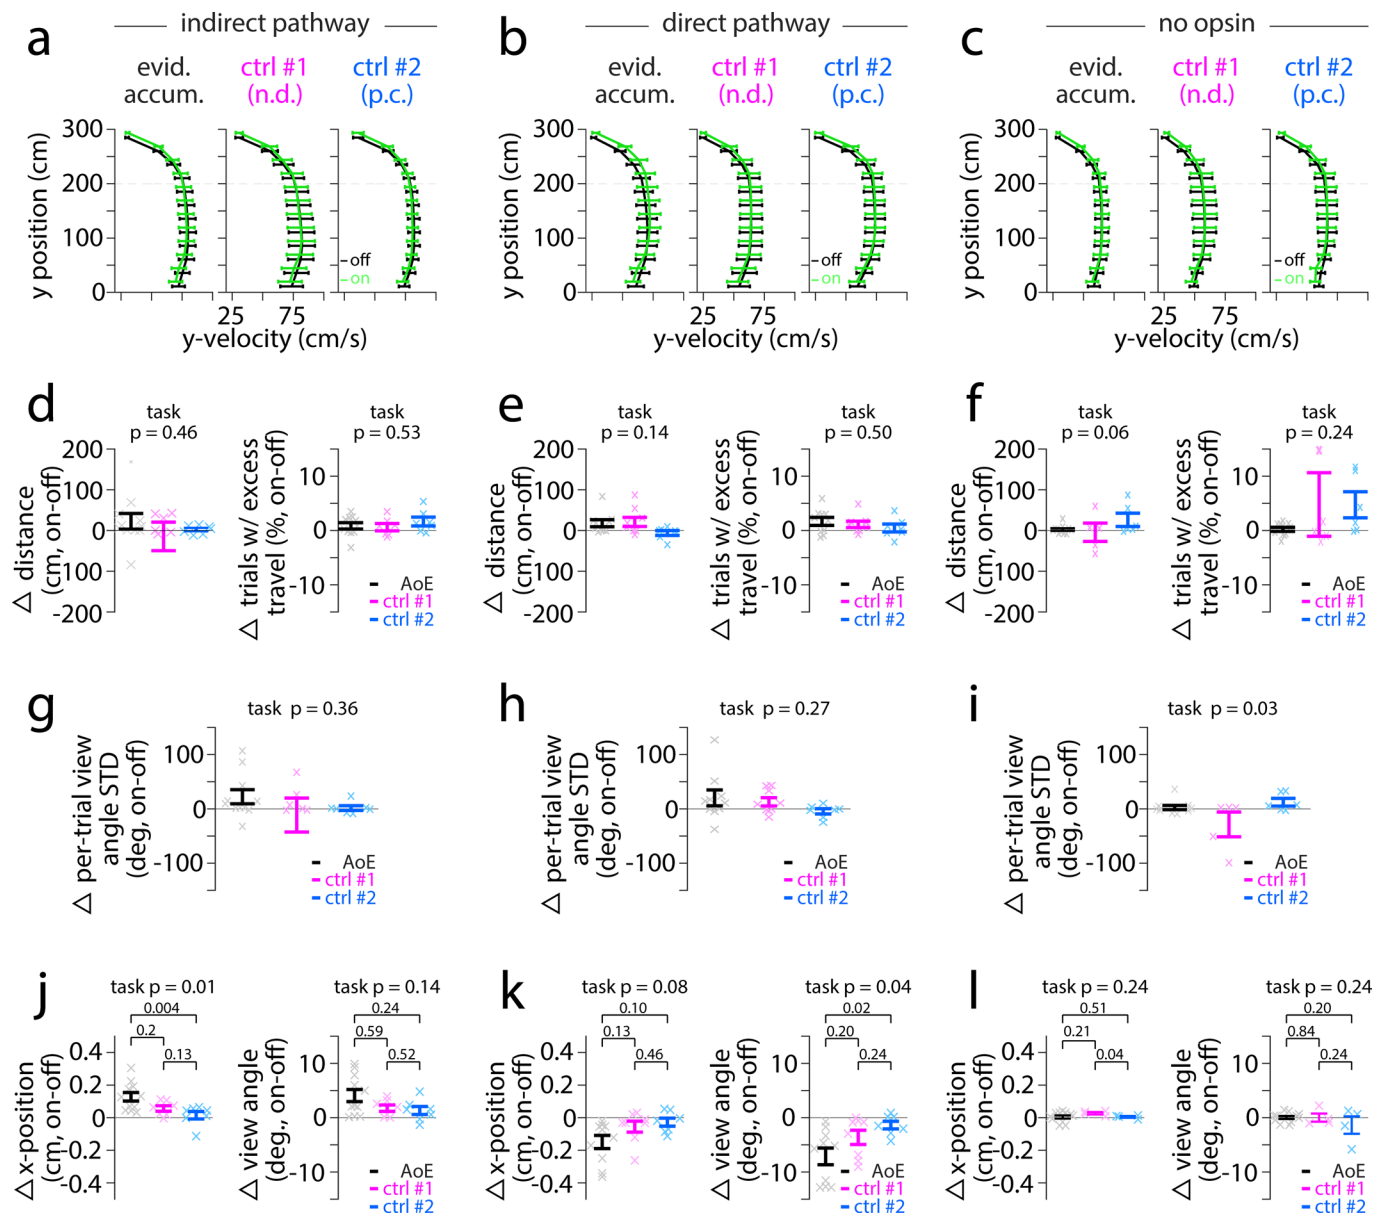
**Extended Data Fig. 4 | Effects of pathway-specific DMS and NAc inhibition on psychometric performance across virtual reality tasks.** (a) Schematic of unilateral indirect pathway DMS inhibition with choice defined ipsilateral or contralateral to the hemisphere receiving 532-nm laser illumination. (b) Schematic of three virtual reality based decision-making tasks (*left*: accumulation of evidence; *middle*: control #1, no distractors; *right*: control #2, permanent cues) and laser illumination restricted to the cue region (0-200 cm). (c) Percent of contralateral choice trials as a function of the difference in sensory cues (contralateral-ipsilateral) binned in increments of 5 from -15 to 15. Transparent lines indicate individual mouse mean during laser off (grey) and on (green) trials for mice receiving indirect-pathway DMS inhibition during the evidence accumulation (black, *left*), no distractors (magenta, ctrl#1, *middle*), or permanent cues (cyan, ctrl#2, *right*). Thick lines indicate mean  $\pm 1$  s.e.m. across mice at each evidence bin during laser off (black) and on (green) trials. (d) Same as a but for mice receiving unilateral direct pathway DMS inhibition. (e) same as b. (f) Same as c but for mice receiving direct pathway DMS inhibition. (g) Same as a but for mice receiving unilateral DMS illumination in the absence of NpHR (no opsin). (h) Same as b. (i) same as c but for mice receiving unilateral DMS illumination in the absence of NpHR (no opsin). (j) Schematic of unilateral inhibition of NAc indirect (*left*) or direct (*middle*) pathway, or NAc illumination in the absence of NpHR (no opsin). (k) Schematic of accumulation of evidence task and delivery of 532-nm light during the cue region (0-200 cm). (l) As in c but for psychometric comparison between groups receiving NAc indirect or direct pathway inhibition, or NAc illumination in the absence of NpHR (no opsin).



Extended Data Fig. 5 | See next page for caption.

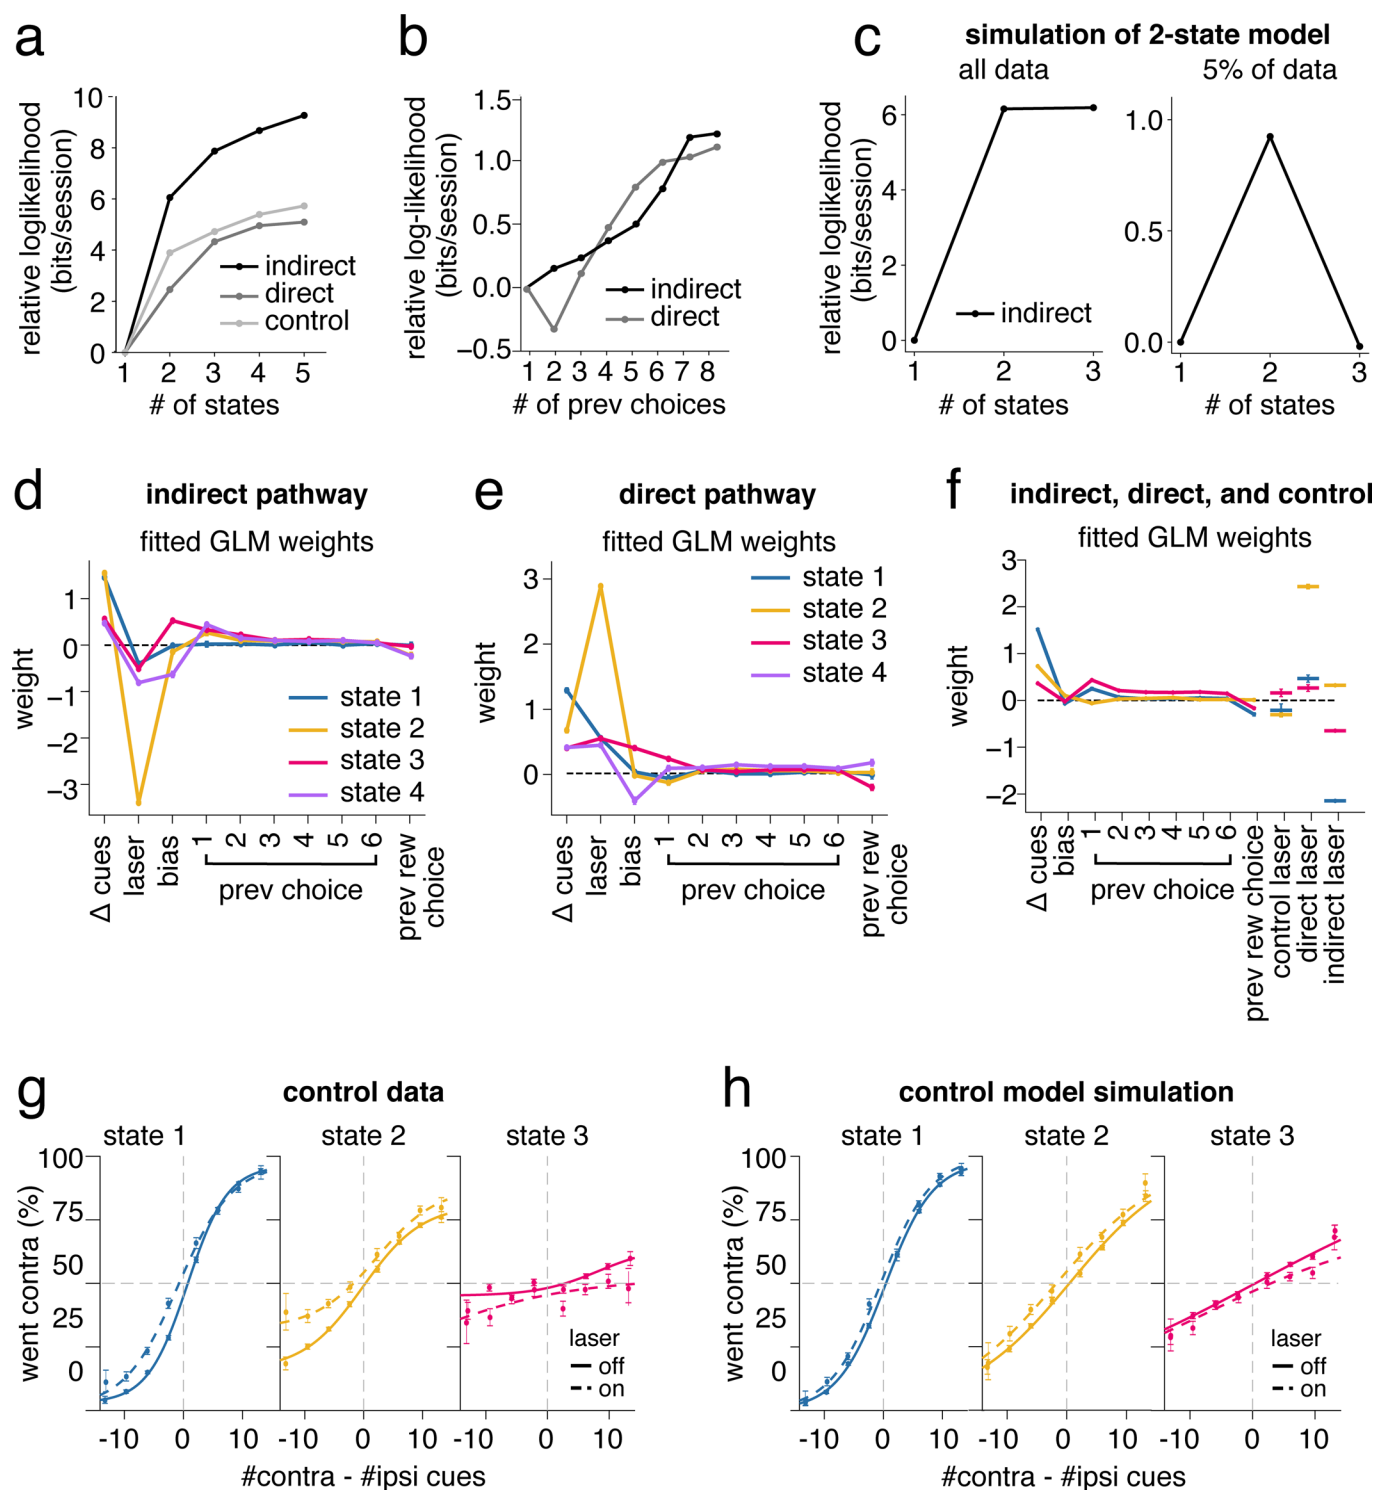
**Extended Data Fig. 5 | Effects of pathway-specific DMS inhibition on choice are larger in the most demanding task, and stronger than effects of pathway-specific NAc inhibition.** **(a)** Schematic of three virtual reality based decision-making tasks (*left*: accumulation of evidence; *middle*: control #1, no distractors; *right*: control #2, permanent cues). **(b)** Schematic of unilateral indirect pathway DMS inhibition with choice defined ipsilateral or contralateral to the hemisphere receiving 532-nm laser illumination (*top*). Difference in choice bias (%; contralateral - ipsilateral) between laser on and off trials (on-off) in mice performing the accumulation of evidence (AoE, black), no distractors (ctrl #1, magenta), or permanent cues (ctrl #2, cyan) tasks. p-value denotes one-way ANOVA of task on delta (on-off) choice bias ( $p = 1.0 \times 10^{-5}$ ,  $F_{2,22} = 20.2$ ). Post-hoc comparisons reflect unpaired, two-tailed Wilcoxon ranksum tests on delta (on-off) choice bias (AoE,  $n = 11$ , vs ctrl #1,  $n = 7$ :  $p = 8.0 \times 10^{-4}$ ,  $z = 3.4$ ; AoE vs ctrl #2,  $n = 7$ :  $p = 0.001$ ,  $z = 3.3$ ). **(c)** Same as **b** but for direct pathway DMS inhibition. p-value denotes one-way ANOVA of task on delta (on-off) choice bias ( $p = 0.001$ ,  $F_{2,23} = 9.4$ ). Post-hoc comparisons reflect two-tailed, unpaired Wilcoxon ranksum tests (AoE,  $n = 10$ , vs ctrl #1,  $n = 9$ :  $p = 0.002$ ,  $z = -3.0$ ; AoE vs ctrl #2,  $n = 7$ :  $p = 0.005$ ,  $z = -2.8$ ). **(d)** Same as **b** but for DMS illumination in the absence of NpHR (no opsin). p-value denotes one-way ANOVA of task on delta (on-off) choice bias ( $p = 0.09$ ,  $F_{2,16} = 2.8$ ). Post-hoc comparisons reflect two-tailed, unpaired Wilcoxon ranksum tests (AoE,  $n = 11$ , vs ctrl #1,  $n = 4$ :  $p = 0.65$ ,  $z = 0.46$ ; AoE vs ctrl #2,  $n = 6$ :  $p = 0.06$ ,  $z = 1.8$ ). **(e)** Schema of evidence accumulation task (*left*), unilateral inhibition of indirect pathway in the DMS (*middle left*) or NAc (*middle right*), and delta (on-off) choice bias in mice receiving indirect pathway DMS ( $n = 11$ ) or NAc ( $n = 9$ ) inhibition (*right*). Statistical comparison reflects two-tailed, unpaired Wilcoxon ranksum test (DMS vs NAc:  $p = 2.6 \times 10^{-4}$ ,  $z = 3.6$ ). **(f)** Same as **e** but for direct pathway DMS ( $n = 10$ ) or NAc ( $n = 10$ ) inhibition. Statistical comparison reflects two-tailed, unpaired Wilcoxon ranksum test (DMS vs NAc:  $p = 1.8 \times 10^{-4}$ ,  $z = -3.7$ ). Throughout solid bars denote mean  $\pm 1$  s.e.m. across mice and transparent 'x' denote individual mouse means.





Extended Data Fig. 6 | See next page for caption.

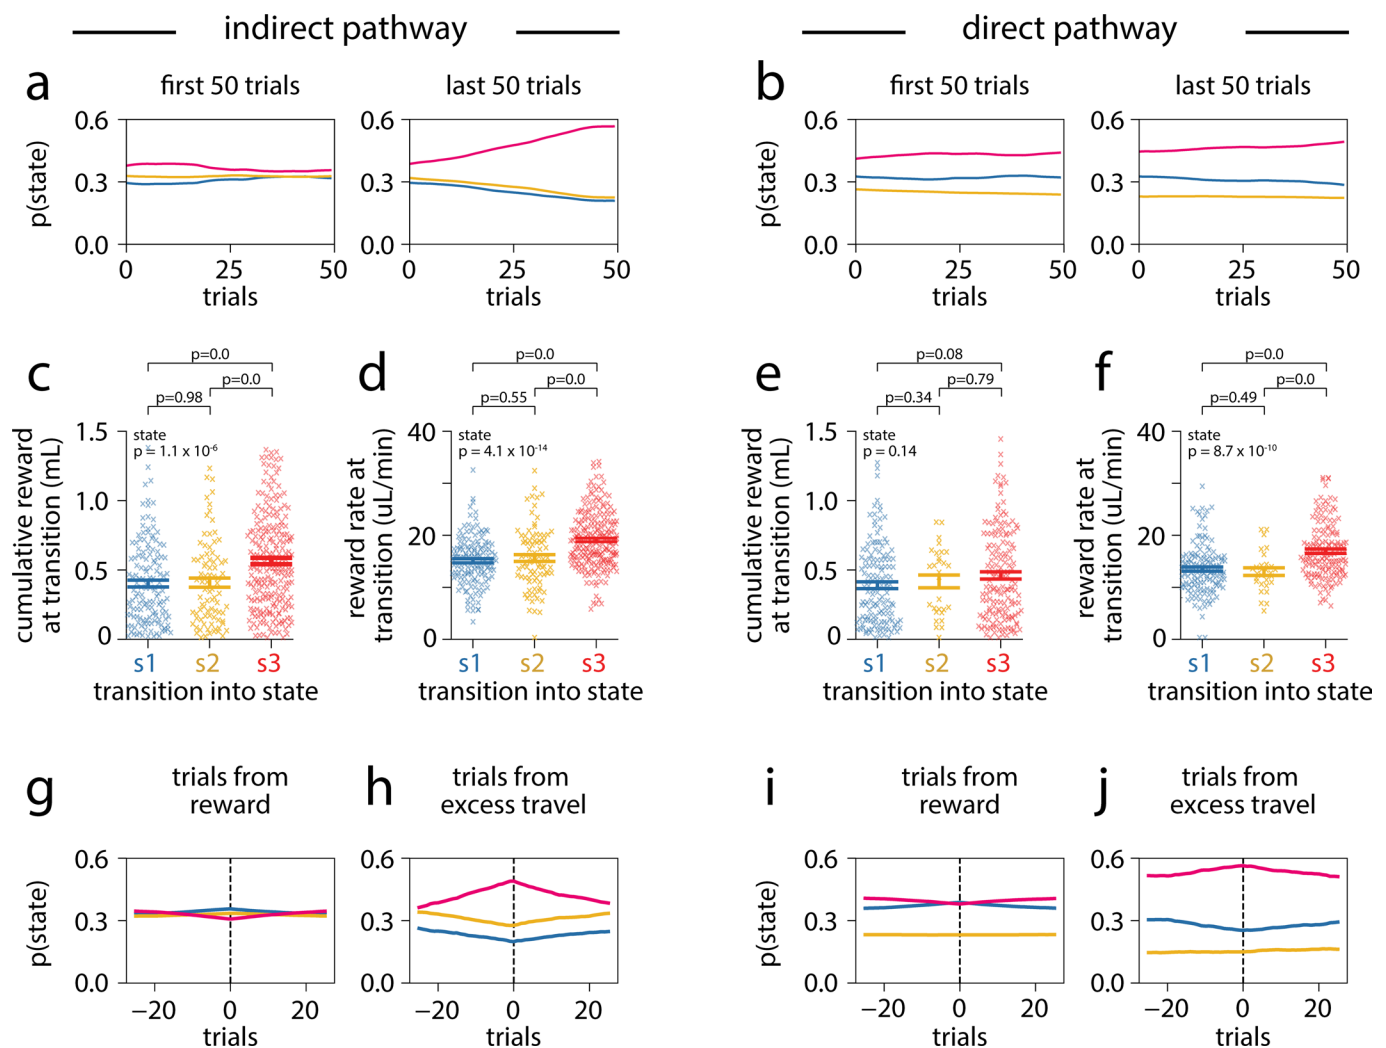
**Extended Data Fig. 6 | Inhibition of DMS pathways has limited impact on motor performance across VR-based decision-making tasks.** (a) Mean  $\pm 1$  s.e.m. y-velocity (cm/s) as a function of y-position (0–300 cm in 25 cm bins) during laser off (black) or laser on (green) trials across mice receiving DMS indirect pathway inhibition during the evidence accumulation (left:  $n = 11$  mice,  $n = 16,935$  laser off and  $n = 3,390$  laser on trials), no distractors (middle, ctrl #1:  $n = 7$  mice,  $n = 13,706$  laser off and  $n = 3,288$  laser on trials) or permanent cues (right, ctrl #2:  $n = 6$  mice,  $n = 4,033$  laser off and  $n = 929$  laser on trials). (b) Same as a but for mice receiving direct pathway inhibition during the evidence accumulation (left:  $n = 10$  mice,  $n = 14,030$  laser off and  $n = 3,103$  laser on trials), no distractors (middle, ctrl #2:  $n = 8$  mice,  $n = 14,647$  laser off and  $n = 3,682$  laser on trials) or permanent cues (right, ctrl #3:  $n = 7$  mice,  $n = 6,061$  laser off and  $n = 1,494$  laser on trials) tasks. (c) Same as a but for mice receiving DMS illumination in the absence of NpHR (no opsin) during the evidence accumulation (left:  $n = 11$  mice,  $n = 21,422$  laser off and  $n = 5,113$  laser on trials), no distractors (middle, ctrl #1:  $n = 4$  mice,  $n = 3,654$  laser off and  $n = 901$  laser on trials), or permanent cues (right, ctrl #2:  $n = 4$  mice,  $n = 3,975$  laser off and  $n = 923$  laser on trials) tasks. (d) Mean  $\pm 1$  s.e.m. in delta (on-off) distance (cm) traveled (left) and delta (on-off) trials (%) with excess travel greater than 10% of maze stem (or  $> 330$  cm) (right) in mice receiving indirect pathway inhibition during the evidence accumulation (black,  $n = 11$  mice,  $n = 22,090$  laser off and  $n = 4,378$  laser on trials), no distractors (magenta,  $n = 7$  mice,  $n = 14,826$  laser off and  $n = 3,591$  laser on trials), or permanent cues ( $n = 6$  mice,  $n = 4,447$  laser off and  $n = 1,050$  laser on trials) tasks. p-value denotes one-way ANOVA of task on delta (on-off) distance ( $p = 0.45$ ,  $F_{2,22} = 0.81$ ) or excess travel ( $p = 0.52$ ,  $F_{2,22} = 0.66$ ). (e) Same as d but for delta (on-off) distance (cm) traveled (left) or delta percent trials with excess travel (right) in mice receiving direct pathway inhibition during the evidence accumulation (black,  $n = 10$  mice,  $n = 20,914$  laser off and  $n = 4,721$  laser on trials), no distractors (magenta,  $n = 9$  mice,  $n = 15,778$  laser off and  $n = 3,992$  laser on trials), or permanent cues ( $n = 7$  mice,  $n = 6,430$  laser off and  $n = 1,591$  laser on trials) tasks. p-value denotes one-way ANOVA of task on delta (on-off) distance ( $p = 0.13$ ,  $F_{2,23} = 2.2$ ) or excess travel ( $p = 0.50$ ,  $F_{2,23} = 0.71$ ). (f) Same as d but for delta (on-off) in distance (cm) traveled (left) or percent trials with excess travel (right) in mice receiving DMS illumination in the absence of NpHR (no opsin) during the evidence accumulation (black,  $n = 11$  mice,  $n = 28,557$  laser off and  $n = 6,772$  laser on trials), no distractors (magenta,  $n = 5$  mice,  $n = 4,118$  laser off and  $n = 1,002$  laser on trials), or permanent cues ( $n = 6$  mice,  $n = 4,360$  laser off and  $n = 1,038$  laser on trials) tasks. p-value denotes one-way ANOVA of task on delta (on-off) distance ( $p = 0.06$ ,  $F_{2,19} = 3.3$ ) or excess travel ( $p = 0.23$ ,  $F_{2,19} = 1.6$ ). (g) Same as d but for delta (on-off) in per-trial standard deviation in view angle in mice receiving DMS indirect pathway inhibition across tasks ( $p = 0.34$ ,  $F_{2,22} = 1.1$ ,  $n$  as in d). (h) Same as g but for mice receiving DMS direct pathway inhibition across tasks ( $p = 0.27$ ,  $F_{2,23} = 1.4$ ,  $n$  as in e). (i) Same as g but for mice receiving DMS illumination (no opsin) in the absence of NpHR ( $p = 0.03$ ,  $F_{2,19} = 4.3$ ,  $n$  as in f). (j) Delta (on-off) x-position (cm) (left) or view angle (degrees) (right) during the cue region (0–200 cm) in mice receiving DMS indirect pathway inhibition during the accumulation of evidence (black), no distractors (control #1, magenta), or permanent cues (control #2, cyan) tasks ( $n$  as in a). One-way ANOVA of task on delta (on-off) x-position ( $p = 0.01$ ,  $F_{2,22} = 5.6$ ). Post-hoc, two-tailed, unpaired Wilcoxon ranksum test on delta (on-off) x-position (AoE v control #1:  $p = 0.2$ ,  $z = 1.3$ ; AoE v control #2:  $p = 0.004$ ,  $z = 2.9$ ; control #1 v control #2:  $p = 0.13$ ,  $z = 1.5$ ). One-way ANOVA of task on delta (on-off) view angle ( $p = 0.14$ ,  $F_{2,22} = 2.2$ ). Post-hoc, two-tailed, unpaired Wilcoxon ranksum test on delta (on-off) view angle (AoE v control #1:  $p = 0.58$ ,  $z = 0.5$ ; AoE v control #2:  $p = 0.24$ ,  $z = 1.78$ ; control #1 v control #2:  $p = 0.52$ ,  $z = 0.6$ ). (k) Same as j but for mice receiving DMS direct pathway inhibition ( $n$  as in b). One-way ANOVA of task on delta (on-off) x-position ( $p = 0.08$ ,  $F_{2,23} = 2.8$ ). Post-hoc, two-tailed unpaired Wilcoxon ranksum test on delta (on-off) x-position (AoE v control #1:  $p = 0.13$ ,  $z = -1.5$ ; AoE v control #2:  $p = 0.1$ ,  $z = -1.6$ ; control #1 v control #2:  $p = 0.46$ ,  $z = -0.7$ ). One-way ANOVA of task on delta (on-off) view angle ( $p = 0.02$ ,  $F_{2,23} = 3.6$ ). Post-hoc, two-tailed, unpaired Wilcoxon ranksum test on delta (on-off) view angle (AoE v control #1:  $p = 0.21$ ,  $z = -1.3$ ; AoE v control #2:  $p = 0.03$ ,  $z = -2.1$ ; control #1 v control #2:  $p = 0.24$ ,  $z = -1.6$ ). (l) Same as j but for mice receiving DMS illumination in the absence of NpHR (no opsin,  $n$  as in c). One-way ANOVA of task on delta (on-off) x-position ( $p = 0.24$ ,  $F_{2,18} = 1.54$ ). Post-hoc, two-tailed, unpaired Wilcoxon ranksum test on delta (on-off) x-position (AoE v control #1:  $p = 0.21$ ,  $z = -1.24$ ; AoE v control #2:  $p = 0.51$ ,  $z = 0.06$ ; control #1 v control #2:  $p = 0.04$ ,  $z = 2.0$ ). One-way ANOVA of task on delta (on-off) view angle ( $p = 0.23$ ,  $F_{2,18} = 1.56$ ). Post-hoc, two-tailed, unpaired Wilcoxon ranksum test on delta (on-off) view angle (AoE v control #1:  $p = 0.84$ ,  $z = 0.19$ ; AoE v control #2:  $p = 0.20$ ,  $z = 1.2$ ; control #1 v control #2:  $p = 0.24$ ,  $z = 1.7$ ). Throughout solid bars denote mean  $\pm 1$  s.e.m. and transparent 'x' indicates individual mouse mean.



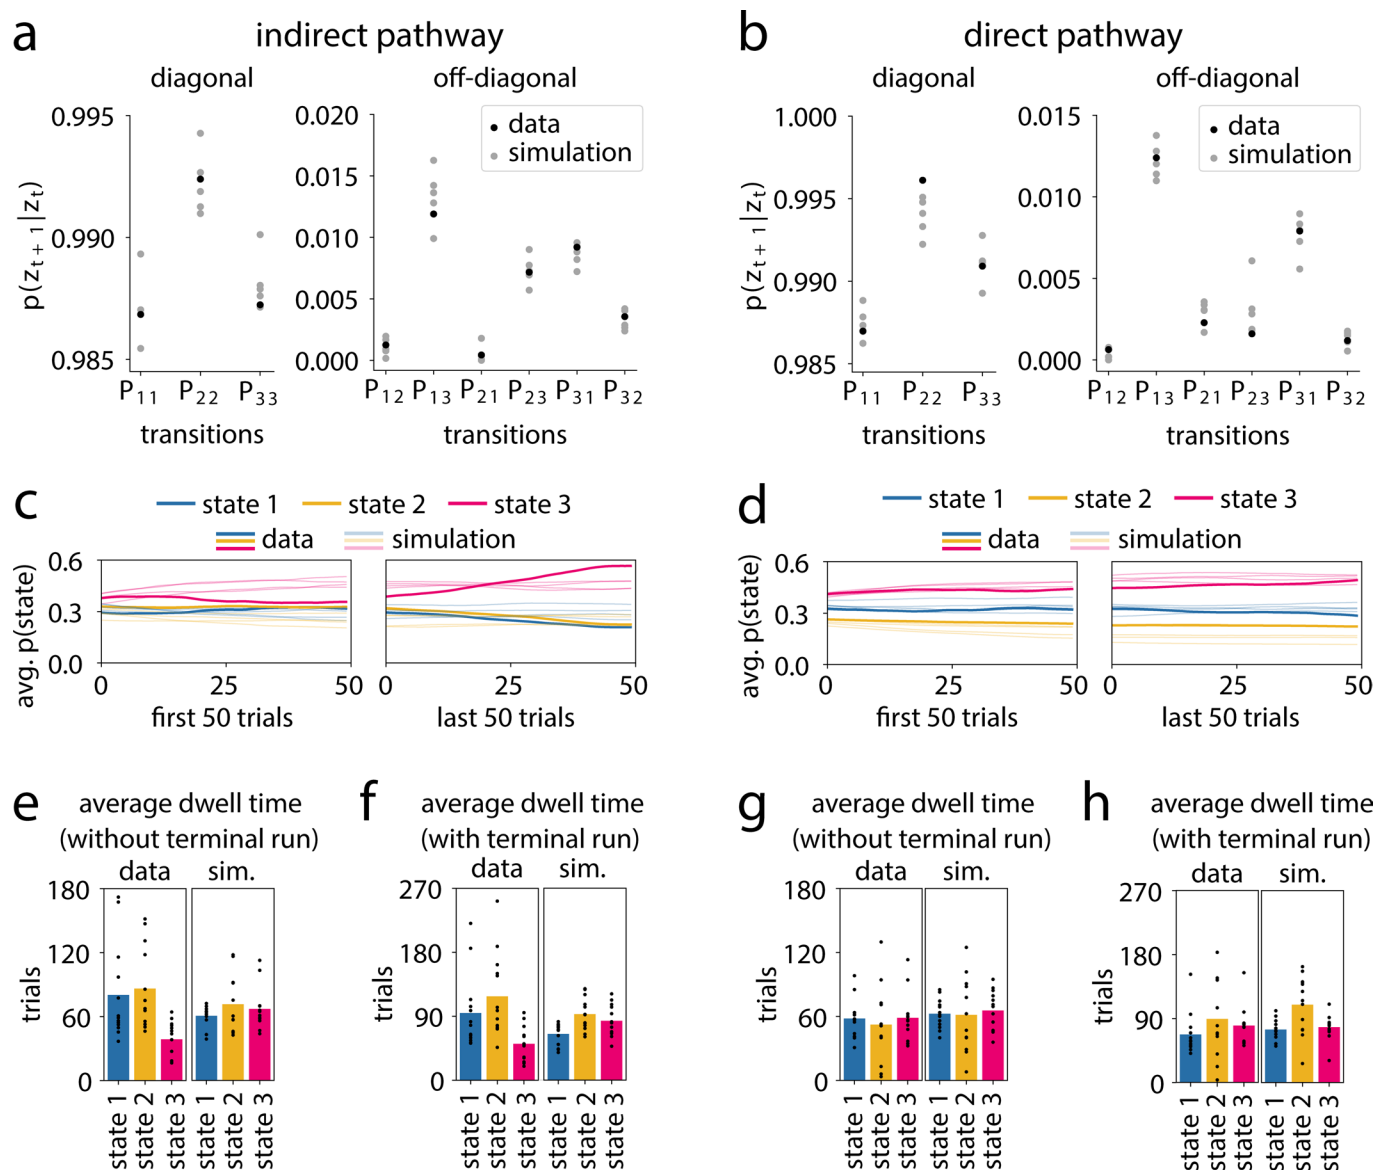
Extended Data Fig. 7 | See next page for caption.

**Extended Data Fig. 7 | Model selection and control data analyses for the GLM-HMM.** **(a)** Comparison of the log-likelihood of the data using GLM-HMMs with different numbers of states for mice inhibited in the DMS direct pathway (dark gray), or indirect pathway (light gray), and mice without DMS opsin (black). All values are relative to the log-likelihood of the standard GLM (1-state GLM-HMM). Values are calculated in bits per session (see **Methods**). Solid curves denote mean  $\pm$ s.e.m. of five different test sets. Held-out data for test sets was selected as a random 20% of sessions, using an approximately equal number of sessions for each mouse. **(b)** Same as **a** but with different numbers of previous choice covariates using a three-state GLM-HMM. **(c)** Comparison of the log-likelihood of simulated data using GLM-HMMs with different numbers of states. Data was simulated from a two-state GLM-HMM that had been fit to data for mice inhibited in the indirect pathway of the DMS and then cross-validation performed either on the entire simulated dataset (~54000 trials, left) or a subset of 5% of the data (2600 trials, right). All values are relative to the log-likelihood of a GLM (one-state GLM-HMM). Values are calculated in bits per session (see **Methods**). Solid curves denote the average of five different test sets. Held-out data for test sets was selected as a random 20% of sessions. Performing cross validation on a small subset of the data serves to demonstrate that the log-likelihood does in fact decrease as the model starts to overfit. This is difficult to see with large datasets where overfitting is less of a concern and therefore the log-likelihood begins to flatten rather than decrease. **(d)** Fitted GLM weights for the four-state GLM-HMM using aggregated data from all mice inhibited in the indirect pathway of the DMS. Error bars denote ( $\pm 1$ ) posterior standard deviation for each weight. The magnitude of the weight represents the relative importance of that covariate in predicting choice, whereas the sign of the weight indicates the side bias. **(e)** Same as **d** but for mice inhibited in the DMS direct pathway. **(f)** GLM weights fitted to a concatenated dataset consisting of the indirect, direct, and control (no opsin) groups. Solid lines on the left connect covariates that are shared across groups. Horizontal marks on the right denote laser weights, which were learned separately for each group. Error bars denote ( $\pm 1$ ) posterior standard deviation of each weight. **(g)** Percent of contralateral choice based on the difference in contralateral versus ipsilateral cues in each trial for mice in the control (no opsin) group. To compute psychometric functions, trials were assigned to each state by taking the maximum of the model's posterior state probabilities on each trial. Error bars denote  $\pm 1$  s.e.m. for light-off (solid) and light-on (dotted) trials. Solid curves denote logistic fits to the concatenated data across mice for light-off (solid) and light-on (dotted) trials. **(h)** Same as **f** but for data simulated from the model fit to mice in the control group (see **Methods**).

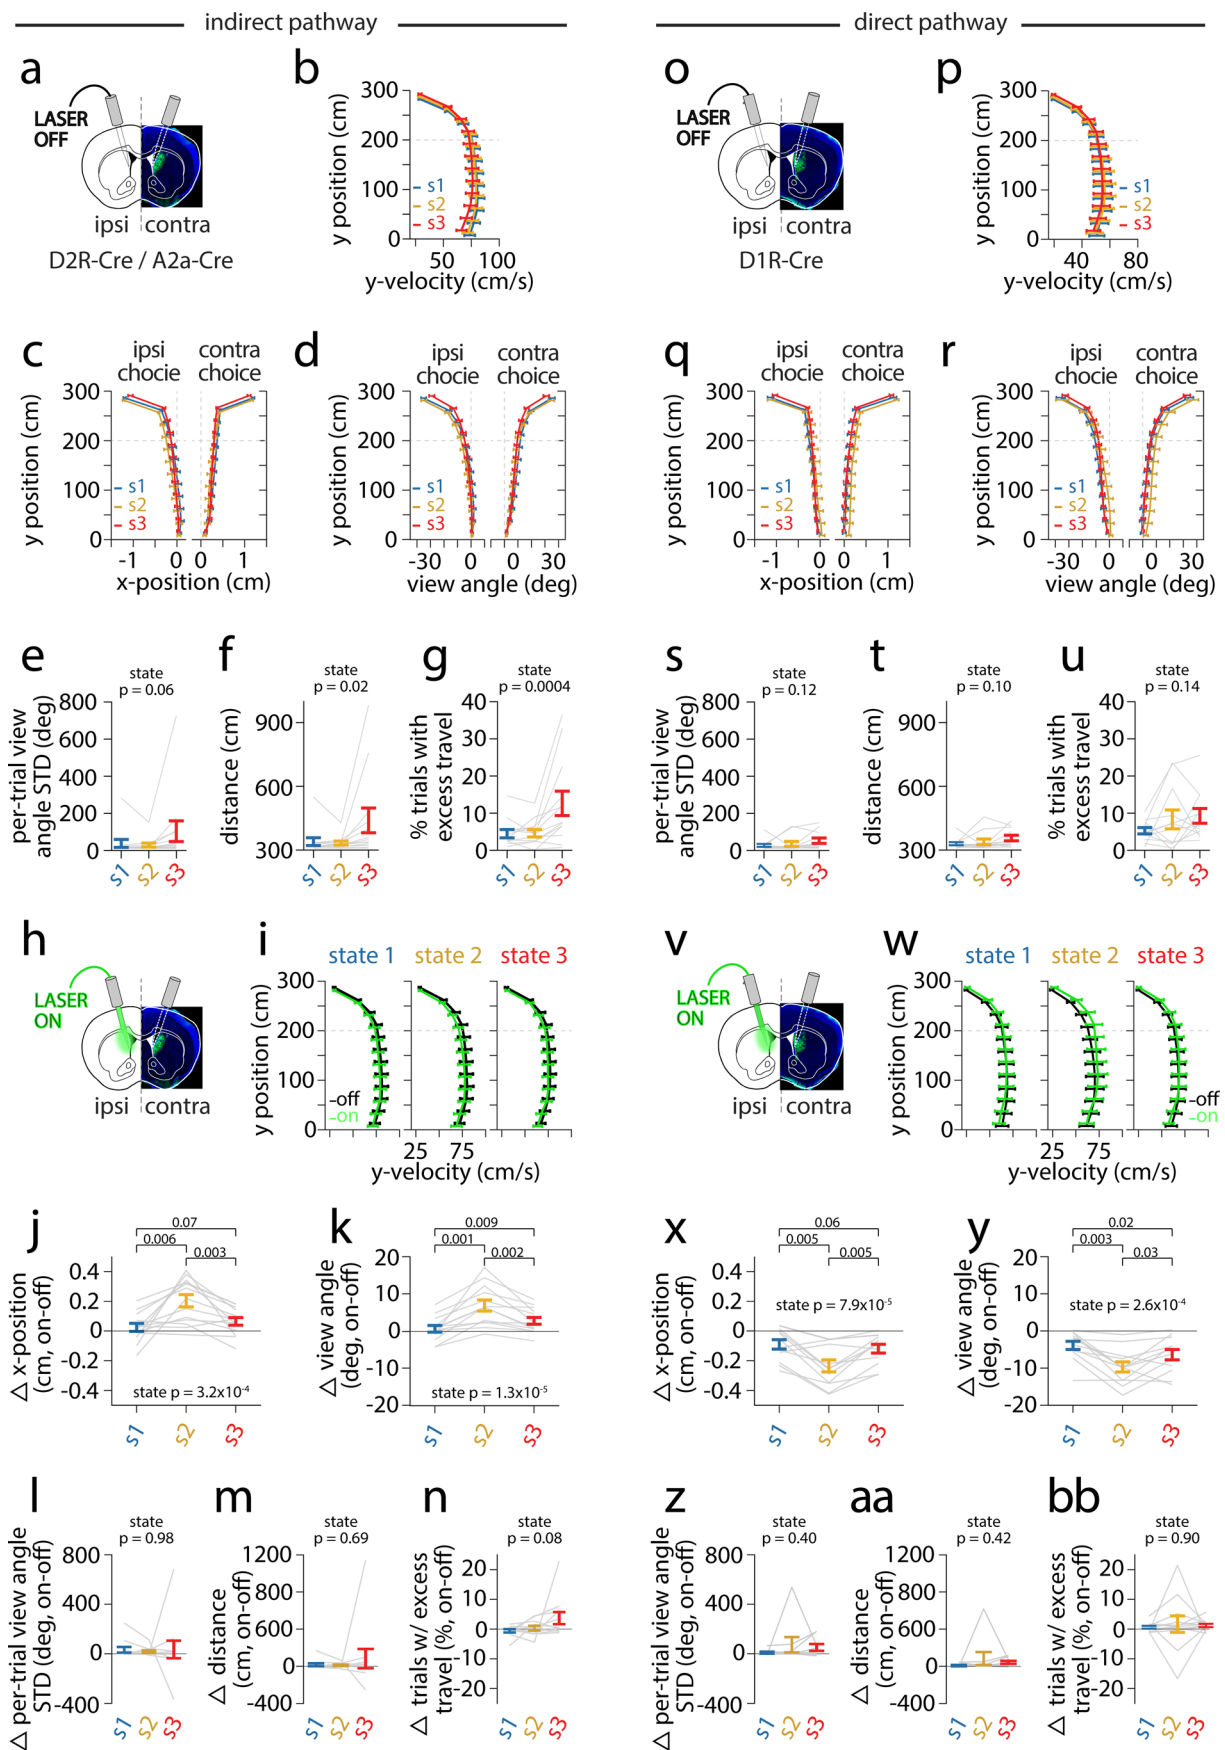




**Extended Data Fig. 8 | GLM-HMM state 3 is associated with indicators of task disengagement.** (a) The mean posterior probability of each state over the first and last 50 trials of a session, averaged across all sessions for mice inhibited in the indirect pathway of the DMS ( $n = 271$  sessions). (b) Same as a but for mice receiving DMS direct pathway inhibition ( $n = 266$  sessions). (c) Mean  $\pm$  s.e.m. of the cumulative reward received in a session prior to transitions into state 1 ( $n = 142$ ), state 2 ( $n = 85$ ), or state 3 ( $n = 237$ ) in the indirect pathway group. One-way ANOVA of transition state on cumulative reward ( $p = 1.0 \times 10^{-6}$ ;  $F_{2,460} = 14.2$ ). Unpaired, two-tailed Wilcoxon ranksum comparison between transition types (state 1 vs 2:  $p = 0.96$ ,  $z = -0.03$ ; state 2 vs 3:  $p = 0$ ,  $z = -3.6$ ; state 1 vs 3:  $p = 0$ ,  $z = -4.5$ ). (d) Mean  $\pm$  s.e.m. of the reward rate (uL/min) in a session prior to transitions into each state for the indirect pathway group. Reward rate was calculated as the sum of reward received from the start of the session up to the transition trial divided by the sum of the duration of all trials from the start of the session up to the transition trial. One-way ANOVA of transition state on reward rate ( $p = 4.1 \times 10^{-14}$ ;  $F_{2,460} = 32.9$ ). Unpaired, two-tailed Wilcoxon ranksum comparison between transition types (state 1 vs 2:  $p = 0.55$ ,  $z = -0.6$ ; state 2 vs 3:  $p = 0$ ,  $z = -4.9$ ; state 1 vs 3:  $p = 0$ ,  $z = -7.4$ ). (e) Same as c but for the direct pathway group (state 1:  $n = 140$ ; state 2:  $n = 29$ ; state 3:  $n = 159$ ). One-way ANOVA of transition state on cumulative reward ( $p = 0.14$ ;  $F_{2,325} = 1.99$ ). Unpaired, two-tailed Wilcoxon ranksum comparison between transition types (state 1 vs 2:  $p = 0.35$ ,  $z = -0.9$ ; state 2 vs 3:  $p = 0.78$ ,  $z = -0.27$ ; state 1 vs 3:  $p = 0.08$ ,  $z = -1.74$ ). (f) Same as d but for the direct pathway group. One-way ANOVA of transition state on reward rate ( $p = 8.7 \times 10^{-10}$ ;  $F_{2,325} = 22.6$ ). Unpaired, two-tailed Wilcoxon ranksum comparison between transition types (state 1 vs 2:  $p = 0.49$ ,  $z = 0.69$ ; state 2 vs 3:  $p = 0.0$ ,  $z = -4.2$ ; state 1 vs 3:  $p = 0.0$ ,  $z = -5.9$ ). (g) The mean posterior probability of each state aligned  $\pm 25$  trials to trials in which reward was received for the indirect pathway group. (h) Same as g but state probability aligned to trials with excess travel (defined as 10% greater than the maze stem, or 330 cm). (i) Same as g but for the direct pathway group. (j) Same as h but for the direct pathway group.



**Extended Data Fig. 9 | Model simulations recapitulate transition and state characteristics of real data.** (a) Transition probabilities of the model fit to data from mice inhibited in the DMS indirect pathway (black) and from five simulated datasets generated from the model fit to mice inhibited in the DMS indirect pathway (gray), shown separately for diagonal (left) and off-diagonal (right) probabilities. (b) Same as a but for mice inhibited in the direct pathway of the DMS. (c) The posterior probability of each state over the first and last 50 trials of a session, averaged across all sessions for mice inhibited in the indirect pathway of the DMS ( $n=271$ ). Dark lines denote average for real data (same as Fig. 7E) and faded lines indicate averages for each of the five simulations. (d) Same as c but for mice inhibited in the direct pathway of the DMS (dark lines are the same as shown in Fig. 7F). (e) Dwell times showing the average consecutive number of trials that mice inhibited in the DMS indirect pathway spent in each state for real data (left; range 39–86 trials, average session length 202 trials, same as shown in Fig. 7g) and one simulated dataset (right; range 60–71 trials, average session length 202 trials). Black dots show averages for individual mice ( $n=13$ ). We removed the last run in each session (including any run that lasted the entire session length) from the analysis, as the termination of the session prematurely truncated the length of those runs. (f) Same as e but without removing the last run in each session for real data (left; range 51–118 trials, average session length 202 trials) and one simulated dataset (right; range 65–93 trials, average session length 202 trials). (g) Same as e but for mice inhibited in the direct pathway of the DMS for real data (left; range 52–59 trials, average session length 185 trials, same as shown in Fig. 7g) and one simulated dataset (right; range 61–66 trials, average session length 185 trials). Black dots show averages for individual mice ( $n=13$ ). (h) Same as g but without removing the last run in each session for real data (left; 67–89 trials, average session length 185 trials) and one simulated dataset (right; range 74–110 trials, average session length 185 trials).



Extended Data Fig. 10 | See next page for caption.

**Extended Data Fig. 10 | Comparison of motor performance across GLM-HMM states with and without pathway-specific DMS inhibition.** (a) Schematic denoting analysis of motor performance across GLM-HMM states on laser off trials only (panels b–g) in mice unilaterally coupled to a fiberoptic for indirect pathway inhibition. (b) Average y-velocity (cm/s) during laser off trials as a function of y-position in the maze (0–300 cm in 25-cm bins) in indirect pathway mice across GLM-HMM states (state 1, blue:  $n=13,394$  trials; state 2, yellow:  $n=13,570$  trials; state 3, red:  $n=16,982$  trials). (c) As in b but for average x-position (cm) on ipsilateral or contralateral choice trials ( $n$  as in b). (d) As in c but for average view angle (degrees) on ipsilateral and contralateral choice trials ( $n$  as in b). (e) Mean per-trial standard deviation in view angle during laser off trials across GLM-HMM states (state 1, blue:  $n=13,854$  trials; state 2, yellow:  $n=14,201$  trials; state 3, red:  $n=18,258$  trials). p-value denotes one-way repeated measures ANOVA of state on view angle deviation ( $p=0.06$ ,  $F_{2,24}=3.2$ ). (f) As in e but for average distance traveled (cm) per trial. p-value denotes one-way repeated measures ANOVA of state on distance ( $p=0.02$ ,  $F_{2,24}=5.0$ ,  $n$  as in e). (g) As in e but for average percent of trials with excess travel. p-value denotes one-way repeated measures ANOVA of state on excess travel ( $p=0.0004$ ,  $F_{2,24}=10.9$ ,  $n$  as in e). (h) Schematic denoting analysis of effects of indirect pathway DMS inhibition on motor performance across GLM-HMM states in i–n. (i) As in b but for average y-velocity on laser off (black) or laser on (green) trials across GLM-HMM states ( $n$  of laser off trials as in b–g,  $n$  of laser on trials: state 1, blue:  $n=2,302$  trials; state 2, yellow:  $n=1,858$  trials; state 3, red:  $n=3,005$  trials). (j) As in c but for delta (on-off) x-position (cm) during the cue region (0–200 cm) across GLM-HMM states in mice with indirect pathway inhibition. p-value denotes one-way repeated measures ANOVA of state on delta x-position ( $p=3.2\times10^{-4}$ ,  $F_{2,24}=11.4$ ,  $n$  as in i). Post-hoc comparisons reflect two-tailed, paired Willcoxon signed rank tests between states (state 1 vs state 3:  $p=0.07$ ,  $z=1.7$ ; state 1 vs state 2:  $p=0.006$ ,  $z=2.7$ ; state 2 vs state 3:  $p=0.03$ ,  $z=2.4$ ). (k) As in j but for delta (on-off) view angle (degrees). p-value denotes one-way repeated measures ANOVA of state on delta view angle ( $p=1.2\times10^{-5}$ ,  $F_{2,26}=18.7$ ,  $n$  as in i). Post-hoc comparisons reflect two-tailed, paired Willcoxon signed rank tests between states (state 1 vs state 3:  $p=0.009$ ,  $z=2.6$ ; state 1 vs state 2:  $p=0.001$ ,  $z=-3.18$ ; state 2 vs state 3:  $p=0.002$ ,  $z=-3.1$ ). (l) Same as e but for delta (on-off) mean per-trial view angle standard deviation across GLM-HMM states in mice with indirect pathway inhibition ( $n$  of laser off trials as in e–g,  $n$  of laser on trials: state 1, blue:  $n=2,887$  trials; state 2, yellow:  $n=2,713$  trials; state 3, red:  $n=2,970$  trials). p-value denotes one-way repeated measures ANOVA of state on delta view angle deviation ( $p=0.97$ ,  $F_{2,24}=0.03$ ,  $n$  as in l). (m) Same as f but for delta (on-off) in mean per-trial distance (cm) traveled across GLM-HMM states with indirect pathway inhibition ( $p=0.68$ ,  $F_{2,24}=0.38$ ,  $n$  as in l). (n) Same as g but for delta (on-off) in percent of trials with excess travel across GLM-HMM states with direct pathway inhibition ( $p=0.08$ ,  $F_{2,24}=2.8$ ,  $n$  as in l). (o) As in a but schematic denoting analysis of motor performance across GLM-HMM states on laser off trials only in mice unilaterally coupled to a fiberoptic for direct pathway inhibition in p–u. (p) As in b but for y-velocity (cm/s) on laser off trials across GLM-HMM states in direct pathway mice (state 1, blue:  $n=12,294$  laser off and  $n=2,302$  laser on trials; state 2, yellow:  $n=9,201$  laser off and  $n=1,858$  laser on trials; state 3, red:  $n=16,239$  laser off and  $n=3,005$  laser on trials). (q) As in c but x-position (cm) for direct pathway mice ( $n$  as in p). (r) As in d but for view angle (degrees) for direct pathway mice ( $n$  as in p). (s) As in e but for mean per-trial view angle standard deviation across GLM-HMM states in direct pathway mice (state 1, blue:  $n=13,403$  laser off and  $n=2,508$  laser on trials; state 2, yellow:  $n=9,555$  laser off and  $n=1,969$  laser on trials; state 3, red:  $n=18,292$  laser off and  $n=3,450$  laser on trials). p-value denotes one-way repeated measures ANOVA of state on per-trial view angle standard deviation ( $p=0.12$ ,  $F_{2,24}=2.3$ ). (t) As in f but for distance (cm) in direct pathway mice ( $p=0.1$ ,  $F_{2,24}=2.5$ ). (u) As in g but for percent trials with excess travel in direct pathway mice ( $p=0.14$ ,  $F_{2,24}=2.1$ ). (v) As in h but schematic denoting analysis of effects of direct pathway DMS inhibition on motor performance across GLM-HMM states in w–bb. (w) As in i but for the mean y-velocity (cm/s) on laser on (green) and off (black) trials across GLM-HMM states in direct pathway mice. (x) As in j but for the delta (on-off) x-position (cm) across GLM-HMM states in direct pathway mice. p-value denotes one-way repeated measures ANOVA of state on delta x-position ( $p=7.9\times10^{-5}$ ,  $F_{2,24}=14.9$ ). Posthoc comparisons reflect two-tailed, paired Willcoxon signed rank tests between states (state 1 vs state 3:  $p=0.06$ ,  $z=1.8$ ; state 1 vs state 2:  $p=0.005$ ,  $z=2.8$ ; state 2 vs state 3:  $p=0.005$ ,  $z=2.8$ ). (y) As in k but for delta (on-off) view angle (degrees) across GLM-HMM states in direct pathway mice. p-value denotes one-way repeated measures ANOVA of state on delta view angle ( $p=2.6\times10^{-4}$ ,  $F_{2,24}=12.3$ ). Posthoc comparisons reflect two-tailed, paired Willcoxon signed rank tests between states (state 1 vs state 3:  $p=0.03$ ,  $z=2.3$ ; state 1 vs state 2:  $p=0.003$ ,  $z=2.98$ ; state 2 vs state 3:  $p=0.03$ ,  $z=2.1$ ). (z) As in l but for delta (on-off) mean per-trial view angle standard deviation (degrees) in direct pathway mice ( $p=0.40$ ,  $F_{2,24}=0.94$ ,  $n$  as in s–u). (aa) as in m but for delta (on-off) in mean distance (cm) traveled in direct pathway mice ( $p=0.43$ ,  $F_{2,24}=0.89$ ). (bb) as in n but for delta (on-off) in percent trials with excess travel in direct pathway mice ( $p=0.90$ ,  $F_{2,24}=0.1$ ). Throughout solid colored bars denote mean  $\pm 1$  s.e.m. while transparent grey lines reflect individual mouse mean.