Strong and opponent contributions of dorsomedial striatal pathways to behavior depends on cognitive demands and task strategy

Scott S. Bolkan^{1,#}, Iris R. Stone^{1,#}, Lucas Pinto¹, Zoe C. Ashwood¹, Jorge M. Iruavedra Garcia¹, Alison L. Herman¹, Priyanka Singh¹, Akhil Bandi¹, Julia Cox¹, Christopher A. Zimmerman¹, Jounhong Ryan Cho¹, Ben Engelhard¹, Sue A. Koay¹, Jonathan W. Pillow^{1,2}, Ilana B. Witten^{1,2}

¹ Princeton Neuroscience Institute; Princeton University; Princeton, NJ 08544; USA

² Department of Psychology; Princeton University; Princeton, NJ 08544; USA

[#] These authors contributed equally to this work

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, has remained unresolved. Here we find that strong, opponent control of behavior by the two pathways of the dorsomedial striatum (DMS) depends on a task's cognitive demands. 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 cognitive demands imposed by a task, as well as the strategy that mice pursue within a task, determine whether DMS pathways provide strong and opponent control of behavior.

Introduction

The striatum is composed of two principal outputs, the direct and indirect pathways, which are thought to exert opposing effects on behavior¹⁻³. In support of this view, a number of influential studies have shown that pathway-specific activation of the striatum produces opposing behavioral biases⁴⁻¹³. For example, activation of the direct pathway increases spontaneous movements⁴⁻⁷, while indirect pathway activation decreases spontaneous movements^{4,5,7}. Similarly, direct or indirect pathway activation oppositely influences whether an animal will spontaneously rotate to the left or right^{4,5}, repeat or cease a stimulation-paired behavior⁸⁻¹⁰, or orient to the left or right to report a value-based decision¹¹.

Despite this pioneering work, it remains unresolved whether the endogenous activity in 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 patterns⁴⁻¹³. A number of these studies have moreover challenged the classic view, reporting either similar or non-opposing behavioral effects of activating each pathway¹⁴⁻²⁰. The imbalance towards reports of activation suggests a wealth of negative results from inhibition, raising questions about the function of the endogenous activity in the two pathways, and whether the endogenous activity 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 towards a movement versus the generation of the movement itself⁴⁻⁸. In contrast, studies of the DMS that were not pathway-specific have instead focused on cognitively demanding behaviors²¹⁻³². Taken together, this leaves open the possibility that the two pathways exert opposing control of movement in the context of decision-making, rather than directly controlling a motor output irrespective of cognition.

Thus, to determine if the contribution of endogenous activity in the two 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 cognitive demands determined the effect of DMS inhibition on behavior. Second, we used a latent state model to identify time-varying strategies within the same task. This allowed us to determine if the contribution of each pathway to behavior depended on the strategy being pursued, 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. However, pathway-specific inhibition produced strong and opposing biases on decisions based on the gradual accumulation of pulsatile sensory evidence into memory in a virtual T-maze³³⁻³⁶. In contrast, we observed significantly smaller effects of pathway-specific inhibition on choice during less cognitively demanding task variants with similar sensory features and identical motor requirements. Our latent state model further revealed that even within the evidence accumulation task, mice pursue different strategies across time that differ in the weighting of sensory evidence and trial history, as well as the extent that DMS pathway inhibition impacts choice. Thus, by comparing the effects of DMS pathway-specific inhibition across behavioral tasks, and across time within a task, we conclude that both task demands and cognitive strategy determine whether or not DMS pathways exert strong and opposing control over behavior.

Results

Pathway specific inhibition of DMS is effective, generating little post-inhibitory rebound or activation during the inhibition period.

We first sought to validate the effectiveness of halorhodopsin¹⁹ (NpHR)-mediated inhibition of indirect and direct striatal pathways in awake, head-fixed mice (**Figure S1A-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) that we verified to have high degrees of specificity and pentrance for the indirect or direct pathways (**Figure S2**). We confirmed that 532-nm (5mW) laser delivery to the DMS through a tapered optical fiber produced rapid, sustained, and reversible inhibition of spiking in mice expressing NpHR in either the indirect pathway (**Figure 1B** and **S1C-E**, $n = 18/60$, 30% of recorded DMS neurons significantly inhibited) or the direct pathway (**Figure 1C** and **S1F-H**, $n = 21/50$, 42% of recorded neurons significantly inhibited). Moreover, we observed (1) minimal excitation during laser delivery, consistent with recent observations^{37,38} (**Figure S1D,G, left**), (2) minimal effects on spiking upon laser offset (**Figure S1D,G, right**), indicating limited post-inhibitory rebound^{39,40}, and (3) stability in the efficacy of inhibition across each recording (**Figure S3**). All together, our findings indicate that NpHR-mediated inhibition of indirect and direct pathway DMS neurons is effective.

Pathway specific inhibition of DMS does not produce detectable changes in motor output during navigation of a virtual corridor.

To determine if the 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 2-dimensional linear corridor in virtuality reality (VR) (**Figure 1D-F**, 6-cm x 330-cm corridor). Rotation of the ball in the anterior-posterior (and medial-lateral) axes of the mouse served to control movements in the y- (and x-) directions in VR (see **Methods** for details). Mice received reward upon reaching the end of the corridor, followed by teleportation back to the start region; unilateral, pathway-specific optogenetic inhibition of the DMS (or DMS illumination alone) was restricted to 0-200 cm (laser on 30% of trials; hemisphere of illumination alternated across days; **Figure 1F**). The parameters of the virtual corridor and inhibition period were selected to closely match the central stem of the VR-based T-maze decision-making tasks that are the focus of the rest of this paper.

We found no detectable impact of indirect or direct pathway inhibition, nor DMS illumination alone, on indicators of motor output during virtual corridor navigation (**Figure 1G-J**). This included measures of velocity (**Figure 1G**), position (**Figure 1H**) or view angle (**Figure 1I**) relative to the laser hemisphere, and distance travelled (**Figure 1J**) (see **Figure S4** 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 assay (**Figure S5**).

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.

Figure 1. Inhibition of DMS pathways has no detectable impact on behavior in mice navigating a virtual corridor. (A) Schematic of viral delivery of Cre-dependent halorhodopsin (NpHR) to the dorsomedial striatum (DMS) of A2a-Cre, D2R-Cre, or D1R-Cre mice (left). Schematic of recording optrode, consisting of tapered optical fiber coupled to 32-channel silicon probe (right). The fiberoptic was attached to a 532-nm laser and light (5-mW) was delivered to the DMS of awake, ambulating mice. (B) Example peristimulus time histograms (PSTH) (top) and raster plot of trial-by-trial spike times (bottom) from a single neuron recorded in the DMS of an awake, ambulating A2a-Cre mouse expressing Cre-dependent NpHR (indirect pathway). Inset at top displays average spike waveform (black) and 100 randomly sampled spike waveforms (grey). A single trial consisted of 5-s without laser (pre, -5 to 0-s), 5-s of 532-nm laser delivery (on, 0 to 5-s), followed by a 10-s ITI (40 total trials). (C) As in B but for a single neuron recorded in the DMS of a D1R-Cre mouse expressing Cre-dependent NpHR (direct pathway). (D) Schematic of bilateral fiberoptic implantation of the DMS and unilateral laser delivery to behaving mice, with example histology from a mouse expressing NpHR in the indirect (left, D2R-/A2a-Cre) or direct (middle, D1R-Cre) pathways, or control mice without opsin (right, no opsin, A2a-/D2R- or D1R-Cre). 532-nm light (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 virtual reality (VR) apparatus allowing 2-D navigation. Displacements of an air-suspended spherical ball in the anterior-posterior (and medial-lateral) axes of the mouse controlled y- (and x-) position movements in a VR environment projected onto a 270° toroidal screen. (F) Schematic of a virtual corridor arena measuring 6-cm in width and 330-cm in total length, consisting of a start region (-10-0cm), a laser delivery region (0-200cm) in which mice received unilateral 532-nm light on a random subset of trials (30%), a reward location (310cm) where mice received reward, and a teleportation location (320cm) where mice

were transported to the start region following a variable ITI with mean of 2-s. **(G)** Average y-velocity (cm/s) across mice as a function of y-position (0-300cm in 25-cm bins) while navigating the virtual corridor on laser off (black) or laser on (green) trials in groups receiving DMS indirect pathway inhibition (*left*, $n = 7$ mice, $n = 1,712$ laser off and $n = 1,288$ laser on trials) or direct pathway inhibition (*middle*, $n = 6$ mice, $n = 1,088$ laser off and $n = 757$ laser on trials), or laser delivery to the DMS in the absence of NpHR (*right*, 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 travelled (cm) to traverse the virtual corridor during laser off (black) or laser on (green) trials for mice receiving DMS indirect pathway inhibition ($n = 7$ mice, $n = 2,109$ laser off and $n = 1,574$ laser on trials), direct pathway inhibition ($n = 6$ mice, $n = 1,330$ laser off and $n = 930$ laser on trials), 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 mean and s.e.m. across mice throughout, and grey lines indicate individual mouse mean.

A set of virtual reality T-mazes have similar sensory features and identical motor requirements but different 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^{33-35,41} that shared identical motor readouts (left or right choice), had highly similar sensory environments, yet differed in their cognitive requirements (**Figure 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 period”, 0-200cm), and mice were rewarded for turning to the maze side with the greater number of cues (**Figure 3A,B**; black, *left*). Thus, mice were required to continually accumulate sensory cues over several seconds into a memory (or motor plan) that guided their left/right decision.

Figure 2. A set of virtual reality T-mazes have similar sensory features and identical motor requirements but different cognitive demands. **(A)** Schematic of three virtual reality (VR)-based T-maze tasks. **(B)** Example VR mouse perspective at the same maze position (-10cm, 120cm, 195cm, and 295cm) from the example trial depicted in **A** of the evidence accumulation (*left*, black), no distractors (*middle*, ctrl #1), or permanent cues (*right*, ctrl #2) tasks. **(C)** Average choice accuracy (% correct) across mice performing the accumulation of evidence (black, $n = 32$ mice, $n = 52,381$ trials), no distractors (magenta, ctrl #1: $n = 34$ mice, $n = 52,381$ trials), or permanent cues (blue, ctrl #2: $n = 34$ mice, $n = 52,381$ trials). task $p = 6.4 \times 10^{-21}$. **(D)** y position (m) vs y-velocity (cm) for the same tasks. **(E)** y position (m) vs x-position (cm) for left and right choices. **(F)** y position (cm) vs view angle (deg) for left and right choices. **(G)** distance (cm) for the tasks. task $p = 0.19$.

= 56,953 trials), or permanent cues (cyan, ctrl #2: $n = 20$ mice, $n = 27,870$ trials) tasks. Solid bars denote mean and s.e.m and transparent 'x' indicate individual mouse mean. p-value denotes one-way ANOVA of task on accuracy ($p = 6.39 \times 10^{-21}$, $F_{2,83} = 87.2$). Asterisks indicate statistical significance of post-hoc unpaired, two-tailed t- (or ranksum) tests for normally (or non-normally) distributed data (top to bottom: $***p = 1.3 \times 10^{-9}$, $t_{50} = -7.4$; $***p = 7.1 \times 10^{-11}$, $z_{62} = -6.5$; $***p = 7.5 \times 10^{-5}$, $z_{50} = 3.9$). (D) Average y-velocity (cm/s) across mice as a function of y-position (0-300 cm in 25cm bins) during performance of each task (colors and n as in C). (E) Same as D but for average x-position (cm) on left and right choice trials. (F) Same as D but for average view angle (degrees) on left and right choice trials. (G) Average distance (cm) travelled to complete a trial across mice (evidence accumulation, $n = 32$ mice, $n = 64,420$ trials; no distractors (control #1): $n = 34$ mice, $n = 61,308$ trials; permanent cues (control #2): $n = 20$ mice, $n = 30,150$ trials). p-value of one-way ANOVA of task on distance ($p = 0.19$, $F_{2,83} = 1.7$).

In two additional control tasks, we made modifications that served 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 non-rewarded arm side were omitted (**Figure 3A-B**; magenta, *middle*). 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 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 that in the evidence-accumulation task, but rather than *transient* visual cue presentation, visual cues were *permanently* visible from trial onset (**Figure 3A-B**; cyan, *right*). 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 (**Figure 2C**, one-way ANOVA of task on accuracy, $p = 6.4 \times 10^{-21}$, $F_{2,83} = 87.2$; post-hoc, unpaired, two-tailed Wilcoxon ranksum test of evidence accumulation vs. no distractors, $p = 7.1 \times 10^{-11}$, $z_{62} = -6.5$; evidence accumulation vs permanent cues, $p = 4.0 \times 10^{-7}$, $z_{50} = -5.0$).

While the motor requirements of a decision were the same across tasks (crossing an x-position threshold at the end of the central stem, see **Methods**), we examined the possibility that the difference in cognitive requirements across tasks altered movement within the central stem of the maze (0-300cm). However, we observed no consistent cross-task differences in the average velocity of mice in the maze stem (**Figure 2D**), the average x-position (**Figure 2E**) or view angle (**Figure 2F**) of mice as they traversed the maze stem on left or right choice trials, nor the total distance travelled to complete a trial (**Figure 2G**; see **Figure S6** for additional measures). We further compared the trial-by-trial relationship between behavior in the central stem of the maze and choice across the three tasks by using a decoder to predict choice based on the trial-by-trial x-position (**Figure S6F**) or view angle (**Figure S6G**) at successive maze positions (0-300cm in 25-cm bins). While we were able to predict choice from either measure above chance levels in all three mazes (consistent with previous studies³³), choice prediction accuracy was statistically indistinguishable across tasks. Together, this indicated that the cross-task differences in cognitive demands did not prompt mice to systematically adopt distinct motor strategies in the performance of each task.

Figure 3. Inhibition of DMS but not NAc pathways has strong and opposing influences on choice during an evidence accumulation task, while having less effect on task variants with weaker cognitive demands.

(A) Schematic of bilateral viral delivery of Cre-dependent NpHR to the dorsomedial striatum (DMS). (B) Schematic of bilateral fiberoptic implantation of the DMS and unilateral laser delivery to behaving mice, with example histology from a mouse expressing NpHR in the indirect (*left*, D2R-/A2a-Cre) or direct (*middle*, D1R-Cre) pathways, or DMS illumination in the absence of NpHR (*right*, no opsin, A2a-/D2R- or D1R-Cre). 532-nm light (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 (0-200cm) on a random subset of trials (10-20%). (D) Average choice bias (% contralateral-ipsilateral) across mice on laser off (black) and laser on (green) trials during the evidence accumulation task in mice receiving unilateral indirect pathway inhibition (*left*, $n = 11$ mice, $n = 16,935$ laser off and $n = 3,390$ laser on trials), unilateral direct pathway inhibition (*middle*: $n = 10$ mice; $n = 14,030$ laser off and $n = 3,103$ laser on trials), or unilateral illumination of the DMS in the absence of NpHR (*right*, $n = 11$ mice, $n = 21,422$ laser off and $n = 5,113$ laser on trials). (E) Difference in contralateral choice bias (% contralateral-ipsilateral) between laser off and on trials (% 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 significance of unpaired Wilcoxon ranksum comparison of indirect to no opsin: $***p = 1.1 \times 10^{-4}$, $z_{20} = 3.9$; direct to no opsin: $***p = 2.2 \times 10^{-4}$, $z_{19} = -3.7$. (F-H) Same as C-E but for the no distractors (control #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 significance of unpaired Wilcoxon ranksum comparison of indirect to no opsin: not significant (n.s.), $p = 0.22$, $z_9 = 1.3$. Direct to no opsin: not significant (n.s.), $p = 0.08$, $z_{11} = -1.8$. (I-K) As in C-E but for the permanent cues (control #2) task. Indirect: $n = 6$ mice, $n = 3,964$ laser off and $n = 916$ 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 significance of unpaired Wilcoxon ranksum comparison of indirect to no opsin: not significant (n.s.), $p = 0.05$, $z_{10} = 2.0$. Direct to no opsin: not significant (n.s.), $p = 0.62$, $z_{11} = 0.5$. (L) As in A but for bilateral viral delivery of Cre-dependent NpHR to the nucleus accumbens (NAc). (M) Same as B but for bilateral fiberoptic implantation of the NAc and unilateral laser delivery to behaving mice, with example histology from a mouse expressing NpHR in the indirect (*left*, D2R-/A2a-Cre) or direct (*middle*, D1R-Cre) pathways, or NAc illumination in the absence of NpHR (*right*, no opsin, A2a-/D2R- 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: $n = 6$ mice, $n = 9,819$ laser off and $n = 1,488$ laser on trials. Asterisks indicate significance of unpaired, two-tailed Wilcoxon ranksum comparison of indirect to no opsin: not significant (n.s.), $p = 0.86$, $z_{13} = 0.18$; direct to no opsin: not significant (n.s.), $p = 0.04$, $z_{14} = 2.0$. To account for multiple group comparisons we considered p-values significant after Bonferroni correction ($\alpha < 0.05/2$ comparisons).

Pathway-specific inhibition in the DMS produces large and opposing choice biases in an evidence accumulation task, while having diminished effects in two control tasks with reduced cognitive demands

We next sought to determine the contribution of endogenous activity in DMS pathways by performing unilateral, pathway-specific inhibition of DMS indirect and direct pathways (or DMS illumination alone) restricted to the cue region (0-200-cm) of each task (laser on 10-20% of trials; hemisphere of illumination alternated across days; **Figure 3A-B**). We found that inhibition of the indirect pathway produced a large bias towards contralateral choices during the accumulation of evidence task (**Figure 3C** and **3D, left**), which was consistent across individual mice, and significantly greater than that observed in control animals (**Figure 3E**, average contralateral bias in indirect pathway: $42.3 \pm 4.4\%$; 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 (**Figure 3D, middle**; average contralateral bias: $-36.8 \pm 8.6\%$), which was consistent across individual mice and significantly greater than that observed in control animals (**Figure 3E**). However, in this case the direction of the choice bias was in the opposite (ipsilateral) direction to that observed with indirect pathway inhibition.

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 both the “no distractors” and “permanent cues” control tasks (**Figure 3F-K**; indirect pathway: evidence accumulation vs no distractors, $p = 8.1 \times 10^{-4}$, $z_{16} = 3.4$; evidence accumulation vs permanent cues, $p = 0.002$, $z_{15} = 3.1$; direct pathway: evidence accumulation vs no distractors, $p = 0.002$, $z_{17} = -3.1$; evidence accumulation vs permanent cues, $p = 0.005$, $z_{15} = -2.8$; unpaired, two-tailed Wilcoxon ranksum test). 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 (**Figure 3H**, for “no distractors”; **Figure 3K**, for “permanent cues”).

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 indirect and direct pathway inhibition had minimal impact on movement in a virtual corridor (**Figure 1** and **S4**), we considered the possibility that pathway-specific DMS inhibition altered motor performance in the T-maze tasks in a manner that depended on task demands. We found no cross-task differences in the effects of pathway-specific inhibition on measures of velocity (**Figure S7A-C**), distance traveled (**Figure S7D-F**), or per-trial standard deviation in view angle (**Figure S7G-I**). However, similar to cross-task effects on choice, and consistent with the tight relationship between x-position and view angle with choice across tasks (**Figure S6F-G**), we found subtle but opposing effects of pathway-specific inhibition on average x-position (**Figure S7J-L**) and view angle (**Figure S7M-O**) in the evidence accumulation task, while such effects tended to be smaller in the control tasks. As the quantitative relationship between x-position or view angle and choice is indistinguishable across tasks in the absence of neural inhibition (**Figure 2E-F**, **S6C-D** and **S6F-G**), cross-task differences in motor strategy does not provide a trivial explanation for these effects. Rather, taken together with the absence of an effect of pathway-specific DMS inhibition on view angle or x-position in the virtual corridor (**Figure 1** and **S4**), these data instead imply that the effects of inhibition on behavior depends on the cognitive demand of a task.

Pathway-specific inhibition of the NAc does not produce large and opposing choice biases

We next sought to determine whether opponent control of choice during the evidence accumulation task was specific to DMS pathways, or if it extended to the ventral striatum. Towards 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 (0-200cm) during the evidence accumulation task (**Figure 3L-P**).

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 (indirect pathway DMS vs NAc: $p = 2.7 \times 10^{-4}$, $z_{18} = 3.6$; direct pathway DMS vs NAc: $p = 1.8 \times 10^{-4}$, $z_{18} = -3.7$; unpaired, two-tailed Wilcoxon ranksum test), and were not significantly different from control animals (**Figure 3O-P**). It is unlikely that this dissociation can be explained by greater co-expression of pathway-specific markers in ventral versus dorsal striatum⁴², as both subregions exhibited equally low co-localization of D1R and D2R receptors (**Figure S2J-L**).

Bernoulli GLM demonstrates that sensory evidence, trial history, and DMS pathway inhibition contribute to choice during evidence accumulation, but cannot 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 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 the possibility that the contributions of the two pathways to behavior could change over time, even within the same task.

Figure 4. A GLM reveals that sensory evidence, DMS 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 (Δ 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 (e.g. 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 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 (+/-) posterior standard deviation 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 on (green) trials, for mice receiving indirect pathway DMS inhibition for the data (left) and for

simulations of the model (right). Error bars denote 95% confidence intervals; 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).

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 (Figure 4A-B). These covariates included the sensory evidence (difference between the number of right and left cues, or “ Δ cues”), the recent choice and reward history, the presence of the 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 (see **Methods** for details).

We fit the GLM to aggregated behavioral data from mice inhibited in each DMS pathway and found that sensory evidence, trial history, and laser all contributed to predicting choice. As expected (**Figure 4C-D**), the effect of laser delivery in the indirect and direct pathways was large and opposite in sign. However, the GLM did not accurately capture the animal's psychometric curve, describing the probability of a rightward choice as a function of the sensory evidence (**Figure 4E-F**). This led us to consider variants of the standard GLM that might better account for choice behavior.

GLM-HMM better explains the choice data than the standard GLM, particularly on DMS inhibition trials

The standard GLM describes choice as depending on a fixed linear combination of sensory evidence, trial history, and laser delivery. However, an alternative possibility is that mice use a weighting function that changes over time. To test this idea, we adopted a latent state model that allowed different GLM weights in different states, using the same external covariates as the standard 1-state GLM (**Figure 4**). The model consists of a Hidden Markov Model (HMM) with Bernoulli Generalized Linear Model (GLM) observations, or GLM-HMM⁴³⁻⁴⁶ (**Figure 5 A-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 (see **Methods** for details).

The GLM-HMM explained the choice data in the evidence accumulation task better than the GLM across multiple measures. We compared the likelihood of each animal's data under the GLM-HMM to the standard Bernoulli GLM using cross-validation with held-out sessions (3-state GLM-HMM in **Figure 5**; also see **Figure S8A-D** for more information on model selection and demonstration that ~3-4 latent states was sufficient to reach a plateau in likelihood). The 3-state GLM-HMM achieved an average of 6.2 bits/session increase in log-likelihood, making an average session ~76 times more likely under the GLM-HMM (**Figure 5C**). Furthermore, the GLM-HMM correctly predicted choice on held-out data more often than the GLM, especially on laser trials (**Figure 5D**; average improvement across mice of 1.6% on all trials, 3.5% on laser trials, and 4.1% on laser trials when considering mice with at least 100 laser trials).

Most interestingly, the GLM-HMM was better able to capture the temporal structure in the effect of laser on choice. Specifically, the choice data contained long runs in which the choice was consistent with the bias direction predicted by the laser, a feature which GLM-HMM simulations recapitulated, but GLM simulations did not (**Figure 5E-F**). Thus, taken together, the GLM-HMM provided a better model of the choice data than a standard GLM, particularly on laser trials.

GLM-HMM identifies multiple task strategies during the evidence accumulation task, differing in their weighting of sensory evidence, choice history, and DMS pathway inhibition

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, laser delivery to DMS pathways (**Figure 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.

Figure 5. A GLM-HMM better explains choice during the evidence accumulation task compared to the GLM, particularly on laser trials. (A) Schematic of GLM-HMM. The model has 3 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 Δ cues, to choice probability. (B) Example sequence of 3 trials, 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 (Δ 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. (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 (left) or on the subset of trials when the laser was on (right). (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 (i.e. a choice contralateral for DMS indirect pathway inhibition). Data (black), GLM simulation (blue), GLM-HMM simulation (pink). For the simulations, data of the same length as the real data was generated 100 times and the resulting histograms averaged. Curves denote smoothed counts using a sliding window average (window size = 3 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 are defined as those ipsilateral to the hemisphere of inhibition.

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 (**Figure 6C-D**). We then analyzed the psychometric curves for trials assigned to each state. In state 3, performance was low (**Figure 6G**) and DMS inhibition had little effect on behavior (**Figure 6C-D**). This is consistent with the high GLM weight on choice history in this state, and low weights on sensory evidence and laser (**Figure 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 (**Figure 2C**) but effects of inhibition are limited (**Figure 3F-K**), 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 (**Figure 6C-D, G**). Interestingly, the effect of laser stimulation 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, low weighting of choice history, but greatly differed in their weighting of the laser (**Figure 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 (**Figure 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 (**Figure 6E-F**). This not only validated our fitting procedure, but also provided additional evidence that a multi-state 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 laser trials (**Figure 6I**), implying that laser delivery itself did not generate transitions between states. In addition, the fraction of trials with laser was equivalent across states (~15% of all trials in each state; **Figure 6H**). This implies that the model did not identify states simply based on the presence of laser trials.

Importantly, we obtained similar states when fitting the model to a combined dataset including both mice receiving DMS indirect and direct pathway inhibition, as well as control mice receiving DMS illumination in the absence of NpHR (**Figure S8E**). As when fitting each cohort separately, the combined model 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. This indicated that the discovery of a state in each inhibition group with a large laser weight was a consequence of the inhibition per se (as opposed to the laser itself, or the analysis).

Figure 6. The GLM-HMM discovers states with different weighting on sensory evidence, DMS pathway inhibition, and choice history. (A) Fitted GLM weights for 3-state model from mice in the indirect pathway DMS inhibition group. Error bars denote (+/-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.g. 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 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). **(B)** Same as **A** but 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, trials were assigned to each state by taking the maximum of the model's posterior state probabilities on each trial. Error bars denote ± 1 SEM 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. **(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 (see **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 pathway (left) and direct pathway (right), 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)** Percentage of laser-on trials that the model assigned to each state for mice receiving DMS inhibition in the indirect pathway (left) and direct pathway (right). 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 ($n=8570$, indirect; $n=7927$, direct).

We also examined the results of fitting the 4-state GLM-HMM (**Figure S8C-D**), given it had a slightly higher cross-validated log-likelihood than the 3-state model (**Figure S8A**). In this case, the weights for states 1 and 2 were very similar to the 3-state model; the key difference was that the choice history state (state 3 from the 3-state model), was further subdivided into two states that differed in having a slight right versus a slight left bias.

Diversity across sessions in the timing and number of GLM-HMM state transitions

The fitted transition matrix revealed a high probability of remaining in the same state across trials (**Figure 7A-B**). These transition probabilities produced a diversity in the timing and number of state transitions across sessions, which we visualized by calculating the posterior probability of being in each state on each trial (**Figure 7C-D**). In some sessions, mice persisted in the same state (with the state on a trial defined as the state with maximum posterior probability), while in many sessions, mice visited two or even all three states (example sessions in **Figure 7C-D**; summaries of state occupancies across sessions in **Figure 7H-J**; summaries of individual mice in **Figure S10**). Average single-state dwell times ranged from 39-86 trials (**Figure 7G**). This was far 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 (with the exception of state 3 for the indirect pathway, which increased in probability towards the end of the session, potentially reflecting a decrease in task engagement related to reward satiety; **Figure 7E-F**). Model simulations recapitulated these state transition characteristics, including dwell times and state occupancies (**Figure S9**).

Given the presence of sessions in which the 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 (**Figure 7K**, fraction of trials in each state for each session for this model). 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 the inhibition by disallowing transitions in and out of that state (**Figure 7L**). Using cross-validation, we found that neither alternative model explained the data as well as a model with unrestricted transitions (**Figure 7M**), indicating that within-session transitions between states was an important feature of the model.

Figure 7. Diversity across sessions in the timing and the number of state transitions. (A) Transition probabilities for the indirect pathway group. (B), Same as (A) but for the direct pathway group. (C) The model's posterior probability of being in each state for each trial for 3 example sessions from a mouse in the indirect pathway group. (D) Same as (C) but for two mice from the direct pathway group. (E) 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). (F) Same as (E) but for mice receiving DMS direct pathway inhibition (n=266). (G) Dwell times showing the average consecutive number of trials that the mice spent in each state for mice with indirect (left; range 39-86 trials, average session length 202 trials) and direct (right; range 52-59 trials, average session length 185 trials) pathway inhibition. Black dots show averages for individual mice (n=13 for both groups). (H). 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 (e.g. blue indicates 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. Grey arrows identify the example sessions shown in C and D. (I) 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 for both groups). (J) Time spent in each state represented as a percentage of total trials for mice inhibited in the indirect pathway (left) and direct pathway (right). Colored bars denote the average state occupancies across all mice. Black dots show averages for individual mice (n=13 for both groups). (K) Same as (H) except state assignments were obtained from a model in which the transition probabilities were restricted to disallow transitions between states (i.e. all off-diagonal transition probabilities equal zero; see Methods). (L) Same as (K) except state assignments were obtained from a model in which transitions were disallowed between state 2 and the other states. (M) Comparison of the cross-validated log-likelihood of the data when fitting GLM-HMMs with the reduced models from K and L, relative to the log-likelihood of the full model, in bits per session.

Motor performance across GLM-HMM states

To provide additional insight into the behavior that characterizes GLM-HMM states, we considered the possibility that the motor performance of mice may differ across states. We found that on trials without laser (**Figure S11A-G** and **S11O-U**), mice exhibited no obvious differences across states in average velocity (**Figure S11B** and **S11P**), average x-position (**Figure S11C** and **S11Q**) or view angle (**Figure S11D** and **S11R**). However, we observed a tendency for increased per-trial standard deviation in view angle (**Figure S11E** and **S11S**) and distance travelled (**Figure S11FG** and **S11T-U**) during state 3 relative to state 1 and 2, which may be consistent with the interpretation of state 3 as a task-disengaged state.

We also considered the possibility that indirect and direct pathway DMS inhibition had state-dependent effects on motor output (**Figure S11H-N** and **S11V-BB**). We observed limited effects of inhibition on velocity (**Figure S11I** and **S11W**), per-trial standard deviation in view angle (**Figure S11L** and **S11Z**), and distance travelled (**Figure S11M-N** and **S11AA-BB**) across all three states. However, similar to our cross-task comparisons (**Figure S7J-O**), we observed a subtle and opposing bias in average x-position (**Figure S11J** and **S11X**) and view angle (**Figure S11K** and **S11Y**) with pathway-specific DMS inhibition, which trended towards being greatest in the state with the largest laser weight (state 2, **Figure 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, albeit noisy, indicators of choice in the context of VR-based T-maze tasks (**Figure S6F-G**).

DISCUSSION

Our findings indicate that while opposing contribution of DMS pathway inhibition to movement is minimal in the absence of a decision (**Figure 1**), the pathways provide large and opponent contributions to decision-making. Moreover, this contribution depends on the cognitive demands of the decision-making task, as the effect of inhibition is much larger in a task that requires gradual evidence accumulation relative to control tasks with weaker cognitive requirements, but similar sensory features and motor requirements (**Figure 2, 3**). The GLM-HMM 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 (**Figure 6**). Thus, our findings imply that DMS pathways provide opposing control of the cognitive process of evidence accumulation, rather than to low level motor output.

Cross-task differences in effects of DMS pathway inhibition

We provide a direct demonstration that *endogenous* activity in direct and indirect pathways of the DMS oppositely controls the decision-making process, rather than providing direct control over the generation of motor output. Previous work supporting the classic view of opposing pathway function has overwhelmingly relied on the synchronous activation, as opposed to inhibition, of striatal pathways. Moreover, some prominent studies employing activations have challenged the classic view, reporting either similar or non-opposing behavioral effects of each pathway^{14–20}, which may suggest limitations in using artificial activation in assessing pathway function. Prior work that has demonstrated opposing control of behavior by activation of the two pathways has not compared effects on motor outputs within the same behavioral framework while only varying cognitive demand, and therefore has not definitively distinguished between motor and cognitive contributions. While DMS pathway activation may be sufficient to bias behaviors such as spontaneous rotations, we observed relatively little impact of inhibition on decisions with diminished cognitive requirements (**Figure 3**) or behaviors in the absence of a decision (**Figure 1, S4, and S5**). The limited contributions of endogenous DMS activity to behavior in these contexts may explain the limited number of reports demonstrating large behavioral effects of pathway-specific inhibition to date.

Our findings also provide new context for the increasingly observed co-activation of striatal pathways during movement^{47–53,21}. Indeed, our results would not necessarily predict opposing correlates of movements in indirect and direct pathways of the DMS, but rather opposing correlates of a decision process^{23,25,54}. The much larger effect of pathway-specific inhibition we observed during the accumulation of evidence task is consistent with a role for the DMS in decision-making and the dynamic comparison of the value of competing options^{22–26,54–58}. Together, our work raises the importance of optogenetic inhibition in complex cognitive settings to probe models of striatal function.

Within-task changes in effects of DMS pathway inhibition

In addition, we reveal the novel insight that mice pursue different strategies within a single task and that the striatal contribution to choice depends on the strategy pursued. The application of a GLM-HMM was critical in uncovering this latent feature of behavior, allowing the unsupervised discovery of behavioral states that differ in how external covariates were weighted to influence a choice^{45,59,60}. This provided three insights.

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 (**Figure 6**). This strengthens our conclusion that arose from the cross-task comparison, which is that DMS pathways have a greater contribution to behavior when actively accumulating evidence towards a decision output.

Second, mice occupied two qualitatively similar task-engaged states that were distinguished most prominently by the influence of DMS inhibition on choice (**Figure 6**). While transitions

between these two states were relatively rare on the same day, there were days that included both states (**Figure 7**). The discovery of these two states leads to the intriguing suggestion that mice are capable of accumulating evidence towards a decision in at least two neurally distinct manners -- one that depends on each pathway (state 2), and another that does not (state 1).

Finally, the GLM-HMM reveals a dissociation between behavioral performance (which was lower in state 3 than state 1 or 2, **Figure 6G**) and the effect of DMS inhibition (which was higher in state 2 than state 1 or 3, **Figure 6C-D**). Taken together with our cross-task comparison (**Figure 3**), where we instead found that the control tasks with higher performance had less DMS dependence, the implication is that performance (or reward rate) alone does not predict the involvement of DMS pathways in behavior. Instead, our results suggest that DMS contributes preferentially to decisions that depend on evidence accumulation, as opposed to decisions guided by choice history (state 3) or by sensory evidence in the absence of a significant memory requirement (“no distractor” and “permanent cues” control tasks).

Thus our findings emphasize the importance of accounting for ongoing behavioral strategy when assessing neural mechanisms⁶¹. Toward 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 settings⁶²⁻⁶⁴.

Supplemental figures and legends

Figure S1. Optogenetic inhibition of indirect and direct pathway neurons in DMS 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) 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_i). Example recording electrode tracks and cre-dependent NpHR expression in an A2a-Cre mouse targeting the indirect pathway of the DMS (B_{ii}). Example recording electrode tracks and cre-dependent NpHR expression in a D1R-Cre mouse targeting the direct pathway of the DMS (B_{iii}). Schematic of silicon optrode recording tip, including tapered optical fiber coupled to a 32-channel silicon probe (B_{iv}). (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 single 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 (or t-test) 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 ($p < 0.00083$). 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: Average z-scored firing rate and s.e.m. 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

Bolkan, Stone, et al.

decreased *post* vs *pre* (red, $n = 4$) neurons recorded from A2a-Cre mice. **(F)** Same as **C** but for two 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 ($p < 0.001$). **(H)** same as **E** but for neurons recorded from the DMS of D1R-Cre mice.

Figure S2: Transgenic mouse lines faithfully report indirect and direct pathways across striatal subregions. (A) Schematic of viral delivery of AAV5-eF1a-DIO-GFP to the dorsomedial striatum (DMS) or nucleus accumbens (NAc) on opposite hemispheres of D1R-Cre mice. Red, blue, and purple squares denote representative areas for stereological quantification of viral co-expression with a *drd1* mRNA probe (RNAScope) in the DMS, NAc core (NAcC), or NAc shell (NAcSh), respectively. (B) Example fluorescent confocal image (63x objective, 5x digital zoom) of the NAc core from a D1R-Cre mouse displaying virally-expressed GFP (green), *drd1* mRNA (red), and DAPI (blue). (C) Percentage of GFP⁺ neurons co-expressing *drd1* mRNA from 2 D1R-Cre mice across the DMS (red; n = 5 sections; 193 GFP⁺ neurons), NAcC (blue; n = 5 sections; 298 GFP⁺ neurons), or NAcSh (purple; n = 4 sections; 312 GFP⁺ neurons). (D) Same as A, but for quantification of viral co-expression with a *drd2* mRNA probe in A2a-Cre mice. (E) Same as B, but for an example image of the DMS from an A2a-Cre mouse and displaying *drd2* mRNA (red). (F) Same as C, but for percentage of virally-expressed GFP⁺ neurons co-expressing *drd2* mRNA in 2 A2a-Cre mice across the DMS (red; n = 4 sections; 312 GFP⁺ neurons), NAcC (blue; n = 4 sections; 326 GFP⁺ neurons), or NAcSh (purple; n = 4 sections; 312 GFP⁺ neurons). (G) Same as A and D, but for quantification of co-expression of *drd2* and *cre* mRNA in 2 D2R-Cre mice. (H) Same as B and E, but for an example image of the NAcSh from a D2R-Cre mouse and displaying *cre* mRNA (green) and *drd2* mRNA (red). (I) *Left*: same as C and F, but for percentage of neurons with *cre* mRNA co-expressing *drd2* mRNA in 2 D2R-Cre mice across the DMS (red; n = 5 sections; 1302 *cre*⁺ neurons), NAcC (blue; n = 5 sections; 1,104 *cre*⁺ neurons), or NAcSh (purple; n = 4 sections; 1,187 *cre*⁺ neurons). *Right*: same as *left* but for neurons with *drd2* mRNA co-expressing *cre* mRNA across DMS (red; n = 5 sections; 1,269 *drd2*⁺ neurons), NAcC (blue; n = 5 sections; 1,055 *drd2*⁺ neurons), or NAcSh (purple; n = 5 sections; 1,114 *cre*⁺ neurons). Solid bars denote mean and s.e.m. throughout. (J) Example fluorescent confocal microscopy image of a coronal section from a DR2-Cre mouse that underwent fluorescent *in situ* hybridization with probes targeting *drd1a* and *drd2* receptor mRNA. *Left*: 20x magnification tilescan spanning dorsal and ventral striatum. *Right top*: 63x confocal images of dorsomedial striatum (DMS, red square) and expression of *drd1a* mRNA (green), *drd2* mRNA (red), and merged image of both (yellow). White triangles indicate co-expression of receptor probes in single neurons. *Right middle*: same as *right top* but for 63x confocal images of nucleus accumbens core (NAcC, blue square). *Right bottom*: same as *right top* but for 63x confocal images of nucleus accumbens shell (NAcSh, purple square). (K) Percentage of *drd2*⁺ neurons co-expressing *drd1a* mRNA from 2 D2R-Cre and 2 D1R-tdTomato mice in the DMS (red; n = 10 sections; 2,423 *drd2*⁺ neurons), NAcC (blue; n = 10 sections; 2,196 *drd2*⁺ neurons), or NAcSh (purple; n = 10 sections; 2,220 *drd2*⁺ neurons). Circles indicate mean overlap from individual sections. (L) Same as K, but for percentage of *drd1a*⁺ neurons co-expressing *drd2* mRNA from 2 D1R-tdTomato mice in the DMS (red; n = 5 sections; 868 *drd1a*⁺ neurons), NAcC (blue; n = 5 sections; 834 *drd1a*⁺ neurons), or NAcSh (purple; n = 5 sections; 874 *drd1a*⁺ neurons).

Figure S3. Indirect and direct pathway inhibition of the DMS is stable across time. (A) Trial-by-trial raster plots of single neuron spiking during laser off baseline (-5 to 0s), 532-nm (5-mW) laser delivery (0 to 5s), and post laser offset (5 to 10s) for all significantly inhibited neurons (n = 18/60) recorded from A2a-Cre mice expressing Cre-dependent NpHR in the DMS. 40 total trials of laser sweeps per recording site (~15 minutes), ordered in time top to bottom. Individual neuron labels indicate: m (mouse), d (day of recording),

Bolkan, Stone, et al.

r/l (right/left hemisphere and penetration number), s (site or depth of recording probe numbered ventral to dorsal), and st (probe stereotrode channel). Number in parenthesis indicates number of spikes sub-sampled for display. Inset displays average (bold) and 100 randomly sampled spike waveforms (grey). **(B)** As in **A** but for all significantly inhibited neurons recorded from D1R-Cre mice expressing Cre-dependent NpHR in the DMS (n =21/50).

Figure S4. Non-significant effects of indirect or direct pathway inhibition compared to non-opsin expressing controls on multiple motor measures during navigation of a virtual corridor.

(A) Schematic of virtual corridor and unilateral delivery of 532-nm light (5-mW) limited to mouse position in the corridor stem (0-200cm). (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 ($p = 0.98$, $F_{2,15} = 0.02$). (C) Same as B but for difference in x-position (cm, on-off) contralateral to the laser hemisphere ($p = 0.60$, $F_{2,15} = 0.53$). (D) Same as C but

for difference in view angle (deg, on-off) contralateral to the laser hemisphere ($p = 0.20$, $F_{2,15} = 1.90$). (E) Same as C but for difference in mean 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 ($p = 0.90$, $F_{2,15} = 0.94$). 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 >110% of maze length, or >363cm) ($p = 0.80$, $F_{2,16} = 0.22$). Solid black lines indicate mean and s.e.m. across mice and transparent 'x' denote individual mouse mean throughout.

Figure S5: No detectable effect of indirect or direct pathway inhibition on spatial preference and speed during a real-time conditioned place preference test.

(A) Schematic of real-time conditioned place preference chamber and bilateral 532-nm laser illumination (5-mW) of the DMS. Left and right sub-chambers of equal size but with repeating vertical or horizontal black-and-white bar patterning distinguished each side, respectively. Mice underwent a 5-min preference test (*left*, Baseline) without any laser illumination, followed by a 20-min preference test (*right*, Test) in which mice received bilateral laser illumination only when occupying one of the two chamber sides (counterbalanced across mice). No illumination (Laser OFF) and illumination (Laser ON) sides during Baseline were defined based on the subsequent Test illumination side. (B) Delta time spent in chamber side (laser OFF - laser ON) during 5-min Baseline (*left*) and 20-min Test (*right*) for mice receiving DMS indirect ($n = 9$ mice) or direct ($n = 20$ mice) pathway inhibition, or DMS illumination alone (no opsin, $n = 9$ mice). Error bars denote mean and s.e.m. Grey transparent 'x' indicates individual mice. p-value denotes one-way ANOVA of group on delta time in the chamber side during Baseline (*left*: $p = 0.73$, $F_{2,27} = 0.31$) or Test (*right*: $p = 0.10$, $F_{2,27} = 2.55$). (C) Average speed when mice occupied laser off (black) or laser on (green) chamber sides during Baseline (*left*) or Test (*right*) for same groups and order as in B. Solid bars indicate mean and s.e.m. Transparent grey lines indicate individual mouse mean. p-value denotes interaction of two-factor (between-subject: group, within-subject: laser) repeated measure ANOVA on speed during Baseline (*left*: group x laser interaction: $p = 0.16$, $F_{1,35} = 2.07$; laser: $p = 0.28$, $F_{1,35} = 1.20$) or Test (*right*: group x laser interaction: $p = 0.07$, $F_{1,35} = 3.6$; laser: $p = 0.10$, $F_{1,35} = 2.8$).

Figure S6: Similar motor performance in three virtual reality T-mazes. (A) Schematic of three virtual reality (VR)-based T-mazes that differ in cognitive requirements. (B) Average y-velocity (cm/s) of mice during the cue region (0-200cm) of the accumulation of evidence task (black, $n = 32$ mice, $n = 52,381$ trials), no distractors (ctrl #1) task (magenta: $n = 34$ mice, 56,953 trials), or permanent cues (ctrl #2) task (cyan: $n = 20$ mice, $n = 27,870$ trials). Solid bars denote mean and s.e.m. across mice while transparent 'x' denote individual mouse mean. p-value denotes one-way ANOVA of task on y-velocity ($p = 0.38$, $F_{2,83} = 0.98$). (C) Same as B but for average x-position (cm) during the cue region (0-200cm) on left and right choice trials. p-value denotes one-way ANOVA of task on x-position (left choice: $p = 0.49$, $F_{2,83} = 0.71$; right choice: $p = 0.42$, $F_{2,83} = 0.89$). (D) Same as B but for average view angle (degrees) during the cue region (0-200cm) on left and right choice trials (left choice: $p = 0.49$, $F_{2,83} = 0.71$; right choice: $p = 0.70$, $F_{2,83} = 0.36$). (E) As in B but for average percent of trials with excess travel (defined as travel >110% of maze length, or >363cm). Accumulation of evidence: $n = 32$ mice, $n = 64,420$ trials; control #2 (no distractors): $n = 34$ mice, $n = 61,308$ trials; control #2 (permanent cues): $n = 20$ mice, $n = 30,150$ trials. p-value denotes one-way ANOVA of task on excess travel ($p = 0.11$, $F_{2,83} = 2.3$). (F) 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 (0-300cm in 25-cm bins). 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 s.e.m. across mice for each task (n as in B). (G) Same as F but for average percent of trials with excess travel based on the trial-by-trial view angle (degrees) of mice (n as in B). (H) As in B but for mean standard deviation in view angle (degrees) per trial. 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 trials for a mouse (n as in E). p-value denotes one-way ANOVA of task on view angle deviation ($p = 0.11$, $F_{2,83} = 2.3$).

Figure S7. Inhibition of indirect and direct DMS pathways has little impact on multiple measures of motor performance across VR-based decision-making tasks, but has small and opponent influence on x-position and view angle. (A) Average y-velocity (cm/s) as a function of y-position (0- 300cm in 25cm 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 = 3,964 laser off and n = 916 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 = 14,533 laser off and n = 3,661 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 = 3,654 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 = 3,752 laser off and n = 866 laser on trials) tasks. (D) Difference in the average distance (cm) travelled (left) and difference in trials (%) with excess travel greater than 110% of maze

length (or >363cm) (right) between laser off and on trials (on-off) for 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, $n = 14,799$ laser off and $n = 3,582$ 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 difference in distance ($p = 0.53$, $F_{2,24} = 0.66$) or excess travel ($p = 0.54$, $F_{2,24} = 0.63$). **(E)** Same as **D** but for difference (on-off) in distance (cm) travelled (left) or 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 $n = 3,991$ 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 difference in distance ($p = 0.10$, $F_{2,25} = 2.5$) or excess travel ($p = 0.45$, $F_{2,25} = 0.82$). **(F)** Same as **D** but for difference (on-off) in distance (cm) travelled (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, $n = 28,556$ laser off and $n = 6,772$ laser on trials), no distractors (magenta, $n = 5$ mice, $n = 4,108$ laser off and $n = 1,001$ laser on trials), or permanent cues ($n = 6$ mice, $n = 4,360$ laser off and $n = 1,037$ laser on trials) tasks. p-value denotes one-way ANOVA of task on difference in distance ($p = 0.18$, $F_{2,21} = 1.9$) or excess travel ($p = 0.24$, $F_{2,21} = 1.5$). **(G)** Same as **D** but for difference (on-off) in per-trial standard deviation in view angle in mice receiving DMS indirect pathway inhibition across tasks ($p = 0.34$, $F_{2,24} = 1.5$). The mean of the standard deviation in view angles sampled in 5-cm steps from 0-300 cm were calculated per trial, and averaged across all laser off (and on) trials for a mouse (n as in **D**). **(H)** Same as **G** but for mice receiving DMS direct pathway inhibition across tasks ($p = 0.28$, $F_{2,25} = 1.3$, n as in **E**). **(I)** Same as **G** but for mice receiving DMS illumination (no opsin) in the absence of NpHR ($p = 0.10$, $F_{2,21} = 2.5$, n as in **F**). **(J)** Difference in x-position (cm) during the cue region (0-200 cm) on laser off and on trials (on-off) when choice was ipsilateral (*left*, ipsi choice) or contralateral (*right*, contra choice) to the laser hemisphere in mice receiving DMS indirect pathway inhibition across tasks (ipsi choice: $p = 0.02$, $F_{2,24} = 4.6$; contra choice: $p = 0.84$, $F_{2,24} = 0.18$, n as in **A**). **(K)** Same as **J** but for mice receiving DMS direct pathway inhibition (ipsi choice: $p = 0.22$, $F_{2,24} = 1.6$; contra choice: $p = 0.02$, $F_{2,24} = 4.9$, n as in **B**). **(L)** Same as **J** but for mice receiving DMS illumination (no opsin) in the absence of NpHR (ipsi choice: $p = 0.18$, $F_{2,18} = 1.9$; contra choice: $p = 0.93$, $F_{2,18} = 0.93$, n as in **C**). **(M)** As in **J** but for difference (on-off) in view angle (degrees) (ipsi choice: $p = 0.21$, $F_{2,24} = 1.7$; contra choice: $p = 0.65$, $F_{2,24} = 0.45$). **(N)** As in **K** but for difference (on-off) in view angle (degrees) (ipsi choice: $p = 0.17$, $F_{2,24} = 1.9$; contra choice: $p = 0.04$, $F_{2,24} = 3.8$). **(O)** As in **L** but for difference (on-off) in view angle (degrees) (ipsi choice: $p = 0.28$, $F_{2,18} = 1.4$; contra choice: $p = 0.73$, $F_{2,18} = 0.32$).

Figure S8. 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 direct pathway of the DMS (dark gray), mice inhibited in the indirect pathway of the DMS (light gray), and mice without 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 the average of five different test sets. Held-out data for test sets was selected as a random 20% of sessions, using the same number of sessions for each mouse. (B) Same as A but with different numbers of previous choice covariates using a 3-state GLM-HMM. (C) Fitted GLM weights for the 4-state model using aggregated data from all mice inhibited in the indirect pathway of the DMS. Error bars denote (+/-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. (D) Same as C but for mice inhibited in the DMS direct pathway. (E) GLM weights fitted to a concatenated data set 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 the posterior standard deviation of each weight. (F) 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 +/-1 SEM 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. (G) Same as F but for data simulated from the model fit to mice in the control group (see Methods).

Figure S9. 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 indirect pathway of the DMS (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).

Figure S10. Individual mice visit multiple types and numbers of states over the course of sessions. (A) The fraction of trials that mice inhibited in the indirect pathway of the DMS spent in each state in each session. Each box represents a different mouse ($n=13$) and each dot in each box represents an individual session for that mouse. Color-coding reinforces the state composition of each session (e.g. blue indicates 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. (B) Same as A but for mice inhibited in the direct pathway of the DMS ($n=13$).

Figure S11. Comparison of motor performance during evidence accumulation across GLM-HMM states with and without indirect and direct pathway inhibition.

(A) Schematic denoting analysis of motor performance across GLM-HMM states on laser off trials only (B-G) in mice unilaterally coupled to a fiberoptic for indirect pathway inhibition. x-position, view angle, and choice are defined relative to the fiberoptic coupled hemisphere on a given session. (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 mice identified as occupying state 1 (blue, n = 12,379 trials), state 2 (yellow, n = 12,168 trials) or state 3 (red, n = 16,939 trials). (C) As in B but for average x-position (cm) on ipsilateral or contralateral choice trials. (D) As in C but for average view angle (degrees) on ipsilateral and contralateral choice trials. (E) Across-mouse average in the per-trial standard deviation in view angle during laser off trials across GLM-HMM states. The mean of the standard deviation in view angles sampled in 5-cm bins from 0-300-cm were calculated for every trial, and then averaged across trials for individual mice (grey lines). Solid bars denote mean and s.e.m. across mice (state 1 (blue): n = 12,781 trials; state 2 (yellow): n = 12,706 trials; state 3 (red): n = 18,194 trials). p-value denotes one-way repeated measures ANOVA of state on view angle deviation (p = 0.10, $F_{2,22} = 2.57$). (F) As in E but for across-mouse average in distance travelled (cm) per trial. p-value denotes one-way repeated measures ANOVA of state on distance (p = 0.03, $F_{2,22} = 3.9$). (G) As in E but for average percent of trials with excess travel (defined as >110% of maze length, or >363-cm). p-value denotes one-way repeated measures ANOVA of state on excess travel (p = 0.0007, $F_{2,22} = 10.2$). (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. (J) As in C but for difference (on-off) in average x-position during the cue region (0-200cm) on ipsilateral and contralateral choice trials. p-value denotes one-way repeated measures ANOVA of state on delta x-position (ipsi choice: p = 0.07, $F_{2,16} = 3.0$; contra choice: p = 0.01, $F_{2,22} = 5.7$; mice with fewer than 5 ipsilateral or contralateral choices in a state were removed from statistical comparison). (K) As in J but for difference (on-off) in average view angle. (L) Same as E but for difference (on-off) in per-trial view angle standard deviation across GLM-HMM states with direct pathway inhibition. p-value denotes one-way repeated measures ANOVA of state on delta view angle deviation (p = 0.96, $F_{2,22} = 0.04$). (M) Same as F but for difference (on-off) in distance (cm) per trial across GLM-HMM states with direct pathway inhibition (p = 0.66, $F_{2,22} = 0.43$). (N) Same as G but for difference (on-off) in percent of trials with excess travel across GLM-HMM states with direct pathway inhibition (p = 0.08, $F_{2,22} = 2.8$). (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 = 10,494 laser off and n = 1,880 laser on trials; state 2, yellow: n = 9,083 laser off and n = 1,827 laser on trials; state 3, red: n = 13,181 laser off and n = 2,392 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 per-trial view angle standard deviation across GLM-HMM states in direct pathway mice (state 1, blue: n = 11,000 trials; state 2, yellow: n = 9,432 trials; state 3, red: n = 14,380 trials). p-value denotes one-way repeated measures ANOVA of state on per-trial view angle standard deviation (p = 0.06, $F_{2,16} = 3.2$). (T) As in F but for distance (cm) in direct pathway mice (p = 0.07, $F_{2,16} = 3.14$). (U) As in G but for percent trials with excess travel in direct pathway mice (p = 0.12, $F_{2,16} = 2.4$). (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 difference (on-off) in y-velocity (cm/s) during GLM-HMM states in direct pathway mice. (X) As in J but for the difference (on-off) in x-position (cm) during GLM-HMM states in direct pathway mice (ipsi choice: p = 0.62, $F_{2,12} = 0.5$; contra choice: p = 0.88, $F_{2,10} = 0.13$; mice with fewer than 5 ipsilateral or contralateral choices in a state were removed from statistical comparison). (Y) As in K but for the difference (on-off) in view angle (degrees) during GLM-HMM states in direct pathway mice (ipsi choice: p = 0.005, $F_{2,12} = 8.3$; contra choice: p = 0.76, $F_{2,10} = 0.28$; mice with fewer than 5 ipsilateral or contralateral choices in a state were removed from statistical comparison). (Z) As in L but for difference (on-off) in per-trial view angle standard deviation (degrees) in direct pathway mice (state 1, blue: n = 11,000 laser off and n = 1,975 laser on trials; state 2, yellow: n = 9,432 laser off and n = 1,938 laser on trials; state 3, red: n = 14,380 laser off and n = 2,657 laser on trials) (p = 0.46, $F_{2,16} = 0.82$). (AA) as in M but for difference (on-off) in distance (cm) in direct pathway mice (p = 0.42, $F_{2,16} = 0.90$). (BB) as in N but for difference (on-off) in percent trials with excess travel in direct pathway mice (p = 0.16, $F_{2,16} = 2.1$).

Figure S12. Behavioral shaping for virtual-reality T-maze tasks. (A) Schema of shaping mazes (top: maze 1-9) and subsequent optogenetic testing mazes (far right) for the accumulation of evidence task. Mazes varied according to the following sensory features: the length of start, cue, and delay regions (green, black, and grey bars and colored text, respectively), whether visual cues were presented 10-cm from cue position (black outline, white filled square) and remained visible (grey square) or disappeared 200-ms after

presentation (black dotted, unfilled square) or were permanently available from trial outset (bold black border, white filled), the presence of left or right whisker air puffs (15-psi, 40-ms) which were delivered upon first instance of being 10-cm from visual cue position (solid vs grey puff symbol), whether a visual guide was located in the rewarded arm (black double square) or if the visual guide was only visible during the cue region (grey double square), the density of cues during the cue region (c.d.), and whether distractor cues occurred on the non-rewarded maze side (side ratio, s.r.: mean density per meter). Following shaping mazes 1-9, optogenetic testing was carried out on mazes 10 and 11 (accumulation of evidence), and maze 12 (no distractors). **(B)** Solid black line depicts the across-mouse median number of sessions spent on each shaping maze (mazes 1-9) until reaching the first testing maze (maze 10) (group median: 22 sessions). Grey transparent lines depict the median number of sessions for individual mice ($n = 87$). **(C)** Cumulative number of sessions to reach each successive maze until the first testing maze (maze 10) (group mean: 23.0 ± 0.8 sessions). Solid black lines depict mean and s.e.m. across mice, and transparent 'x' denote individual mice. **(D)** Total number of sessions spent on each shaping maze. Solid black lines depict mean and s.e.m., and transparent 'x' denote individual mice at each respective shaping maze (maze 1-9). **(E)** Percent correct performance across shaping mazes. Solid black lines depict mean and s.e.m., and transparent 'x' denote individual mice at each respective shaping maze (maze 1-9). **(F)** Same as **A** but for shaping (left, maze 1-6) for the permanent cues task. Following shaping mazes 1-6, optogenetic testing was carried out on mazes 7 and 8 (permanent cues) and maze 12 (no distractors). **(G)** Same as **B** but for permanent cues shaping (group median: 17 days; $n = 20$ mice). **(H)** Same as **C** but for permanent cues shaping (group mean: 18.9 ± 1.5 sessions). **(I)** Same as **D** but for permanent cues shaping. **(J)** Same as **E** but for permanent cues shaping.

Figure S13. Histological confirmation of viral expression and fiberoptic placement. (A) *Top*: two individual mouse examples of cre-dependent NpHR-GFP expression in and fiberoptic targeting of the dorsomedial striatum (DMS) in D2R-Cre or A2a-Cre mice. *Bottom left*: schematic representation of the minimum (dark grey) and maximum (light grey) spread of NpHR expression in all mice targeting the indirect pathway of the DMS (DMS::Indirect, n = 20 mice). *Bottom right*: summary of tapered fiberoptic tip location and angled track for all mice targeting the indirect pathway of the DMS. (B) Same as A but for all experiments targeting the direct pathway of the DMS (DMS::Direct, n = 23 mice). (C) Same as A but for fiberoptic targeting only for all control mice receiving DMS illumination in the absence of NpHR (DMS::NoOpsin, n = 17 mice). (D) Same as A but for experiments targeting the indirect pathway of the nucleus accumbens (NAc::Indirect, n = 12 mice). (E) Same as A but for experiments targeting the direct pathway of the nucleus accumbens (NAc::Direct, n = 10 mice). (F) Same as C but for all control mice receiving NAc illumination in the absence of NpHR (NAc::NoOpsin, n = 9).

Methods

Animals

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 = 23, ER44Gsat, MMRRC-UCD), and *A2a-Cre* (n = 18, KG139Gsat, MMRRC-UCD). An additional 4 *Drd1-Cre* mice, 3 *A2a-Cre* mice, and 2 *Drd2-Cre* mice were used for electrophysiological characterization of halorhodopsin (NpHR)-mediated inhibition, or fluorescent *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-hour light – 12-hour 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 μ 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 (5mg/kg, I.M.), as well as analgesia pre-operatively and 24-hours later in the form of meloxicam injections (2mg/kg, S.C.). A microsyringe pump controlling a 10 μ L glass syringe (Nanofill) 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: AAV5-eF1a-DIO-eNpHR3.0-EYFP-WPRE-hGH (UPenn, 1.3×10^{13} parts/mL) or AAV5-eF1a-DIO-eNpHR3.0-EYFP-WPRE-hGH (PNI Viral Core, 2.2×10^{14} parts/mL, 1:5 dilution). For fluorescence *in situ* hybridization experiments, AAV5-eF1a-DIO-EYFP-hGHpA (PNI Viral Core, 6.0×10^{13} parts/mL) was used to label D1R⁺ and D2R⁺ neurons in D1R-Cre and A2A-Cre transgenic lines. In all experiments, virus was delivered at a rate of 0.2 μ L/min for a total volume of 0.3-0.7 μ L in the DMS, or 0.3-0.4 μ L in the NAc. To accommodate patch fiber coupling, optical fibers were implanted at angles (DMS: 15°, 0.74 mm anterior, 1.1 mm lateral, -3.6 mm ventral; NAc: 10°, 1.3 mm anterior, 0.55 mm lateral, -5.0 ventral) and were then fixed to the skull using dental cement. Mice were allowed to recover and closely monitored for 5 days 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- μ M) craniotomies made using a carbide drill (**Figure S2A**). The craniotomies were filled with a small amount of silicon adhesive (Kwik-Sil, World Precision instruments) and then covered with UV-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 (**Figure S2B**). 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 μ M, 0.22 NA) was stereotaxically inserted under visual guidance of a stereoscope and allowed to stabilize for ~30 minutes. 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 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 forty deliveries of 5-s light (5-mW, measured from fiber tip), separated by 15-s intertrial intervals. From 1-3 recordings were made at different depths within a single probe penetration (minimum separation of 300- μ M), with each hemisphere receiving 1-3 penetrations at different medial-lateral or anterior-posterior 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).

VR 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 (**Figure 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 ViRMEn⁶⁵ 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 anterior-posterior (and medial-lateral) 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 was given by $\sqrt{dY^2/dt}$, where dt was 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: $\omega = \text{atan2}(-dX \cdot \text{sign}(dY), |dY|)$. Then the rate of change of view angle (θ) for each sampling of the sensor was given by:

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

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 in order to allow sharp turns into the maze arms³³.

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 μ l of 10% sweetened condensed milk in water v/v) to a lick tube, or the release of air flow (40-ms, 15 psi) to 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, 200mW).

Behavioral shaping. Following post-surgical recovery, over the course of 4-7 days 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^{33,66} (**Figure S12A**) or the permanent cues (control #2) task (**Figure S12F**).

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 (**Figure S12A,F**). In mazes 5-7 of the evidence accumulation shaping procedure (**Figure S12A**), 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 (**Figure S12F**). 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⁻¹). 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 maze 7 and 8. In all evidence accumulation shaping mazes (maze 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 (**Figure S12A**, maze 10) to 7.7::2.3 m⁻¹ (**Figure S12A**, maze 11). In contrast, throughout all shaping (maze 1-6) and testing mazes (maze 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 +/- 0.8 (**Figure S12B-C**). Mice typically spent between 2-5 sessions on each shaping maze before progressing to the next, with performance increasing or remaining stable throughout (**Figure S12D-E**; maze 9: 74.1 +/- 9.8 percent 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 +/- 1.5 (**Figure S12G-J**). Mice typically spent between 2-4 sessions on each shaping maze before progressing to the next, with performance increasing or remaining largely stable throughout (**Figure S12G-J**; maze 6: 87.0 +/- 4.3 percent 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 (**Figure 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 were +/- 25-cm relative to the center of the maze stem). Mice received reward (~4-8 µL of 10% v/v 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 m^{-1} for the rewarded side and 1.6 m^{-1} for the non-rewarded side (**Figure S12A**, maze 10) or $7.3::2.3 \text{ m}^{-1}$ (**Figure S12A**, maze 11). Note that only maze 10 data was used for cross-task comparisons of optogenetic effects with permanent cues and no distractors control tasks in order to precisely match cue presentation statistics (**Figure 2, 3, S6, S7**). 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 (**Figure S12A**, maze 4), with mice progressing to the evidence accumulation testing maze after 10 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 (**Figure S12A**, maze 7). These “easy blocks” were limited to 10 trials, after which mice returned to the main testing maze regardless of performance. Behavioral sessions lasted for ~1-hour and typically consisted of ~150-200 trials.

All features of the “no distractors” (control #1) task (**Figure 2B**, magenta, *middle*; **Figure S12A** or **S12G**, maze 12) were identical to the evidence accumulation task (**Figure S12A**, maze 10) except: (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 (**Figure 2B**, cyan, *right*; **Figure S12G**, maze 8) were identical to the evidence accumulation task except: (1) 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 (maze 7 and maze 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 (**Figure 2, 3, S6, S7**).

To discourage side biases, in all tasks we used a previously implemented debiasing algorithm²⁶. This was achieved by changing the underlying probability of drawing a left or a right trial according to a balanced method described in detail elsewhere⁶⁷. In brief, the probability of drawing a right trial, p_R , is given by:

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

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 $[0.15, 0.85]$. Because the empirical fraction of drawn right trials could significantly deviate from p_R , particularly when the number of trials is small, we applied an additional pseudo-random drawing prescription to p_R . Specifically, if the empirical fraction of right trials (calculated using a $\sigma = 60$ trials half-Gaussian weighting window) is above p_R , right trials were drawn with probability $0.5 p_R$, whereas if this fraction is below p_R , right trials were drawn with probability $0.5 (1 + p_R)$.

Virtual corridor. Following shaping in the behavioral tasks above 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 (**Figure 1E-F**). This included a start region (-10 to 0-cm), a reward location (310-cm) in which mice received 4 μ L of 10% v/v 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 over the course of ~70 minute sessions. The virtual environment was controlled by the ViRMEn software engine, with real-to-virtual world movement transformations as described above.

Optogenetics during VR behavior

According to a previously published protocol²⁵, optical fibers (200 μ M, 0.37 NA) were chemically etched using 48% hydrofluoric acid to achieve tapered tips of lengths 1.5-2 mm (DMS-targeted) or 1-1.5 mm (NAc-targeted). Following behavioral shaping in VR (and >6weeks 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 multi-mode fiberoptic patch cable (PFP, 62.5 μ 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 (**Figure 1** and **S4**), or the cue region (0-200cm) of each T-maze decision-making task (**Figure 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 (**Figure S5**). The CPP apparatus consisted of a rectangular Plexiglass box with two chambers (29-cm x 25-cm) separated by a clear portal in the center. The same grey, 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 minutes. 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 analysis.

Behavior analyses

Data selection. For cross-task comparisons of optogenetic inhibition (**Figure 2, 3, S6, S7**) we analyzed only trials from evidence accumulation maze 10 (**Figure S12A**), “no distractors” maze 12 (**Figure S12A** or **S12G**), and “permanent cues” maze 8 (**Figure S12G**), which each followed matching cue probability statistics (except for the by-design removal of distractors in the “no distractors” control task). For model-based analyses of the evidence accumulation task (**Figure 4-7, S8-11**) both maze 10 and maze 11 data were included, which differed only modestly in the side ratio of reward to non-reward side cues (**Figure 12A**, ~50% of trials were maze 10 or 11). In all tasks and all analyses throughout we removed initial warm-up blocks (**Figure S12A**, maze 4, approximately 5-15% of total trials). For model-based analyses of the evidence accumulation task (**Figure 4-7, S8-10**), we included interspersed “easy blocks” capped at 10 trials in length (**Figure S12A**, maze 7, see description above). These trial blocks comprised approximately ~5% of total trials, were included to avoid gaps in trial history, and were treated identical to the main

evidence accumulation mazes by the models. These trials were removed from cross-task comparisons of optogenetic inhibition (**Figure 2, 3, S6, S7**).

For analysis of optogenetic inhibition during virtual corridor navigation (**Figure 1 and S4**) we removed trials with excess travel >110% of maze length (or >363-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; exact trial N in figure legends), but 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 larger 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 in order to detect potential effects of pathway-specific DMS inhibition (or DMS illumination alone) on these measures (**Figure 1 and S4**).

Similarly, for all cross-task comparisons of optogenetic inhibition (**Figure 2, 3, S6, S7**) 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. No trials were excluded for cross-task comparisons of laser effects on measures of distance, per-trial view angle standard deviation, and trials with excess travel (exact trial N in figure legends).

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 (**Figure 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 on trials as the difference in choice accuracy (% correct) between trials where sensory evidence indicated a contralateral reward versus when sensory evidence indicated an ipsilateral reward (contralateral-ipsilateral, positive values indicate greater contralateral choice bias) (**Figure 3D,G,J,O**). Delta choice bias was calculated as the difference in contralateral choice bias between laser off and on trials (on-off, positive values indicate laser-induced contralateral choice bias) (**Figure 3E,H,K,P**).

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 (#contra-#ipsi). Psychometric curves were fit to the following 4-parameter sigmoid:

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

where λ and γ are the right and left lapse rates, respectively, σ is the offset, μ is the slope, and δ is the difference in the number of contralateral and ipsilateral cues on a given trial. In **Figure 4E-F**, we took the difference in the number of contralateral and ipsilateral cues and mouse choice on each trial and used maximum likelihood fitting⁶⁸ to fit all the data (due to the relatively small number of trials per mouse per state) to the same 4-parameter sigmoid (equation 3). For the individual points plotted in **Figure 4E-F**, we binned the difference in cues in increments of 4 from -16 to 16 and calculated the percentage of contralateral choice trials for each bin.

Motor performance indicators. Y-velocity (cm/s) was calculated on every sampling iteration (120 Hz, or every ~8ms) 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 from 0-300-cm in 25-cm bins were then averaged 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 standard error of the mean across mice reported throughout (**Figure 1G, 2D, S7A-C, S11B, S11P, S11I, and S11W**; averaged across y-position 0-200cm in **Figure S4B and S6B**).

X-position trajectory (cm) as a function of y-position was calculated per trial by first taking the x-position at y-positions 0-300cm 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 across y-position bins of 25-cm from 0 to 300cm. 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 standard error of the mean across mice was reported throughout (**Figure 1H, 2E, S11C, and S11Q**; averaged across y-position 0-200cm in **Figure S4, S6, S7J-L, S11J, S11X**). Average view angle trajectory (degrees) was calculated in the same manner as x-position (**Figure 1I, 2F, S7J-L, S11D, and S11R**; average across y-positions 0-200cm in **Figure S4D, S6D, S7M-O, S11K, and S11Y**). View angle standard deviation was calculated by first sampling the per-trial view angle from 0-300cm 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 standard error of the mean across mice are reported throughout (**Figure S4E, S6H, S7G-I, S11E, S11S, S11L, and S11Z**). This measure sought to capture unusually large deviations in single trial view angles, which would be indicative of excessive turning or rotations.

Distance 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 was then averaged across trials for individual mice and the average and standard error of the mean across mice was reported throughout (**Figure 1J, 2G, S4F, S7D-F left, S11F, S11T, S11M and S11AA**). Excess travel was defined as the fraction of trials with total distance travelled per trial (calculated as above) greater than 110% of maze length (or >363cm). The average and standard error of the mean across mice was reported throughout (**Figure S4G, S6E, S7D-F right, S11G, S11U, S11N and S11BB**).

Decoding of choice based on the trial-by-trial x-position (**Figure S6F**) or view angle (**Figure S6G**) of mice was carried out by performing a binomial logistic regression using the MATLAB function *glmfit*. The logistic regression was fit separately for individual mice at successive y-positions in the T-maze stem (0-300cm 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. 5-fold cross-validation (re-sampled for new folds 10 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 decoded choices matching actual mouse choice, 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>.

General statistics. We performed one-way ANOVAs of the factors task (three levels: evidence

accumulation/“no distractors”/“permanent cues”) for effects on choice (**Figure 1C**), distance travelled (**Figure 1G** and **S7D-F left**), y-velocity (**Figure S6B**), x-position on left or right choice trials (**Figure S6C** and **S7J-L**), view angle on left or right choice trials (**Figure S6D** and **S7M-O**), trials with excess travel (**Figure S6E** and **S7D-F, right**), and per-trial standard deviation in view angle (**Figure S6H** and **S7G-I**). We performed one-way ANOVAs of the factor group (three levels: indirect pathway inhibition, direct pathway inhibition, no opsin illumination) for effects on y-velocity (**Figure S4B**), x-position (**Figure S4C**), view angle (**Figure S4D**), per-trial standard deviation in view angle (**Figure S4E**), distance travelled (**Figure S4F**), and fraction of trials with excess travel (**Figure S4G**). We performed a repeated-measure one-way ANOVA on the within-subject factor state (three levels: GLM-HMM state 1/ state 2/ state 3) on x-position during ipsilateral and contralateral choice trials (**Figure S11J** and **S11X**), view angle during ipsilateral and contralateral choice trials (**Figure S11K** and **S11Y**), per-trial standard deviation in view angle (**Figure S11E, S11L, S11S, and S11Z**), distance travelled (**Figure S11F, S11M, S11T, and S11AA**), and fraction of trials with excess travel (**Figure S11G, S11N, S11U, and S11BB**). For post-hoc comparisons between indirect or direct pathway inhibition groups and no opsin control mice (**Figure 3E, 3H, 3K and 3P**), we used non-parametric, unpaired, and two-tailed Wilcoxon rank sum tests. Due to multiple group comparisons we only considered a p-value below 0.016 significant (or 0.05/3 groups).

Bernoulli GLM

Coding of covariates and choice output. We coded the external covariates (referred to as inputs in **Figure 4B**) and output (the mouse’s choice) on each trial as follows:

Δ cues: an integer value from -16 to 16, divided by the standard deviation of the Δ 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 (e.g. 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 generalized linear model (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 1-state GLM-HMM (**Figure S8**). 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:

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

We then 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 $\mathbf{Y} = y_1, \dots, y_T$ given all the external covariates $\mathbf{X} = x_1, \dots, x_T$, considered as a function of the model parameters:

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

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, which regularizes by penalizing large weight values. We computed the posterior standard deviation of the fitted GLM weights (shown as error bars in **Figure 4C-D**) by taking the diagonal elements of the inverse negative Hessian (matrix of second derivatives) of the log-likelihood at its maximum^{69,70}.

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 as follows:

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

Each GLM has a unique set of weights w_k that map the external covariates x_t (coded as described in the section **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 equation 4

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

where now the choice probability on each trial is dependent on both the external covariates (inputs) and the state on that trial and is determined by state-dependent GLM weights⁴³⁻⁴⁵. We refer to these as observation probabilities, which can be abbreviated as follows:

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

Fitting. We fit the GLM-HMM to the data using the expectation-maximization (EM) algorithm⁴⁴. 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 optimum⁷⁰.

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

$$\log(p(\mathbf{Y}|\mathbf{X}, \theta)) = \log\left(\sum_{\mathbf{Z}} p(\mathbf{Y}, \mathbf{Z}|\mathbf{X}, \theta)\right) \quad (9)$$

The set of parameters θ 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 from a Dirichlet distribution with a larger concentration parameter over the diagonal entries ($\alpha_{ii} = 5$, $\alpha_{ij} = 1$), reflecting the fact that the probability of staying in the same state from one trial to the next should be larger than the probability of transitioning to a different state. 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 (see **Bernoulli GLM**) to find w_0 . Then, since we didn't want the initial weights to be the same in each state, we initialized $w_k = w_0 + \epsilon_k$ where $\epsilon \sim \mathcal{N}(0, 0.2)$.

The goal of the E-step of the EM algorithm is to compute $p(\mathbf{Z}|\mathbf{X}, \mathbf{Y}, \theta)$, 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 or the Baum-Welch algorithm⁴⁴. 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

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

where $c_t = p(y_t|y_{1:t-1})$ is a scale factor, $\hat{\alpha}(z_0)$, the prior probability over states before any data are observed, is given as a uniform distribution over states, and K is the total number of states. Note that this is a normalized version of the algorithm that avoids underflow errors (see Bishop Chapter 13).

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

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

where $\beta(z_T) = 1$.

From these two conditional probabilities we can calculate the marginal posterior probabilities of the latent states:

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

which was the goal of the E-step. We can also compute the joint posterior distribution of two successive latents:

$$\xi(z_t, z_{t+1}) = \hat{\alpha}(z_t) \hat{\beta}(z_{t+1}) P \phi \quad (13)$$

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 of the latents and uses them to update the values of the model parameters (equations 6-8) by maximizing the expression for P and w . Since the transition probabilities are fixed, we can compute their updates using the closed form solution:

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

This closed-form update can be derived by applying the appropriate Lagrange multipliers to the complete-data log-likelihood function⁷⁰.

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

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

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-likelihood. We alternated E and M steps until the difference between the log-likelihoods over ten iterations was smaller than a given tolerance (tol = 1e-3). Because the EM algorithm only guarantees that the log-likelihood will converge upon a local optimum⁷⁰, 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 +/- 0.05) in order 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 **Figure 6A-B** and **Figure S8C-E**) by taking the inverse Hessian of the optimized log-likelihood.

Model selection. In **Figure S8A**, we performed cross-validation on the data from both the indirect and direct pathway inhibition groups. 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 units of bits per session (bps), defined:

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

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:

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

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_{bps} from five different test sets. We followed the same procedure as above in **Figure S8B**, selecting the optimal number of previous choices using a 3-state GLM-HMM under parameterizations of 1-8 previous choices while holding the number of all other external inputs (Δ cues, laser, bias, and previous rewarded choice) constant.

Testing. In **Figure 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):

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

Here, $\hat{L}_m$ is the optimized log-likelihood of the model in question (either the GLM or 3-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 **Figure 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 equation 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, labelling 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 equations 7 & 8, and then calculated the predictive distribution as follows:

$$p(y_t | x_{1:t-1}, y_{1:t-1}) = \sum_{k=1}^3 p(y_t | x_t, z_t = k) p(z_t = k | x_{1:t-1}, y_{1:t-1}) \quad (19)$$

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 (**Figure 6C-I, 7G,J, S8F-G, S9E-H**), we assigned each trial to the maximum posterior probability over states given the inputs and choice data:

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

Simulating data. For the analyses in **Figure 6E-F, Figure S8G, and Figure S9**, we evaluated the ability of the 3-state GLM-HMM to predict choices and state transitions that matched the animals' actual behavior in each state. For the covariates for the simulation, we kept the evidence (Δ 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 (equations 7 & 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 fitted the model to the simulated data using the same procedure described previously (see *Fitting*) to obtain the posterior probability over states. For **Figure 6E-F** and **Figure S8G**, 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 (**Figure 7K-M**), we fit the 3-state GLM-HMM using the same general procedure as described above. However, in the case where we disallowed transitions during a session, (**Figure 7K**), the transition matrix was fixed to the identity matrix and we only fit the state-dependent GLM weights. In the case where we disallowed transitions in and out of state 2 (**Figure 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 (**Figure 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 **Figure 7M**, solid curves denote the average log-likelihood for five different test sets. Held-out data for test sets was selected as a random 20% of sessions, using the same number of sessions for each mouse.

Fluorescent *in situ* hybridization and stereological quantification

In situ hybridization (**Figure S2**) 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 fluorescent signal in these lines. We therefore relied on Cre-dependent viral expression of AAV5-DIO-EYFP to report Cre⁺ neurons alongside Drd1a and Drd2 probes in sections from 2 D1R-Cre and 2 A2R-Cre mice, but used all three probes in sections from 2 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 DMSO according to Perkin Elmer instructions and diluted 1:1200 in TSA buffer included in the RNAscope kit. Post *in situ* hybridization slides were cover-slipped using Fluoromount-G containing DAPI (SouthernBiotech).

We then obtained 20x confocal z-stacks from the DMS, NAcCore, and NAcShell in all lines and manually quantified specificity, penetrance, and D1R⁺/D2R⁺ overlap using LASX software (Leica). Specificity was determined as the percentage of the following: GFP⁺ neurons co-expressing Drd1 in D1R-Cre mice (DMS, n = 5 sections, 193 cells; NAcCore, n = 5 sections, 298 cells; NAcShell, n = 5 sections, 363 cells), GFP⁺ neurons co-expressing 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⁺ neurons co-expressing 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⁺ neurons co-expressing 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⁺ neurons. Quantification of D1R⁺/D2R⁺ overlap in striatal regions was carried out on 2 D2R-Cre mice and/or 2 D1R-tdTomato mice and measured as both the percentage of D1R⁺ neurons that were D2R⁺ (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⁺ neurons that were D1R⁺ (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 (i.p.) and transcardially perfused with 1x phosphate-buffered saline (PBS), followed by fixation with 4% paraformaldehyde (PFA). Whole brains with intact fiberoptic implants were post-fixed in 4% PFA for 1-3 days, followed by brain dissection and another 24 hours of post-fixation in PFA. For optogenetic experiments, brains were then transferred to PBS for coronal sectioning (50 μ M) on a vibratome. Viral expression and fiberoptic placements were assessed under slide-scanning (NanoZoomer, Hamamatsu) or single slide (Leica) epifluorescent microscopes. For FISH experiments, post-fixation dissected brains were transferred through a sucrose gradient: 10% sucrose in PBS for 6-8 hours, 20% sucrose in PBS overnight, and 30% sucrose in PBS overnight. Coronal sections (18 μ M) containing the DMS and NAc were made using a cryostat, mounted uncoverslipped on Superfrost plus slides (Fisher), and stored at -80° prior to the FISH protocol. After the FISH protocol, tile-scanning and cellular resolution images of cover-slipped slides were acquired using a confocal microscope (Leica TCS SP8).

Code Availability

Code used for analysis of the data that support the findings of this study is available on github upon publication.

Data Availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

1. Albin, R. L., Young, A. B. & Penney, J. B. The functional anatomy of basal ganglia disorders. *Trends Neurosci.* **12**, 366–375 (1989).
2. Alexander, G. E. & Crutcher, M. D. Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends Neurosci.* **13**, 266–271 (1990).

3. Mink, J. W. THE BASAL GANGLIA: FOCUSED SELECTION AND INHIBITION OF COMPETING MOTOR PROGRAMS. *Prog. Neurobiol.* **50**, 381–425 (1996).
4. Kravitz, A. V. *et al.* Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature* **466**, 622–626 (2010).
5. Roseberry, T. K. *et al.* Cell-Type-Specific Control of Brainstem Locomotor Circuits by Basal Ganglia. *Cell* **164**, 526–537 (2016).
6. Bartholomew, R. A. *et al.* Striatonigral control of movement velocity in mice. *Eur. J. Neurosci.* **43**, 1097–1110 (2016).
7. Bakhurin, K. I. *et al.* Opponent regulation of action performance and timing by striatonigral and striatopallidal pathways. *bioRxiv* 2019.12.26.889030 (2019)
doi:10.1101/2019.12.26.889030.
8. Lobo, M. K. *et al.* Cell Type–Specific Loss of BDNF Signaling Mimics Optogenetic Control of Cocaine Reward. *Science* **330**, 385–390 (2010).
9. 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).
10. Yttri, E. A. & Dudman, J. T. Opponent and bidirectional control of movement velocity in the basal ganglia. *Nature* **533**, 402–406 (2016).
11. 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).
12. 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.e5 (2018).
13. Lee, J., Wang, W. & Sabatini, B. L. Anatomically segregated basal ganglia pathways allow parallel behavioral modulation. *Nat. Neurosci.* **23**, 1388–1398 (2020).

14. Soares-Cunha, C. *et al.* Activation of D2 dopamine receptor-expressing neurons in the nucleus accumbens increases motivation. *Nat. Commun.* **7**, 1–11 (2016).
15. 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).
16. 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).
17. 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).
18. Geddes, C. E., Li, H. & Jin, X. Optogenetic Editing Reveals the Hierarchical Organization of Learned Action Sequences. *Cell* **174**, 32–43.e15 (2018).
19. 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).
20. 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).
21. London, T. D. *et al.* Coordinated Ramping of Dorsal Striatal Pathways preceding Food Approach and Consumption. *J. Neurosci.* **38**, 3547–3558 (2018).
22. 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).
23. 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).
24. Lau, B. & Glimcher, P. W. Value Representations in the Primate Striatum during Matching

- Behavior. *Neuron* **58**, 451–463 (2008).
25. 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).
 26. Shadlen, M. N. & Shohamy, D. Decision Making and Sequential Sampling from Memory. *Neuron* **90**, 927–939 (2016).
 27. Hikosaka, O., Kim, H. F., Yasuda, M. & Yamamoto, S. Basal ganglia circuits for reward value-guided behavior. *Annu. Rev. Neurosci.* **37**, 289–306 (2014).
 28. 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).
 29. Yin, H. H. *et al.* Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nat. Neurosci.* **12**, 333–341 (2009).
 30. Liljeholm, M. & O'Doherty, J. P. Contributions of the striatum to learning, motivation, and performance: an associative account. *Trends Cogn. Sci.* **16**, 467–475 (2012).
 31. Scimeca, J. M. & Badre, D. Striatal contributions to declarative memory retrieval. *Neuron* **75**, 380–392 (2012).
 32. Akhlaghpour, H. *et al.* Dissociated sequential activity and stimulus encoding in the dorsomedial striatum during spatial working memory. *eLife* **5**, e19507 (2016).
 33. Pinto, L. *et al.* An Accumulation-of-Evidence Task Using Visual Pulses for Mice Navigating in Virtual Reality. *Front. Behav. Neurosci.* **12**, (2018).
 34. Engelhard, B. *et al.* Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. *Nature* **570**, 509–513 (2019).
 35. Pinto, L. *et al.* Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. *Neuron* **104**, 810–824.e9 (2019).

36. Koay, S. A., Thiberge, S., Brody, C. D. & Tank, D. W. Amplitude modulations of cortical sensory responses in pulsatile evidence accumulation. *eLife* **9**, e60628 (2020).
37. Owen, S. F., Liu, M. H. & Kreitzer, A. C. Thermal constraints on in vivo optogenetic manipulations. *Nat. Neurosci.* **22**, 1061–1065 (2019).
38. Cruz, B. F., Soares, S. & Paton, J. J. *Striatal circuits support broadly opponent aspects of action suppression and production*. <http://biorxiv.org/lookup/doi/10.1101/2020.06.30.180539> (2020) doi:10.1101/2020.06.30.180539.
39. Raimondo, J. V., Kay, L., Ellender, T. J. & Akerman, C. J. Optogenetic silencing strategies differ in their effects on inhibitory synaptic transmission. *Nat. Neurosci.* **15**, 1102–1104 (2012).
40. Mahn, M., Prigge, M., Ron, S., Levy, R. & Yizhar, O. Biophysical constraints of optogenetic inhibition at presynaptic terminals. *Nat. Neurosci.* **19**, 554–556 (2016).
41. Koay, S. A., Thiberge, S. Y., Brody, C. D. & Tank, D. W. *Neural Correlates of Cognition in Primary Visual versus Neighboring Posterior Cortices during Visual Evidence-Accumulation-based Navigation*. <http://biorxiv.org/lookup/doi/10.1101/568766> (2019) doi:10.1101/568766.
42. 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).
43. Bengio, Y. & Frasconi, P. An Input Output HMM Architecture. *Adv. Neural Inf. Process. Syst.* **7**, 427–234 (1994).
44. 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).
45. Calhoun, A. J., Pillow, J. W. & Murthy, M. Unsupervised identification of the internal states that shape natural behavior. *Nat. Neurosci.* **22**, 2040–2049 (2019).

46. Ashwood, Z. C. *et al.* Mice alternate between discrete strategies during perceptual decision-making. *bioRxiv* 2020.10.19.346353 (2021) doi:10.1101/2020.10.19.346353.
47. Cui, G. *et al.* Concurrent activation of striatal direct and indirect pathways during action initiation. *Nature* **494**, 238–242 (2013).
48. Barbera, G. *et al.* Spatially Compact Neural Clusters in the Dorsal Striatum Encode Locomotion Relevant Information. *Neuron* **92**, 202–213 (2016).
49. 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).
50. Markowitz, J. E. *et al.* The Striatum Organizes 3D Behavior via Moment-to-Moment Action Selection. *Cell* **174**, 44–58.e17 (2018).
51. Parker, J. G. *et al.* Diametric neural ensemble dynamics in parkinsonian and dyskinetic states. *Nature* **557**, 177–182 (2018).
52. 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).
53. Meng, C. *et al.* Spectrally Resolved Fiber Photometry for Multi-component Analysis of Brain Circuits. *Neuron* **98**, 707–717.e4 (2018).
54. Donahue, C. H., Liu, M. & Kreitzer, A. C. Distinct value encoding in striatal direct and indirect pathways during adaptive learning. *bioRxiv* 277855 (2018) doi:10.1101/277855.
55. Pisupati, S., Chartarifsky-Lynn, L., Khanal, A. & Churchland, A. K. *Lapses in perceptual decisions reflect exploration*. <http://biorxiv.org/lookup/doi/10.1101/613828> (2019) doi:10.1101/613828.
56. Peters, A. J., Steinmetz, N. A., Harris, K. D. & Carandini, M. Striatal activity reflects cortical activity patterns. *bioRxiv* 703710 (2019) doi:10.1101/703710.

57. 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).
58. 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. *bioRxiv* 675850 (2020) doi:10.1101/675850.
59. 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).
60. Ahilan, S. *et al.* Forgetful inference in a sophisticated world model. *bioRxiv* (2018) doi:10.1101/419317.
61. Engel, T. A. *et al.* Selective modulation of cortical state during spatial attention. *Science* **354**, 1140–1144 (2016).
62. Juavinett, A. L., Erlich, J. C. & Churchland, A. K. Decision-making behaviors: weighing ethology, complexity, and sensorimotor compatibility. *Curr. Opin. Neurobiol.* **49**, 42–50 (2018).
63. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
64. Dayan, P. & Daw, N. D. Decision theory, reinforcement learning, and the brain. *Cogn. Affect. Behav. Neurosci.* **8**, 429–453 (2008).
65. Aronov, D. & Tank, D. W. Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System. *Neuron* **84**, 442–456 (2014).
66. Pinto, L. *et al.* Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. *Neuron* **104**, 810–824.e9 (2019).
67. Hu, F., Zhang, L.-X. & He, X. Efficient Randomized-Adaptive Designs. *Ann. Stat.* **37**, 2543–2560 (2009).

Bolkan, Stone, et al.

68. Wichmann, F. A. & Hill, N. J. The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept. Psychophys.* **63**, 1293–1313 (2001).
69. 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).
70. Bishop, C. M. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. (Springer-Verlag, 2006).

Acknowledgements

We would like to thank the entire BRAINCoGs team as well as the Witten and Pillow labs 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 F32MH118792 (SSB), U19 NS104648-01 (JWP, IBW), F32NS101871 (LP), K99MH120047 (LP), 1R01MH106689 (IBW), and the New York Stem Cell Foundation (IBW). IBW is a NYSCF—Robertson Investigator.

Affiliations

Princeton Neuroscience Institute, Princeton University Princeton, NJ, USA

Scott S. Bolkan, Iris R. Stone, Lucas Pinto, Zoe C. Ashwood, Jorge M Iravedra-Garcia, Alison L. Herman, Priyanka Singh, Ben Engelhard, Sue A. Koay, Jonathan W. Pillow & Ilana B. Witten

Department of Psychology, Princeton University Princeton, NJ, USA

Jonathan W. Pillow & Ilana B. Witten

Contributions

S.S.B. performed the experiments with support from J.M.I., A.L.H., and P.S. I.R.S. and J.W.P. developed the GLM-HMM, and S.S.B. and I.R.S. analyzed the data. L.P., Z.C.A., B.E. and S.A.K. provided technical and analysis support. 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.

Ethics declarations

Competing Interests

Bolkan, Stone, et al.

The authors declare no competing financial interests.