SUPPLEMENTARY NOTE

We undertook a multipronged analytical approach to assess the effects of unilateral superior colliculus (SC) inactivation on perceptual decision-making in monkeys. First, we evaluated changes in psychometric functions and mean reaction times (RT) pre- and post-muscimol for decisions to and away from the inactivated field (IF; Fig. 2). Second, we evaluated changes in the shapes of RT distributions pre- and post-muscimol for decisions to and away from the IF as the shapes of RT distributions provide reliable ways to distinguish changes in different decision processes[1, 2] (Fig. 5 and Extended Data Fig. 7). Third, we simulated DDMs with changes only in specific parameters from pre- to post-muscimol to assess and compare predicted changes in psychophysical performance and mean RTs that may occur if different aspects of decision-making were changed post-muscimol (Fig. 5, Extended Data Fig. 5). Fourth, we fit a hierarchical drift-diffusion model (HDDM), and estimated parameters to see what parameters changed after SC inactivation (Extended Data Fig. 6). Fifth, to determine which parameter best explained the changes observed with muscimol in the SC, we fitted HDDM variants to the data allowing different model parameters to vary individually and compared these to the HDDM with all parameters allowed to vary to determine which model variant best explained the results of the muscimol inactivation (Supplementary Table 4 and Extended Data Fig. 7). We also fitted a non-hierarchical DDM, thus we used two different methods of parameter estimation (hierarchical Bayesian estimation and quantile maximum products estimation, QMPE) to ensure robustness (Extended Data Fig. 6). Finally, to determine whether an altered urgency signal explained the effects of unilateral inactivation of the SC during decision-making, we also fit UGM variants to the pre- and post-muscimol data to see which parameters changed due to the muscimol. We then compared UGM variants with the urgency signal or the drift rate parameters free to vary to determine which

model variant best explained the post-muscimol data (Supplementary Table 4 and Extended Data Fig. 7).

Signal detection theory (SDT) analysis

In addition to the logistic fits described in the Methods and shown in Fig. 2, we also calculated *d'* and *c* to assess whether unilateral inactivation of the SC with muscimol impaired perceptual sensitivity or decision criterion, respectively. To calculate *d*', we used the following equation modified from Macmillan[3];

$$d' = \frac{1}{\sqrt{2}}\begin{pmatrix} invcdf(P(away\ IF\ choice \mid away\ IF\ trial)) \\ -invcdf(P(away\ IF\ choice \mid to\ IF\ trial) \end{pmatrix} \quad \text{(S1)}$$

To calculate the position of the decision criterion we also used the equation from Macmillan[3];

$$c = -0.5\begin{pmatrix} invcdf(P(away\ IF\ choice \mid away\ IF\ trial)) \\ +invcdf(P(away\ IF\ choice \mid to\ IF\ trial) \end{pmatrix} \quad \text{(S2)}$$

Extended Data Fig. 3 shows the *d'* and *c* results for both monkeys. Overall, there was a significant change in criterion after muscimol that produced a decision bias away from the IF, whereas sensitivity, *d'*, showed no significant change after muscimol. Neither *d'* nor *c* changed after saline injections.

Non-parametric reaction time (RT) analysis

To visualize and analyze the chronometric functions from the RT version of the decision task (seven in monkey S and two in monkey B for muscimol injections, and four in monkey S for saline injections), we used least squares regression to fit the RT data with a linear function of the form:

$$RT = b_0 + b_1 k \quad \text{(S3)}$$

We fitted the RT function separately for the positive (toIF) and negative (awayIF) coherences (k). We extracted the slope parameter (b_1), to determine coherence dependent changes in RT from pre- to post-muscimol and recovery, and the intercept parameter (b_0), to assess non-coherence dependent changes in RT from pre- to post-muscimol and recovery. The pre- and post-muscimol data and results appear in Fig. 2 and the recovery data appear in Extended Data Fig. 4. Coherence dependent changes in RT occurred for toIF decisions and not for awayIF decisions, indicating that the SC inactivation affected decision and not purely motor processes.

Drift-diffusion model (DDM) simulations

DDMs assume that decision-makers accumulate a representation of noisy sensory evidence over time. An evidence path on one trial reflects an internal representation of orientation evidence obtained by visually sampling the Glass pattern until a fixed amount of evidence is reached for a decision, reported by a saccade, to or away from the IF. Different aspects of the decision-making processes are reflected in different parameters of the DDM as described in the Methods[2] and seen in Fig. 5 and Extended Data Fig. 5. The DDM process was defined and simulated by the following equations:

$$E(t) = \delta_k \Delta t + \varsigma(\sqrt{\Delta t})\, N(0,1) \quad \text{(S4)}$$

$$x(t) = x(t\text{-}1) + E(t) \quad \text{(S5)}$$

$$x(t=0) = w*a \quad \text{(S6)}$$

$$x(t) \geq a \rightarrow \textit{choose toIF} \quad \text{(S7)}$$

$$x(t) \leq 0 \rightarrow \textit{choose awayIF} \quad \text{(S8)}$$

The evidence at time (t), denoted by $E(t)$, depends on the drift rate (δ_k) for each condition (k) with an accumulation noise (ς), for every infinitely small time step Δt. In simulations of DDMs and in model fitting using QMPE (see below), the small time step is estimated with $\Delta t = 1$ ms and the noise constant is set to $\varsigma = 1$ evidence unit. The evidence at each time point is determined by the drift rate, the elapsed time from Glass pattern onset, and within trial noise. The resulting decision variable $x(t)$, progresses over time until one of two boundaries is reached (a or 0). The proportional start point parameter (w) is defined as the proportion of the distance between the two bounds (a) such that a start point of $w = 0.5$ is the midpoint between the upper (toIF) and lower (awayIF) bound. $w > 0.5$ indicates the start point is closer to the upper (toIF) bound, and $w < 0.5$ indicates the start point is closer to the lower (awayIF) bound.

There are at least six possible, non-mutually exclusive ways that unilateral SC inactivation could alter decision-making, which we show in simulation using the DDM (Fig. 5, Extended Data Fig. 5). The simulations yield different expectations for the psychometric and chronometric functions and RT distributions depending on which parameters change and how. Five hundred trials were simulated for each possible scenario to make the predictions of choice performance and RT distributions, and the first 20 trial simulations of the evidence paths are shown in Fig. 5 and Extended Data Fig. 5. All these simulations of performance and RT data can be used for visual comparison with the actual data obtained from monkey S and B shown in the shaded regions of Fig. 5q-x.

The first possibility is that muscimol impairs or alters the drift rates (δ) in a way that changes the proportionality factor between coherence and drift rate. Changes in the proportionality factor between coherence and drift rate change how the internal strength of the sensory information for orientation is converted into drift rate and therefore, the accumulation of evidence. The change in

the proportionality factor can be considered analogous to a change of sensitivity in SDT (assuming fixed noise, Fig. 4a,b, in addition to a symmetric change in boundary *a*, discussed below). We simulated a DDM (Fig. 5a-d, Extended Data Fig. 5a-d) with only a proportionality factor decrease achieved through a multiplicative decrease in drift rates of 0.5 (pre-muscimol δ_{1k} spanned -3 to 3 with steps of 0.5 away with mean $\Delta_1 = 0$ and post-muscimol δ_{2k} spanned -1.5 to 1.5 with steps of 0.25 away with mean $\Delta_2 = 0$) and no other parameter changes ($a = 1.4$, $\tau = .45$ sec, $w = 0.5$, $\lambda = 0$). Fig. 5a and Extended Data Fig. 5a show the representation of a sequential sampling of relative evidence towards one of two boundaries with 20 simulated evidence paths for a simulated pre-muscimol injection condition by faded grey lines and a simulated post-muscimol injection condition by faded orange lines in the 0% coherence condition. A decision is made once the relative path reaches one of two boundaries, then a saccade to the IF or saccade away from the IF occurs. Simulated RT distributions from the 0% coherence condition are shown on their respective boundaries for each choice (Fig. 5a, Extended Data Fig. 5a; shaded orange RT curves from the post-muscimol condition in front of shaded grey RT curves from the pre-muscimol condition). Drift rates are the mean rate of evidence accumulation within a trial and as such, are the mean rate of evidence accumulation for individual coherences and one direction (shorter arrows Fig. 5a and Extended Data Fig. 5a for the 36%, 10%, and 0% coherence conditions for to and away from the IF directions). The drift rate is not to be confused with the *mean drift rate across all trials, toIF and awayIF directions, and all coherences*, denoted by the longer arrows in Fig. 5a and Extended Data Fig. 5a. The mean drift rate across direction and coherence is also known as the drift rate offset, the distance from the drift criterion[1] and the drift bias[4]. Note that the drift rate offset is a parameter explicitly fit and estimated in the HDDM, but not in the non-hierarchical DDM and UGM, where we calculated the drift rate offset as the mean of the estimated drift rate parameters.

We then compared the change in drift rate offsets before and after muscimol (c.f., the longer black and orange lines in Fig. 5m and Extended Data Fig. 5m). As we see in Extended Data Fig. 5a-d, a decrease in the proportionality factor between coherence and drift rate (with the same drift rate offset of zero) would be evident as an increase in the RT distribution tails for both decisions to and away from the IF (Fig. 5a, Extended Data Fig. 5a), a change in the slope of the psychometric function (Fig. 5b, Extended Data Fig. 5b; dots indicate simulated data and curves show the fit to the simulated data, with black indicating pre-muscimol and orange indicating post-muscimol), and an increase in mean RTs for correct toIF and awayIF decisions (Fig. 5c,d, Extended Fig. 5c,d; black indicating pre- and orange indicating post-muscimol). However, we did not observe significant changes in the slope of the psychometric function (Fig. 2f, Fig. 5r,v, shaded) nor did we observe consistent mean RT increases for correct awayIF decisions in at least one monkey (monkey S; Fig. 5s, shaded), making the possibility of a decrease in the proportionality factor between coherence and drift rate unlikely to explain the post-muscimol data.

The second possibility is that muscimol alters the proportional starting point of the evidence accumulation (w) leading to a bias in choosing one or the other option (as a proportion of the boundary parameter, a). A change in starting point before evidence accumulation is one way to change the decision criterion in SDT, the second being a change in drift rate offset (see Fig. 4 of the main text). Drift rate offset and starting point changes can be differentiated in the DDM framework by comparing RT distributions. We simulated a DDM (Fig. 5e-h, Extended Data Fig. 5e-h) with only a decrease in start-point of the evidence accumulation path in proportion to the increasing boundary (pre-muscimol w_1 was 0.5 and post-muscimol w_2 was 0.25) with the same initial parameters and boundary height change as the other simulations. Note that absolute changes in starting points of evidence accumulation cannot be differentiated from asymmetric changes in

evidence boundaries to and away from the IF with behavioral data alone. However, changes in the proportional starting point (*w*) can be differentiated from symmetric changes in boundaries (*a*) (see parameter recovery in Supplementary Fig. 2j-m). As we see in Fig. 5e-h and Extended Data Fig. 5e-h of the simulation with a change in start point, the predictions show large changes in the leading edge of the awayIF RT distribution and shorter mean error toIF RTs, along with a decision bias. Looking at the choice performance prediction, a start point change may explain the effects of muscimol (Fig. 5r,v, shaded). However, when we look at the mean error toIF RT's for both monkeys (Fig. 5t,x, shaded), we do not observe a decrease, as suggested by the simulation, but rather an increase. We also observe a slight shortening of the leading edge of the toIF RT distribution for monkey S, but not for monkey B (Fig. 5q,u, shaded). A change in starting point could explain the post-muscimol choice performance, but is inconsistent with the changes observed in RT distributions.

A third possibility is that only the toIF boundary (*aU*) increases (i.e., more evidence is required for a toIF decision after muscimol in the SC) without an absolute start point *(z)* change (Fig. 5i-l, Extended Data Fig. 5i-l). Note that an increase in only one decision boundary (perhaps due to inactivation of only one SC) would be reflected as an increase in symmetric boundaries (*a*) and a decrease in proportional start point (*w*) away from the IF with no changes in drift rate offset in the HDDMs and DDMs we used. We simulated an upper boundary change with parameters (δ_1 spanned -3 to 3 with steps of 0.5 away with mean $\Delta = 0$, $\tau = 0.45$ sec, $z = 0.5$, $\lambda = 0$) with an increase in top boundary from $aU_1 = 0.7$ to $aU_2 = 1$ (which could also be reflected by a symmetric boundary increase of $a_1 = 1.4$ to $a_2 = 1.7$ and proportional start point decrease away from the IF $w_1 = .50$ to $w_2 = 0.41$) shown in Fig. 5i and Extended Data Fig. 5i. A single bound change also predicts a shift in the decision bias (Fig. 5j and Extended Data Fig. 5j), like those seen with a

proportional start point (*w*) change, a drift rate offset change (*Δ)*, and also in the actual data (see Fig. 5r,v, shaded). For the RT predictions, the single bound change predicts an increase in toIF RTs for correct and incorrect trials (Fig. 5k,l and Extended Data Fig. 5k,l), which matches the RT change we see in monkey S (Fig. 5s,t, shaded) and monkey B (Fig. 5w,x, shaded). However, the change in starting point cannot explain the magnitude of the psychometric function shift we observed in the data (Fig. 5r,v, shaded) with similar changes in simulated and observed mean RTs (Fig. 5s,t,w,x, shaded).

The fourth possibility is that muscimol alters or impairs the drift rate offset. We simulated a drift rate offset change in the DDM by adding an evidence independent constant to all the drift rates, which would be shown as an offset in all the drift rates and also the mean of the drift rates across toIF and awayIF directions and coherences. However, note that in our model fitting, we defined the drift rate offset as the mean of all the drift rates across toIF and awayIF directions and coherences. Although the drift rate offset parameter was explicitly fit in the HDDM, the drift rate offset in the DDM and UGM was calculated from the estimated drift rates per direction and coherence. Changing the drift rate offset is one way to change the decision criterion in SDT, and alterations in the drift rate offset would appear as decision biases (in addition to start point biases, discussed above; see Fig. 4 of the main text). We simulated a DDM (Fig. 5m-p, Extended Data Fig. 5m-p) with a change in drift rate offset (*Δ)*, achieved by a shift in drift rates of 1 evidence unit away from the IF (pre-muscimol δ_{1k} spanned -3 to 3 with steps of 0.5 away with mean $\Delta_1 = 0$ and post-muscimol δ_{2k} spanned -4 to 2 with steps of 0.5 away with mean $\Delta_2 = -1$) with the same initial parameters as the other simulations. A shift in the drift rate offset favoring awayIF decisions results in a laterally shifted psychometric function (Fig. 5n, Extended Data Fig. 5n). The resulting chronometric curve from a drift rate offset is distinct, similar to the one we see in the data for

monkey S (Fig. 5s, shaded). The drift rate offset change simulation also predicts an increase in correct toIF RTs (Fig. 5o,p, Extended Data Fig. 5o,p), which we observed in both monkeys (Fig. 5s,t,w,x, shaded). Correct awayIF RTs are expected to decrease; an observation made in monkey S but not monkey B. Overall, the drift rate offset favoring awayIF decisions explains the effects of muscimol in the SC on the choice performance and RT data across monkeys and experiments most consistently.

The fifth possibility is that muscimol alters or impairs decision thresholds symmetrically which control the amount of evidence required to reach a decision. We simulated only an increase in symmetric evidence boundaries (pre-muscimol a_1 was 1.4 and post-muscimol a_2 was 1.7) with the same initial parameters as the other simulations (Extended Data Fig. 5q-t). An increase in the symmetric boundary parameter (a) always results in an increase of RTs and increases in accuracy whereas a decrease in the symmetric boundary parameter (a) always results in a decrease of RTs and decreases in accuracy. Note that a change in sensitivity in SDT could be reflected by a change in symmetric boundary (a) or a change in the proportionality factor between coherence and drift rate. Comparing these symmetric boundary simulation predictions of choice performance and RT to the actual data (Fig. 5q-x, shaded), we see a lack of a decision bias prediction, but rather only a slight sensitivity change prediction (Extended Data Fig. 5r). Since we observed a large decision bias in the data, the possibility of a change in symmetric boundaries is unlikely to explain the effects of muscimol in the SC on the choice performance and RT data.

A sixth possibility is that muscimol inactivation of the SC alters the time for processes within a trial not related to evidence accumulation, ie., non-decision time (τ), such as visual encoding of the Glass pattern orientation or saccade execution. We simulated increased non-decision time (pre-muscimol τ_1 was 0.45 sec and post-muscimol τ_2 was 0.65) with the same initial parameters as the

other simulations (Extended Data Fig. 5u-x). A non-decision time change does not predict any changes in SDT because no choice behavior is affected, and therefore also is unlikely to explain the post-muscimol data.

Note that for the simulations we assumed that drift rates would change across coherence based on a simple linear function of coherence conditions of the Glass Pattern task, with the explicit ranges of drift rates given above. The relationship of drift rate changes in DDMs to SDT (Fig. 4) and the simulation results however, are not dependent upon this assumption. In fact, we show that the data from both monkeys are explained well by a HDDM with an explicit function that describes each drift rate (δ) as the comparison of two accumulators. In this HDDM model, the decrease in the gain (G) of one accumulator would change the mean drift rate, i.e. drift rate offset Δ, significantly (analogous to Fig. 4c,d, Fig. 5m-p, Extended Data Fig. 5m-p) with less (but some) impact on the mean change in proportionality factors across drift rates, whereas a decrease in the gain (G) of both accumulators by the same amount would decrease the proportionality factors on each drift rate by the same amount (analogous to Fig. 4a,b, Fig. 5a-d, Extended Data Fig. 5a-d) with no impact on the drift rate offset.

These different simulation predictions lay out ways in which SC inactivation can affect RT distributions and decision performance, which we can compare visually with the actual data in Fig. 5q-x shaded in grey. However, for a formal analysis comparing predictions of different models, we fit a hierarchical DDM (HDDM) with individual parameters free to vary to see what parameter change best explained the effects of muscimol in the SC.

Hierarchical drift-diffusion model of decision-making (HDDM)

We estimated parameters of hierarchical drift-diffusion models (HDDMs) fit to pre, post, and recovery data. Hierarchical DDMs often yield better estimates of session parameters due to "shrinkage" towards the mean parameters[5, 6]. This is in contrast to fitting a model per session of data which could lead to overfitting and misestimation[7, 8].

The hierarchical drift-diffusion model (HDDM) can be written as a series of statistical relationships between parameters and data. In particular, the HDDM can be written as a set of prior distributions of parameters, statistical parameter relationships, and likelihoods of observed data. Prior distributions provide initial uncertainty about parameters before data are observed. Here, we let only the data influence the shape of the posterior distributions by using wide (less informative) priors for the hierarchical parameters. Notation "~" refers to a probability distribution of parameters or data. The five probability distributions that were used in the hierarchical model are the Normal distribution (N) with mean and standard deviation parameters, the Truncated Normal distribution with truncation between c and d indicated by $\in(c, d)$, the Gamma distribution (Γ) with rate and shape parameters, the Uniform distribution (U) with boundary parameters, and the Wiener likelihood for the drift-diffusion model with parameters: non-decision time (τ), symmetric boundary (a), proportional start point (w), and drift rate (δ). Similar to other DDM model fitting procedures[2, 9, 10], a lapse process defined by the proportion of trials (λ) is found by assuming a mixture between the Wiener likelihood and a uniform distribution over all possible positive and negative RTs less than three seconds (where positive and negative RTs differentiate toIF and awayIF respectively using the Wiener module[11]). This lapse process thus captures a proportion of trials where the monkey did not complete the task and instead made a Bernoulli (i.e. "coin flip") decision either toIF or awayIF with a random RT between 0 and 3 seconds. The hierarchical mean parameters (μ) varied by experimental injection condition (e) and monkey (m), which affected

parameters that varied by experimental session (s) with some variance given by hierarchical variance parameters (σ). The drift rate offset parameter (Δ), in each experimental session was defined as the mean of the drift rates (δ) across positive and negative coherence values (k). The drift rate offset is a parameter that was explicitly fitted in the HDDM, in contrast to the non-hierarchical DDM and UGM where it was calculated. The overall hierarchical model (HDDM, Supplementary Fig. 1a) of choice and RT data (vector y) per trial (n) is defined by the following equations:

$$\tau_s \sim N\left(\mu_{(\tau)em}, \sigma^2_{(\tau)}\right) \in (0,1); \mu_{(\tau)em} \sim N(0.5, 0.25^2); \sigma_{(\tau)} \sim \Gamma(0.3,1) \quad \text{(S9)}$$

$$w_s \sim N\left(\mu_{(w)em}, \sigma^2_{(w)}\right) \in (0,1); \mu_{(w)em} \sim N(0.5,0.25^2); \sigma_{(w)} \sim \Gamma(0.3,1) \quad \text{(S10)}$$

$$\lambda_s \sim N\left(\mu_{(\lambda)em}, \sigma^2_{(\lambda)}\right) \in (0,1); \mu_{(\lambda)em} \sim N(0.3,0.15^2); \sigma_{(\lambda)} \sim \Gamma(0.3,1) \quad \text{(S11)}$$

$$a_s \sim N\left(\mu_{(a)em}, \sigma^2_{(a)}\right) \in (0,3); \mu_{(a)em} \sim N(1, 0.5^2); \sigma_{(a)} \sim \Gamma(1,1) \quad \text{(S12)}$$

$$\Delta_s \sim N\left(\mu_{(\Delta)em}, \sigma^2_{(\Delta)}\right); \mu_{(\Delta)em} \sim N(0, 2^2); \sigma_{(\Delta)} \sim \Gamma(1,1) \quad \text{(S13)}$$

$$\delta_{sk} \sim N\left(\Delta_s, \sigma^2_{(\delta)}\right) \in (-9,9); \sigma_{(\delta)} \sim \Gamma(1,1) \quad \text{(S14)}$$

$$y_n \sim (1-\lambda_s) DDM(\tau_s, w_s, a_s, \delta_{sk}) + \lambda_s U(-3,3) \quad \text{(S15)}$$

Note that we fit a simple HDDM without trial-to-trial variability in the model parameters, specifically trial-to-trial variability in non-decision time, drift rate, and initial bias. We assume that some across-trial variability in these parameters exist in the data due to there being intrinsic variability in the stimulus display and the monkeys' performance. However simple DDM models often estimate the mean DDM parameters well even in the presence of across-trial variability[12].

We therefore simplified the model fitting procedure by excluding across-trial variability parameters, which are often difficult to estimate[13].

We tested the prior sensitivity of the HDDM by changing the prior distributions of the hierarchical mean parameters (μ) and keeping the other equations the same. We observed no changes greater than a few percentage points in posterior distributions (see Supplementary Table 5), although changes in BFs did occur as we previously discussed in the Methods. We fitted an additional 3 models; a "Shifted" HDDM with shifted hierarchical mean priors from the original priors above, $\mu_{(\tau)em} \sim N(1.0, 0.25^2)$, $\mu_{(w)em} \sim N(0.25, 0.25^2)$, $\mu_{(\lambda)em} \sim N(0.4, 0.15^2)$, $\mu_{(a)em} \sim N(2, 0.5^2)$, and $\mu_{(\Delta)em} \sim N(1, 2^2)$; a "Narrow" HDDM with half the width hierarchical mean priors compared to the original priors above, $\mu_{(\tau)em} \sim N(0.5, 0.125^2)$, $\mu_{(w)em} \sim N(0.5, 0.125^2)$, $\mu_{(\lambda)em} \sim N(0.3, 0.075^2)$, $\mu_{(a)em} \sim N(1, 0.25^2)$, and $\mu_{(\Delta)em} \sim N(0, 1^2)$; a "Wide" HDDM with double the width hierarchical mean priors compared to the original priors above, $\mu_{(\tau)em} \sim N(0.5, 0.5^2)$, $\mu_{(w)em} \sim N(0.5, 0.5^2)$, $\mu_{(\lambda)em} \sim N(0.3, 0.3^2)$, $\mu_{(a)em} \sim N(1, 1.0^2)$, and $\mu_{(\Delta)em} \sim N(0, 4^2)$.

Hierarchical parameter estimates of pre, post, and recovery data from the HDDM appear in Extended Data Fig. 6a-j. Also presented are 2.5th and 97.5th percentiles of the posterior distributions of hierarchical parameters, which provide 95% credible intervals. These credible intervals provide 95% certainty about the parameter estimates given the likelihood of the data and prior distributions. For both monkeys, we see a significant change (defined as greater than 95% probability estimated by the posterior distributions) in the drift rate offset towards the awayIF decisions in the post-muscimol data after unilateral inactivation of the SC (Extended Data Fig. 6a,b). We observed a significant decrease in start point for monkey S (Extended Data Fig. 6c) but not for monkey B (Extended Data Fig. 6d). Increases in non-decision time and symmetric bound

for monkey B were also observed (Extended Data Fig. 6f,h). The changes in non-decision time (τ) and symmetric bound (a) for monkey S were not very probable (the 95% credible interval of post-data includes the estimate of the pre- data, Extended Data Fig. 6e,g). We did not observe any significant changes in the lapse proportion (λ) (Extended Data Fig. 6i,j) for either monkey.

The graphical results of the full HDDM (Fig. 5q,u, shaded) were generated by estimating the median posterior samples of each hierarchical parameter (μ), the parameter estimates that are shown in Extended Data Fig. 6a-j. The graphical DDM for monkey S (Fig. 5q, shaded) was generated with median posteriors of -0.06 evidence units per second drift rate offset (μ_v) pre-muscimol, -0.64 drift rate offset (μ_v) post-muscimol, 0.54 proportion of decision evidence as the start-point (μ_w) pre-muscimol, 0.49 proportion of decision evidence as the start-point (μ_w) post-muscimol, 1.50 evidence units boundary (μ_a) pre-muscimol, 1.55 evidence units boundary (μ_a) post-muscimol, 408 ms non-decision time (μ_τ) pre-muscimol, 433 ms non-decision time (μ_τ) post-muscimol. The median posteriors of 3% lapse rate (μ_λ) pre-muscimol and 4% lapse rate (μ_λ) post-muscimol for monkey S are not shown in the graphical results (Fig. 5q, shaded). The graphical DDM for monkey B was generated with mean posteriors of 0.09 drift rate offset (μ_v) pre-muscimol, -0.85 drift rate offset (μ_v) post-muscimol, 0.52 proportion of decision evidence as the start-point (μ_w) pre-muscimol, 0.55 proportion of decision evidence as the start-point (μ_w) post-muscimol, 1.33 evidence units boundary (μ_a) pre-muscimol, 1.55 evidence units boundary (μ_a) post-muscimol, 543 ms non-decision time (μ_τ) pre-muscimol, 597 ms non-decision time (μ_τ) post-muscimol. The median posteriors of 5% lapse rate (μ_λ) pre-muscimol and 5% lapse rate (μ_λ) post-muscimol for monkey B are not shown in the graphical results (Fig. 5u, shaded). We also estimated the median posterior samples of drift rates for each monkey in each session by combining the posterior distributions across sessions. The estimates were for the -36%, -10%, 0%, 10%, and 36%

were used to plot drift rate estimates as arrows in the graphical representation (Fig. 5q,u, shaded) and were the following: (-1.82, -0.86, -0.16, 0.71, 1.96) pre-muscimol in monkey S, (-1.89, -1.29, -0.90, -0.02, 1.3) post-muscimol in monkey S, (-2.74, -1.13, 0.20, 1.71, 3.29) pre-muscimol in monkey B, and (-3.20, -1.87, -0.87, 0.02, 2.16) post-muscimol in monkey B.

Given that there are multiple parameter changes from pre-muscimol to post-muscimol, we tested which parameter best described the observed behavioral changes in choice performance and RT distributions after muscimol injection in the SC by fitting the HDDMs that fixed all parameters across injection conditions besides the parameter of interest. This resulted in four additional variants of the hierarchical DDM that were fitted to the data: HDDM with the drift rate offset free to vary (HDDM-Δ), HDDM with the start point free to vary (HDDM-w), HDDM with non-decision time free to vary (HDDM-τ), and HDDM with the start point and bound free to vary (HDDM-a,w). In each variant, hierarchical parameters (μ), besides hierarchical parameters of the variable of interest (Δ or w or τ), only varied by monkey (m) and *not* by injection condition (e). We modeled the HDDM with the free-to-vary drift rate offset by letting all individual drift rates per k condition and the drift rate offset parameter free to vary. With this model, we assessed whether the drift rate offset changes could explain the post-muscimol data. Hierarchical parameters of the variable of interest could vary by both monkey (m) and injection condition (e). Sessions (s) were collapsed across the pre- and post-muscimol and recovery sessions for each separate injection into a parameter "day" (d), where one "day" collapsed across three sessions of data that constituted two sessions (a pre- session occurring immediately preceding a post-muscimol or saline session) and a recovery session within ~24 hours. Parameters besides the variable of interest could vary by day (d), while the parameter of interest varied by session (s). Thus, each parameter besides the parameter of interest could only explain variance in the data due to the monkey (B or S) and across

unique “days”. Only the parameter of interest could explain variance in the data due to the effect of an injection changing across injection types (*e*) and separate sessions (*s*). Similar non-informative prior distributions were used as in the full HDDM. The overall hierarchical model with the drift rate offset (Δ) as the variable of interest (HDDM-Δ, Supplementary Fig. 1b) of choice and RT data (vector *y*) per trial (*n*) is defined by the following equations:

$$\tau_d \sim N\left(\mu_{(\tau)m}, \sigma^2_{(\tau)}\right) \in (0,1); \mu_{(\tau)m} \sim N(0.5, 0.25^2); \sigma_{(\tau)} \sim \Gamma(0.3,1) \quad \text{(S16)}$$

$$w_d \sim N\left(\mu_{(w)m}, \sigma^2_{(w)}\right) \in (0,1); \mu_{(w)m} \sim N(0.5, 0.25^2); \sigma_{(w)} \sim \Gamma(0.3,1) \quad \text{(S17)}$$

$$\lambda_d \sim N\left(\mu_{(\lambda)m}, \sigma^2_{(\lambda)}\right) \in (0,1); \mu_{(\lambda)m} \sim N(0.3, 0.15^2); \sigma_{(\lambda)} \sim \Gamma(0.3,1) \quad \text{(S18)}$$

$$a_d \sim N\left(\mu_{(a)m}, \sigma^2_{(a)}\right) \in (0,3); \mu_{(a)m} \sim N(1, 0.5^2); \sigma_{(a)} \sim \Gamma(1,1) \quad \text{(S19)}$$

$$\Delta_s \sim N\left(\mu_{(\Delta)em}, \sigma^2_{(\Delta)}\right); \mu_{(\Delta)em} \sim N(0, 2^2); \sigma_{(\Delta)} \sim \Gamma(1,1) \quad \text{(S20)}$$

$$\delta_{sk} \sim N\left(\Delta_s, \sigma^2_{(\delta)}\right) \in (-9,9); \sigma_{(\delta)} \sim \Gamma(1,1) \quad \text{(S21)}$$

$$y_n \sim (1-\lambda_d) DDM(\tau_d, w_d, a_d, \delta_{sk}) + \lambda_d U(-3,3) \quad \text{(S22)}$$

As in the full HDDM, in a hierarchical model with start point (*w*) as the variable of interest (HDDM-*w*, Supplementary Fig. 1c), the start point of evidence accumulation (*w*) was assumed to be fixed across coherence values before evidence accumulation in a trial. The model variant is defined by the following equations:

$$\tau_d \sim N\left(\mu_{(\tau)m}, \sigma^2_{(\tau)}\right) \in (0,1); \mu_{(\tau)m} \sim N(0.5, 0.25^2); \sigma_{(\tau)} \sim \Gamma(0.3,1) \quad \text{(S23)}$$

$$w_s \sim N\left(\mu_{(w)em}, \sigma^2_{(w)}\right) \in (0,1); \mu_{(w)em} \sim N(0.5, 0.25^2); \sigma_{(w)} \sim \Gamma(0.3,1) \quad \text{(S24)}$$

$$\lambda_d \sim N\left(\mu_{(\lambda)m}, \sigma^2_{(\lambda)}\right) \in (0,1); \mu_{(\lambda)m} \sim N(0.3, 0.15^2); \sigma_{(\lambda)} \sim \Gamma(0.3,1) \quad \text{(S25)}$$

$$a_d \sim N\left(\mu_{(a)m}, \sigma^2_{(a)}\right) \in (0,3); \mu_{(a)m} \sim N(1,0.5^2); \sigma_{(a)} \sim \Gamma(1,1) \quad \text{(S26)}$$

$$\Delta_d \sim N\left(\mu_{(\Delta)m}, \sigma^2_{(\Delta)}\right); \mu_{(\Delta)m} \sim N(0,2^2); \sigma_{(\Delta)} \sim \Gamma(1,1) \quad \text{(S27)}$$

$$\delta_{dk} \sim N\left(\Delta_d, \sigma^2_{(\delta)}\right) \in (-9,9); \sigma_{(\delta)} \sim \Gamma(1,1) \quad \text{(S28)}$$

$$y_n \sim (1-\lambda_d) DDM(\tau_d, w_s, a_d, \delta_{dk}) + \lambda_d U(-3,3) \quad \text{(S29)}$$

Unlike in the full HDDM, non-decision time (ω) was allowed to be variable across coherence values (k) with a hierarchical non-decision time (τ) as the variable of interest (HDDM-τ, Supplementary Fig. 1d). The model variant is defined by the following equations:

$$\tau_s \sim N\left(\mu_{(\tau)em}, \sigma^2_{(\tau)}\right); \mu_{(\tau)em} \sim N(0.5,0.25^2); \sigma_{(\tau)} \sim \Gamma(0.3,1) \quad \text{(S30)}$$

$$w_d \sim N\left(\mu_{(w)m}, \sigma^2_{(w)}\right) \in (0,1); \mu_{(w)m} \sim N(0.5, 0.25^2); \sigma_{(w)} \sim \Gamma(0.3,1) \quad \text{(S31)}$$

$$\lambda_d \sim N\left(\mu_{(\lambda)m}, \sigma^2_{(\lambda)}\right) \in (0,1); \mu_{(\lambda)m} \sim N(0.3, 0.15^2); \sigma_{(\lambda)} \sim \Gamma(0.3,1) \quad \text{(S32)}$$

$$a_d \sim N\left(\mu_{(a)m}, \sigma^2_{(a)}\right) \in (0,3); \mu_{(a)m} \sim N(1, 0.5^2); \sigma_{(a)} \sim \Gamma(1,1) \quad \text{(S33)}$$

$$\Delta_d \sim N\left(\mu_{(\Delta)m}, \sigma^2_{(\Delta)}\right); \mu_{(\Delta)m} \sim N(0,2^2); \sigma_{(\Delta)} \sim \Gamma(1,1) \quad \text{(S34)}$$

$$\delta_{dk} \sim N\left(\Delta_d, \sigma^2_{(\delta)}\right) \in (-9,9); \sigma_{(\delta)} \sim \Gamma(1,1) \quad \text{(S35)}$$

$$\omega_{sk} \sim N\left(\tau_s, \sigma^2_{(\omega)}\right) \in (0,1); \sigma_{(\omega)} \sim \Gamma(0.3,1) \quad \text{(S36)}$$

$$y_n \sim (1-\lambda_d) DDM(\omega_{sk}, w_d, a_d, \delta_{dk}) + \lambda_d U(-3,3) \quad \text{(S37)}$$

To answer the question whether *any* boundary change (including a symmetric boundary or either single boundary change) could explain the effects of unilateral inactivation of the SC with muscimol, we fit a model variant with only the symmetric boundary (a) and start point (w) free to

vary across experimental conditions (e) and sessions (s). This final model variant is defined by the following priors and likelihood equations:

$$\tau_d \sim N\left(\mu_{(\tau)m}, \sigma^2_{(\tau)}\right) \in (0,1); \mu_{(\tau)m} \sim N(0.5, 0.25^2); \sigma_{(\tau)} \sim \Gamma(0.3,1) \quad \text{(S38)}$$

$$w_s \sim N\left(\mu_{(w)em}, \sigma^2_{(w)}\right) \in (0,1); \mu_{(w)em} \sim N(0.5, 0.25^2); \sigma_{(w)} \sim \Gamma(0.3,1) \quad \text{(S39)}$$

$$\lambda_d \sim N\left(\mu_{(\lambda)m}, \sigma^2_{(\lambda)}\right) \in (0,1); \mu_{(\lambda)m} \sim N(0.3, 0.15^2); \sigma_{(\lambda)} \sim \Gamma(0.3,1) \quad \text{(S40)}$$

$$a_s \sim N\left(\mu_{(a)em}, \sigma^2_{(a)}\right) \in (0,3); \mu_{(a)em} \sim N(1, 0.5^2); \sigma_{(a)} \sim \Gamma(1,1) \quad \text{(S41)}$$

$$\Delta_d \sim N\left(\mu_{(\Delta)m}, \sigma^2_{(\Delta)}\right); \mu_{(\Delta)m} \sim N(0, 2^2); \sigma_{(\Delta)} \sim \Gamma(1,1) \quad \text{(S42)}$$

$$\delta_{dk} \sim N\left(\Delta_d, \sigma^2_{(\delta)}\right) \in (-9,9); \sigma_{(\delta)} \sim \Gamma(1,1) \quad \text{(S43)}$$

$$y_n \sim (1 - \lambda_d) DDM(\tau_d, w_s, a_s, \delta_{dk}) + \lambda_d U(-3,3) \quad \text{(S44)}$$

When comparing the in-sample prediction of all five model variants, the full HDDM explains the change in decision bias the best in both monkeys (98.3% of the data for monkey S, 99.3% for monkey B) as shown in the predicted psychometric functions (Extended Data Fig. 7; Supplementary Table 4). The HDDM-Δ captures the change in decision bias almost as well as the full HDDM in both monkeys (97.6% for monkey S, 98.3% for monkey B), followed by HDDM-a,w (94.6% for monkey S, 94.0% for monkey B), then HDDM-w (94.5% for monkey S, 93.5% for monkey B), and then HDDM-τ (89.9% for monkey S, 91.3% for monkey B). The full HDDM also explains the RT distribution best for both monkeys (Supplementary Table 4) as seen by the visual comparisons of the predictions of the HDDM and the actual data (Extended Data Fig. 7). While HDDM-τ fails to capture any decision bias from pre to post data, it does capture RT data almost as well as the full model (Extended Data Fig. 7; Supplementary Table 4). Overall, the full HDDM

best predicts the choice performance and the RT data and the results of this model are discussed in the Main text.

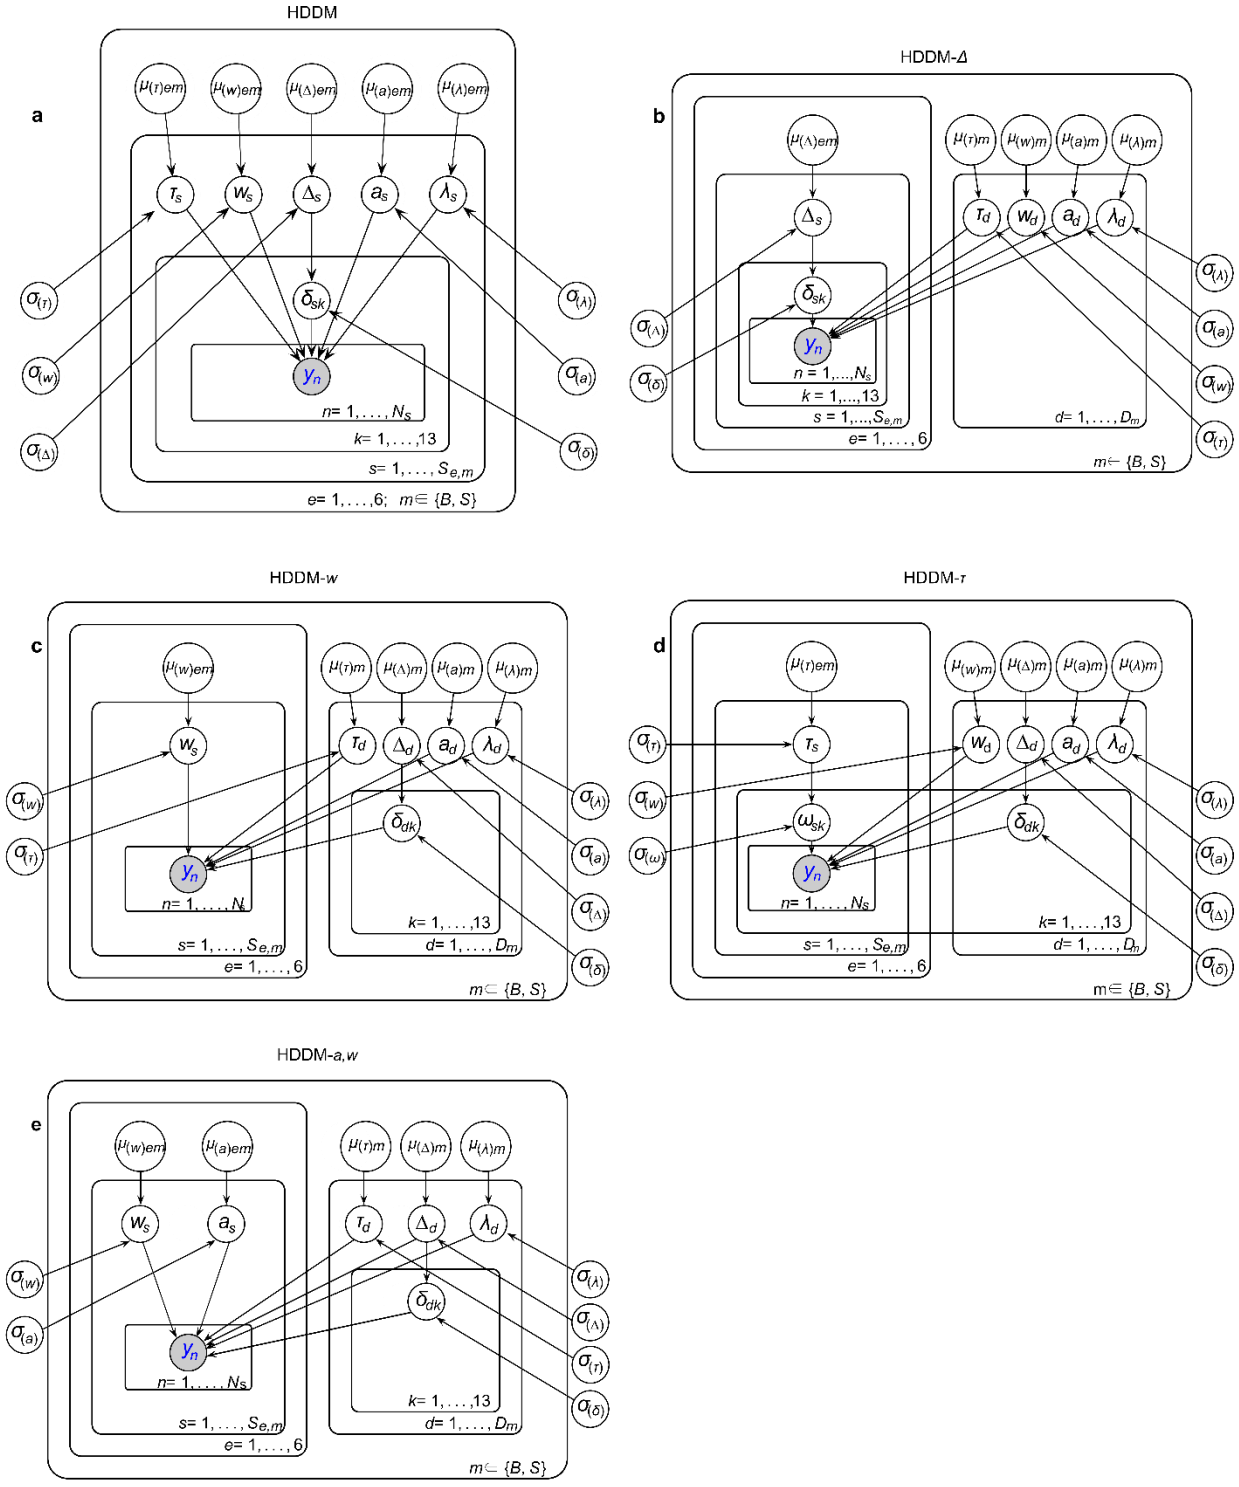

Supplementary Fig. 1. Graphical model descriptions. (Associated with Fig. 5 of the main text) Hierarchical models represented by the graph structure of **a** full HDDM **b** HDDM-Δ **c** HDDM-w **d** HDDM-τ **e** HDDM-a,w, given by the prior and likelihood equations in the Supplementary Note (S9-44). Arrows represent dependencies of variables on hierarchical variables, including hierarchical means (μ) and hierarchical variance parameters (σ). Hierarchical means vary by either both monkey (m) and injection condition (e) or just monkey (m). Other variables vary by session (s), day (d), coherence condition (k). The observed joint decision-RT data ($\boldsymbol{y}$) per trial (n) is described by a Wiener process depending upon variables in the lowest level of the hierarchy.

To better understand the mechanism of unilateral inactivation of the SC on evidence accumulation, we fit one final HDDM variant with the same parameters and priors as the full HDDM but with an enforced equation that describes each drift rate (δ) per experimental session (s) and coherence (k) as the linear comparison of two accumulators. One can interpret this model as the addition of a gain element that impacts how evidence is accumulated in one of two accumulators, i.e., a biologically plausible implementation of a change in drift rate offset[14, 15]. To achieve model identifiability, we assumed that the injected SC accumulated a probability for an IF decision with probability (θ) that varied with coherence (k) with the other SC accumulating with probability (1 - θ). The injected SC was also influenced by gain effects ($G_{(toIF)}$) and the unaffected SC was influenced by gain effects ($G_{(awayIF)}$) that both varied by injection condition (e). For model identifiability, we could only gauge the relative influence of injection conditions on the gain effects, and thus $G_{(awayIF)}$ was set to 1 in each injection condition and $G_{(toIF)}$ was set to 1 in the pre-muscimol injection condition. We explored the amount of gain on toIF decisions ($G_{(toIF)}$) in the other five experimental conditions: post-muscimol (e=2), recovery from muscimol (e=3), pre-saline (e=4), post-saline (e=5), and recovery from saline (e=6) with Bayes Factors (see Methods). Finally, the linear accumulator comparison was scaled by a parameter in seconds (R) that varied

per monkey (*m*). Thus, this final model variant (HDDM-*G*) is fully described by equations S9 through S12 as well as the following equations:

$$\theta_k \sim N(0.5, 0.25^2) \in (0,1) \tag{S45}$$

$$G_{(awayIF)e} = 1; G_{(toIF)1} = 1; G_{(toIF)e} \sim N(1, 3^2) \forall e \in \{2,3,4,5,6\} \tag{S46}$$

$$R_m \sim N(0.5, 0.25^2) \in (0.1,4) \tag{S47}$$

$$\delta_{sk} = \frac{1}{R_m}\left(\theta_k G_{(toIF)e} - (1-\theta_k) G_{(awayIF)e}\right) \tag{S48}$$

$$y_n \sim (1-\lambda_s) DDM(\tau_s, w_s, a_s, \delta_{sk}) + \lambda_s U(-3,3) \tag{S49}$$

All HDDMs were fitted for 52,000 original samples in each of six Markov chains for each parameter. After removing the first 2000 samples as a “warm-up" and then keeping only every 10th sample, i.e., using a “thinning” parameter of 10, this resulted in 5,000 posterior samples in each Markov chain for 5,000 * 6 = 30,000 samples from the estimated posterior distributions for each parameter. To assess model convergence, the Gelman-Rubin statistic and the number of effective samples were calculated[7]. The Gelman-Rubin statistic assesses the convergence of MCMC samplers by comparing the between-chain variance to the within-chain variance of each parameter, with Gelman-Rubin statistic > 1.1 thought to be a necessity for convergence. The “effective number of samples” equation scales the total sample number for each parameter posterior by autocorrelation in the chains in order to estimate an independent number of samples[7]. We also implemented the recommendation by [7] (see footnote in the 3rd Edition on page 283) to split the chains in half before calculating the Gelman-Rubin statistic in order to account for non-stationary chains. Larger effective numbers of samples for each parameter in the model are better. The chains for parameters with the largest Gelman-Rubin statistics and smallest effective number

of samples were also visually inspected to ensure convergence. Model HDDM converged with Gelman-Rubin statistics of <= 1.03 and a number of effective samples >=222 for all parameters. Model HDDM-Δ converged with Gelman-Rubin statistics of <= 1.01 and a number of effective samples of >=882 for all parameters, and model HDDM-*w* converged with Gelman-Rubin statistics of <= 1.03 and a number of effective samples >=31 for all parameters. Model HDDM-τ had some non-decision time variables reach multiple-peaked posterior distributions, and so did not converge with Gelman-Rubin statistics of <= 1.12 and a number of effective samples >=37 for all parameters. The model HDDM-*a*,*w* converged with Gelman-Rubin statistics of <= 1.06 and a number of effective samples >=22 for all parameters. Model HDDM-*G* converged with Gelman-Rubin statistics of <= 1.02 and a number of effective samples >=337 for all parameters. The goodness of fit measure, R^2_{pred} for these sets of parameter estimates for HDDM to the data are shown in Supplementary Table 4.

Non-hierarchical drift-diffusion model of decision-making (DDM)

We also fitted the non-hierarchical DDM, to the pre-muscimol and post-muscimol data using quantile maximum products estimation (QMPE)[16] in the ChaRTr package[17], along with the HDDM to ensure that our parameter estimations were robust against different modeling methods. Model fitting of the non-hierarchical DDM was performed on the RT task data from monkey S and monkey B (seven for monkey S and two for monkey B), separately for each monkey for pre- and post-muscimol data, and pooled across sessions. This model fit to the data with nine RT quantiles that summarize the RT distribution of each of the 11 trial conditions (24%, 17% 10%, 5%, 3%, 0%) for toIF and awayIF decisions with the 0% collapsed for to and awayIF decisions. The 36% conditions in the toIF and awayIF data were excluded for both monkeys because there were not enough data for fitting error RT quantiles. For all the fits using the models in the ChaRTr package,

five iterations were fitted for each dataset with different random seeds to avoid getting stuck in local minima when optimizing the goodness of fit measure, the QMP statistic. The best out of the five fits (according to AIC/BIC scores calculated from the QMP statistic) was used for model comparisons and reported for parameter estimation (Supplementary Table 4 and Extended Data Fig. 6).

The main differences between the HDDM and the non-hierarchical DDM are that 1) they employ a different modeling method (QMPE for non-hierarchical DDM and Bayesian Estimation for HDDM), 2) the HDDM has a hierarchical structure that allows for the parameters of individual experimental sessions to be pushed towards the mean parameters of all the sessions in an effort to better estimate parameters, 3) the inclusion of a lapse rate in the HDDM and 4) the drift rate offset was a parameter explicitly fit in the HDDM whereas in the DDM, the drift rate offset was calculated from fitted individual drift rate parameters. However, for model fitting, the definition of the drift rate offset, the mean of all the drift rates across directions and coherences, is the same in the HDDM, DDM, and UGM. For our non-hierarchical DDM, the drift rate for each condition, start point (w), non-decision time (τ), and the symmetric boundary (a) were free parameters that were estimated for pre- and post-muscimol data.

The parameter estimates of the DDM from fitting pre and post sessions showed the same relative changes from pre- to post-muscimol in all the parameters for both monkeys that we also saw using the HDDM parameter estimates, including the same change in the drift rate offset (Extended Data Fig. 6k,l), indicating that our parameter estimation results are robust against different modeling methods.

Urgency-gating models (UGMs)

Urgency-gating Models (UGMs) are another class of accumulation to bound model where the sensory evidence is 1) low-pass filtered to prioritize more recent evidence and 2) multiplied by a linearly growing urgency signal during a decision[18, 19]. We found evidence that the UGM may also reasonably explain the behavior of monkey B in the pre-muscimol sessions, whereas data from monkey S were better explained by DDMs (Supplementary Table 4), although it is impossible to distinguish DDMs from UGMs using tasks with constant evidence during a trial, such as the Glass pattern decision task[20]. Nevertheless, under the assumption that the monkeys' behavior could be explained by the UGM, and given our finding that a change in the drift rate offset (Δ) best explained the effect of the muscimol when assuming a DDM (Supplementary Table 4), we tested whether unilateral SC inactivation could be explained by a change in drift rate offset, or a change in the urgency signal of a UGM. Changes in the drift rate offset parameter signals a shift in the drift rates that are directly involved in the computation of evidence, biasing the evidence accumulation to one direction. In contrast, changes in the urgency signal affect a time-varying gain factor that influences decisions after the evidence is computed, rather than being involved in the computation of evidence itself. To assess whether the effects of SC inactivation were better explained by a change in the drift rate offset or an urgency signal, we fitted UGM variants to the data from monkey S and monkey B by letting either the drift rate offset or the urgency signal free to vary in post-muscimol data while keeping other parameters fixed to their pre-muscimol parameter estimate values. The UGM variant with the drift rate offset free to vary was modeled by letting all individual drift rates per coherence and direction free to vary, therefore allowing drift rate offset to change.

This UGM process that we used for our model is defined by the following equations modified from the ChaRTr package:

$$\alpha = T/(T + \Delta t) \quad \text{(S50)}$$

$$E(t) = \alpha\, E(t\text{-}1) + (1\text{-}\alpha)\,(\delta_k\, \Delta t + \varsigma\, N(0,1)) \quad \text{(S51)}$$

$$u(t) = i + mt \quad \text{(S52)}$$

$$x(t) = E(t)u(t) \quad \text{(S53)}$$

$$x(t=0) = w*a \quad \text{(S54)}$$

$$x(t) \geq a \rightarrow choose\ toIF \quad \text{(S55)}$$

$$x(t) \leq 0 \rightarrow choose\ awayIF \quad \text{(S56)}$$

Evidence during a trial, $E(t)$, is low pass filtered with a weight for previous evidence (α) that is controlled by the filter constant set to T = 100 ms and a weight for incoming evidence (1-α) that is influenced by the drift rate (δ) for each coherence condition (k), the time step (Δt = 1 ms), and the noise constant (ς), set to 100 evidence units. Evidence is calculated and influenced by the drift rate, elapsed time since the Glass pattern onset, and the within-trial noise. Evidence in the UGM is low-pass filtered to prioritize more recent evidence, similar to a leaky integrator. Also in the UGM, unlike the DDM, the current state of evidence, $E(t)$, is multiplied by the linear urgency signal $u(t)$ with a slope (m) and an intercept (i), fixed to 0. This linear urgency signal is a time-varying gain parameter that is multiplied by the evidence. In the UGM, ramping activity of decision-making neurons is driven by the urgency signal rather than evidence accumulation as in DDM.

The slope of the linearly growing urgency signal (m) in the UGM was free to vary to assess whether the effects of muscimol were best explained by a change in the urgency signal. The symmetric

boundary (a) was fixed to estimate the urgency slope (m) as a free parameter because parameter recovery results of the UGM with both a free boundary and a free urgency slope suggested that there was an overlapping role of the bound height and the urgency slope in explaining the choice and reaction time data (Supplementary Fig. 2a-e). Because we were only interested in assessing whether the effects of muscimol in the SC were best explained by a change in the urgency signal or the drift rate offset, we fixed the symmetric boundary to accurately estimate the change in urgency slope from pre- to post-muscimol. For fitting the UGM with both the pre- and post-muscimol data, the boundary value was fixed at the value of the parameter estimate of the boundary from fitting the UGM with the only the boundary (a), drift rates (δ_k), non-decision time (τ), and start point (w) allowed to vary in the pre-muscimol data. Therefore, in our full UGM, the drift rate for each condition (δ_k), non-decision time (τ), start point (w), and the slope of the urgency signal (m) were free parameters that were estimated for pre- and post-muscimol data for both monkeys.

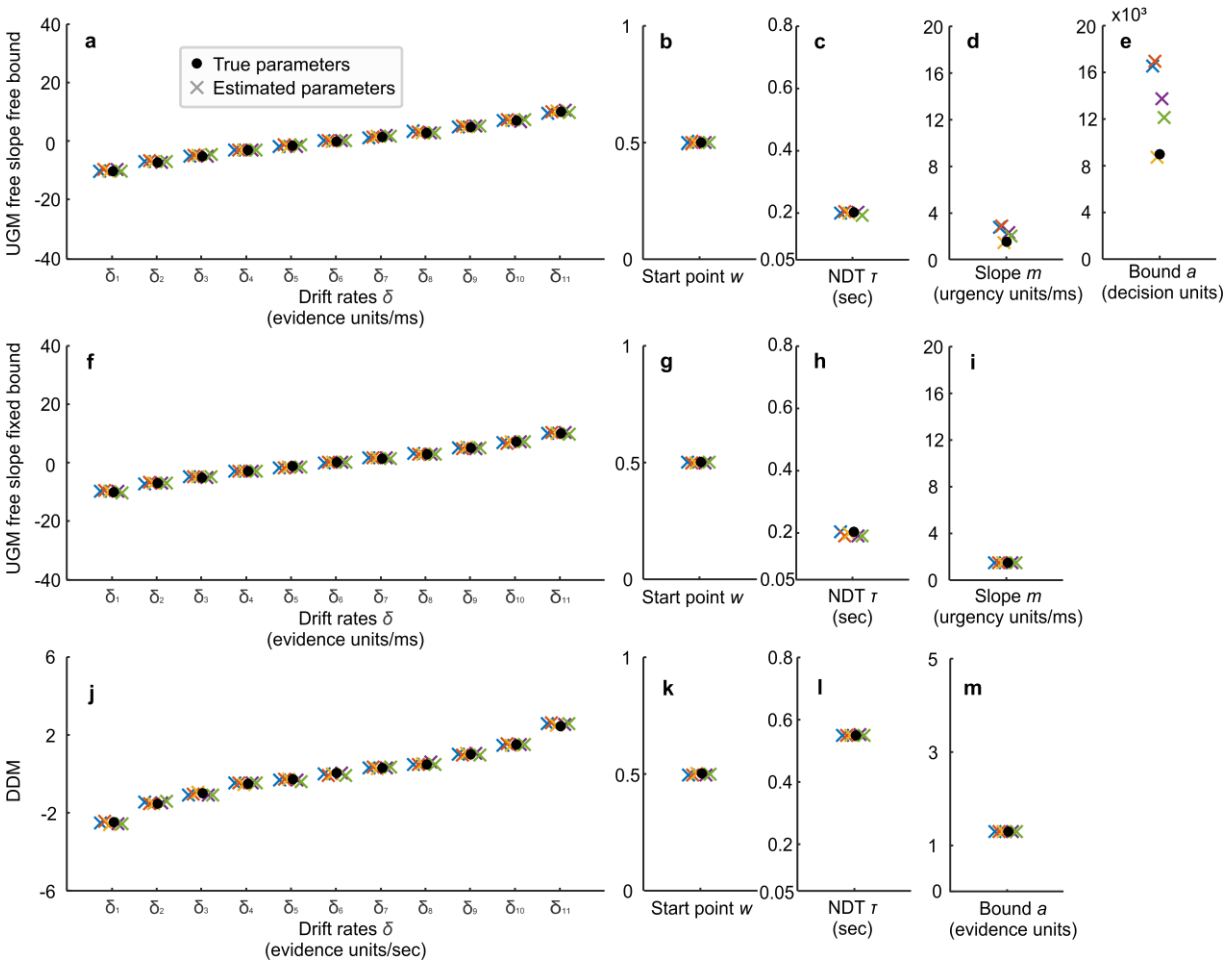

Supplementary Fig. 2. Parameter Recovery for UGM and DDM The set of parameter estimates for each of the five iterations of fits used for parameter recovery (denoted as x) are shown along with the known parameters used to simulate the data (denoted as black circles). The y-axis scale for each parameter was determined by the lower restraint and the upper restraint when estimating the parameter during fitting. **a-e** The parameter recovery estimates for the UGM with free proportional start point, free drift rates, free non-decision time (NDT), free urgency slope, and a free bound. Visually, the parameter recovery estimates of all 5 iterations of the drift rate, proportional start point, and non-decision time match the true parameter estimate. However, parameter estimates from four out of the five fits of the urgency slope and the bound are overestimated. The magnitude of the overestimation of the urgency slope and the bound even seem correlated, where the fits with a higher slope also have a higher bound. Because we were unable to recover the slope reliably when the bound was free and because we wanted to assess whether a change in urgency slope best explained the effect of muscimol in the SC, we used a version of the UGM where the urgency slope, drift rate, proportional start point, and non-decision time were free, but the bound was fixed. **f-i** To confirm that the UGM

with the free urgency slope but fixed bound (free proportional start point, free drift rates, free non-decision time, free urgency slope, fixed bound) was able to estimate parameters accurately, we show the parameter recovery estimates for this UGM variant. All parameter recovery estimates of all five fits recover the true parameter for this model. We defined this as our "full UGM" and used for modeling the actual data to recover slope estimates. **j-m** We also made sure that the non-hierarchical DDM could recover parameters accurately with free proportional start point, free drift rate, free bound, and free non-decision time. All parameters from all of the five fits match the true parameters.

Fitting the UGM for parameter estimation of pre- and post-muscimol data and also fitting UGM variants with only either the drift rates or the urgency slope free, we used the same pre- and post-muscimol data pooled across sessions and modeling method described in the non-hierarchical DDM section, using the ChaRTr package[17]. To test the robustness of this procedure of fitting pooled data, we also fitted all the pre-muscimol sessions individually across experimental sessions, with only 9 conditions: 17%, 10%, 5%, 3% each for toIF and awayIF, and 0% condition collapsed across toIF and awayIF. The 24% and 36% conditions were excluded because there were not enough error trials in each individual session for either monkey. We also used the results of fitting individual pre-muscimol session data, for model comparison between the HDDM, DDM, and UGM for each monkey (model fits denoted with ** in Supplementary Table 4). The parameter estimates for fitting the DDM and the UGM using the pooled and individual data were similar (Supplementary Fig. 3).

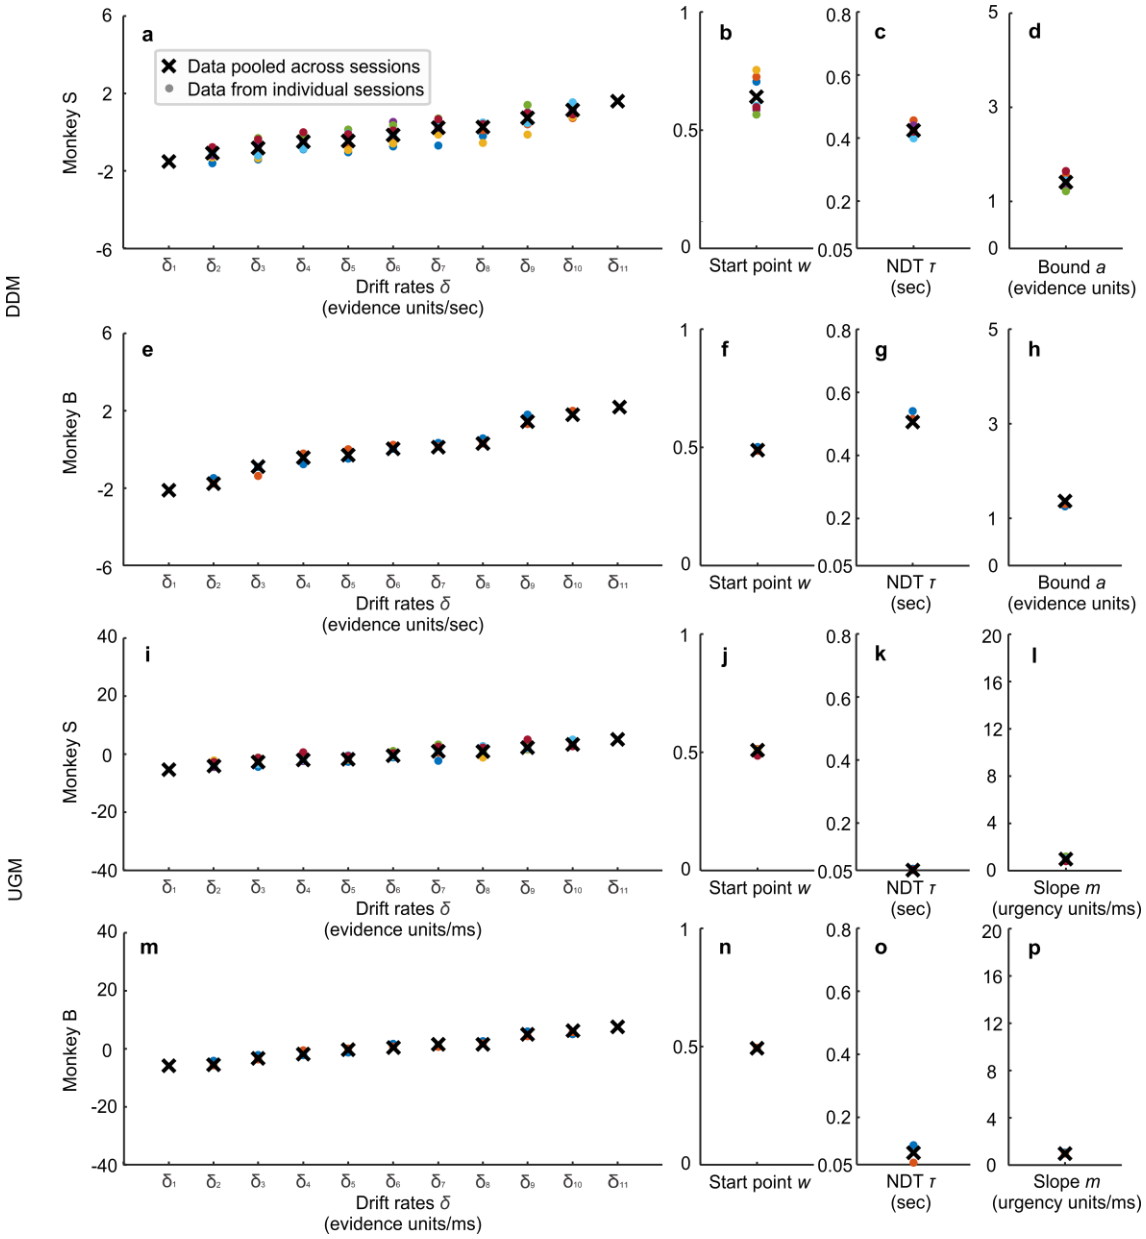

Supplementary Fig. 3. Parameter estimates of fits using pre-muscimol data from individual data versus pooled data a-d DDM parameter estimates fitted to pre-muscimol data that were pooled across all muscimol injections (black x) are shown in comparison to the DDM parameter estimates fitted to the pre-muscimol data from individual muscimol injection sessions for monkey S (n=9 injections, circles). **e-h** Same as in a-d for monkey B (n=2 injections, circles). **i-l** Same as in a-d for the UGM for monkey S (n=9 injections, circles). **m-p** Same as in i-l for monkey B (n=2 injections, circles). For both monkeys and for both models, although there is some variability in parameter estimates between

injection sessions, especially in DDM parameter estimates for monkey S, the parameter estimates obtained using pooled data for fitting the models summarize the parameters from all the individual sessions reasonably well.

The parameter estimates of the UGM from fitting pre- and post-muscimol sessions show the relative changes from pre- to post-muscimol for the drift rate offset as seen with the HDDM and DDM, along with changes in urgency slope for both monkeys, slightly for monkey S and more so for monkey B (Extended Data Fig. 6k,l). However, when fitting the UGM with only the slope free to vary, the model failed to capture the post-muscimol decision bias data (Supplementarty Table 4 and Extended Data Fig. 7b,d in UGM-*m* row). In contrast, the UGM with the drift rates free to vary captured the choice data almost as well as the full UGM (Supplementary Table 4 and Extended Data Fig. 7b,d in UGM-δ row), indicating that the change in drift rate offset best explains the effect of muscimol on decision bias. This result confirmed the results obtained using the HDDM and DDM.

Model Comparison

To compare across model variants, we found the amount of variance explained for in-sample and out-of-sample prediction of choice-RT data (Supplementary Table 4). In-sample prediction refers to how well each model performed at predicting proportion of choices and RT quantiles of the data that were used to fit model parameters. We used 80% of the in-session trials (randomly picked with a uniform distribution), to fit model parameters and determine in-sample prediction statistics. Out-of-sample predictions were obtained using the 20% of the data that was not used to fit the data. Bayesian Information Criteria (BIC) and Akaike Information Criteria (AIC) were also calculated for the DDM and UGM models in ChaRTr, which are measures of in-sample prediction collapsed across data types (choice performance and RT) that include penalties for model

complexity. BIC more often favors models that match the ground truth while AIC more often favors models that will be more predictive of new data[17].

R^2_{pred} measures the amount of variance in the data (either in-sample or out-of-sample), explained by the prediction of the model. Obtaining estimates of variance explained for RT statistics (mean, 25th percentiles, median, and 75th percentiles) and proportions of choice are useful to understand how one model performs better than another model. R^2_{pred} ranges from -∞ to 1 and can be multiplied by 100 for an estimate of -∞ to 100% percentage of variance explained by the data statistic (Q) by prediction. Note that values can be negative indicating that the amount of variance in prediction is greater than the variance of the statistic (Supplementary Table 4). However, the relative values from one model to the next are still informative about the improvement in prediction. R^2_{pred} is defined as one minus the mean squared error of prediction of a certain statistic (Q) scaled by the estimated variance of a statistic using a number of observations J equal to the total number of coherence conditions (K) across experimental sessions (S) for a total $J = K*S$.

For calculating the R^2_{pred} for the non-hierarchical DDM and UGM that were fit to individual pre-muscimol sessions for model comparison between the HDDM, DDM, and UGM, we used the individual session's predictions (with the 24% and 36% conditions excluded, see UGMs section) in calculating the mean squared error of prediction for R^2_{pred}, while the R^2_{pred} of the HDDM was also calculated from fitting the exact pre-muscimol data for each monkey without the 36% and 24% coherence conditions (model fits denoted with ** in Supplementary Table 4). For all other R^2_{pred} for the non-hierarchical DDM and the UGMs (i.e. DDM, UGM, UGM-δ, UGM-m), we found the difference between the predictions from those pooled fits to each individual session's actual data to calculate the error of prediction.

Parameter recovery

We performed parameter recovery to ensure that QMPE would return accurate estimates of free-to-vary UGM and DDM parameters. For our process of parameter recovery, we first simulated choice and RT data for the DDM and UGM, using parameter values approximated from previous fits of our pre-muscimol data for monkey B using the ChaRTr. We then fitted the UGM and DDM to the simulated data to determine whether the parameter estimates from the fits matched the parameters that generated the simulations. All five sets of parameter estimates from the five fits demonstrated that the models used (DDM and UGM with fixed bound and free slope) can recover parameters accurately. We found that the parameter estimates closely matched the known parameters for the DDM and the UGM (Supplementary Fig. 2).

Supplementary Note references

1. Ratcliff, R. & McKoon, G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput* **20**, 873-922 (2008).

2. Ratcliff, R., Smith, P.L., Brown, S.D. & McKoon, G. Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences* **20**, 260-281 (2016).

3. Macmillan, N.A. & Creelman, C.D. *Detection theory: A user's guide* (Lawerence Erlbaum Associates, Inc., Mahwah, New Jersey, 2004).

4. Kloosterman, N.A., *et al.* Humans strategically shift decision bias by flexibly adjusting sensory evidence accumulation. *eLife* **8**, e37321 (2019).

5. Ratcliff, R. & Childers, R. Individual differences and fitting methods for the two-choice diffusion model of decision making. *Decision (Wash D C)* **2015** (2015).

6. Vandekerckhove, J., Tuerlinckx, F. & Lee, M.D. Hierarchical diffusion models for two-choice response times. *Psychol Methods* **16**, 44-62 (2011).

7. Gelman, A., *et al. Bayesian data analysis* (CRC Press, 2014).

8. Lee, M.D. & Wagenmakers, E.-J. Bayesian cognitive modeling. in *A Practical Course* (Cambridge University Press, Cambridge, 2014).

9. Drugowitsch, J., Moreno-Bote, R., Churchland, A.K., Shadlen, M.N. & Pouget, A. The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience* **32**, 3612-3628 (2012).

10. Nunez, M.D., Gosai, A., Vandekerckhove, J. & Srinivasan, R. The latency of a visual evoked potential tracks the onset of decision making. *Neuroimage* **197**, 93-108 (2019).

11. Wabersich, D. & Vandekerckhove, J. Extending JAGS: a tutorial on adding custom distributions to JAGS (with a diffusion model example). *Behav Res Methods* **46**, 15-28 (2014).

12. Lerche, V. & Voss, A. Model complexity in diffusion modeling: benefits of making the model more parsimonious. *Front Psychol* **7**, 1324 (2016).

13. Boehm, U., *et al.* Estimating across-trial variability parameters of the Diffusion Decision Model: Expert advice and recommendations. *Journal of Mathematical Psychology* **87**, 46-75 (2018).

14. Ratcliff, R., Hasegawa, Y.T., Hasegawa, R.P., Smith, P.L. & Segraves, M.A. Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J. Neurophysiol.* **97**, 1756-1774 (2007).

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

16. Heathcote, A., Brown, S. & Mewhort, D.J. Quantile maximum likelihood estimation of response time distributions. *Psychon Bull Rev* **9**, 394-401 (2002).

17. Chandrasekaran, C. & Hawkins, G.E. ChaRTr: An R toolbox for modeling choices and response times in decision-making tasks. *J Neurosci Methods* **328**, 108432 (2019).

18. Cisek, P., Puskas, G.A. & El-Murr, S. Decisions in changing conditions: the urgency-gating model. *J Neurosci* **29**, 11560-11571 (2009).

19. Evans, N.J., Hawkins, G.E., Boehm, U., Wagenmakers, E.J. & Brown, S.D. The computations that support simple decision-making: A comparison between the diffusion and urgency-gating models. *Sci Rep* **7**, 16433 (2017).

20. Thura, D. & Cisek, P. On the difference between evidence accumulator models and the urgency gating model. *Journal of Neurophysiology* **115**, 622-623 (2016).

Supplementary Tables

Exp. #	Monkey	Treatment	Volume (uL)	Conc. (ug/uL)	Side of injection	Decision task type	Rf_x (°)	Rf_y (°)
1	S	Muscimol	0.75	0.5	Left	Delay	16.7	-14.9
2	S	Muscimol	0.5	0.5	Left	Delay	13.2	-4.0
3	S	Muscimol	0.5	0.5	Left	Delay	12.8	-15.7
4	B	Muscimol	0.5	4.0	Right	Delay	-19.1	-9.6
5	B	Muscimol	0.5	4.0	Right	Delay	-19.9	-13.0
6	B	Muscimol	0.5	4.0	Right	Delay	-21.1	-9.6
7	B	Muscimol	0.5	4.0	Right	Delay	-20.9	-7.6
8	B	Muscimol	0.5	4.0	Left	Delay	9.7	-18.1
9	B	Muscimol	0.5	4.0	Left	Delay	15.2	-19.5
10	S	Muscimol	2.0	2.0	Right	Delay	-18.4	3.0
11	S	Muscimol	1.0	2.0	Right	Delay	-15.0	8.4
12	B	Muscimol	0.5	4.0	Left	Delay	18.6	-19.6
13	B	Saline	0.5	9.0	Left	Delay	19.3	-17.9
14	B	Saline	0.5	9.0	Left	Delay	19.1	-18.6
15	B	Muscimol	0.5	4.0	Left	Delay	15.9	7.1
16	B	Muscimol	0.5	4.0	Right	Delay	-15.7	15.4
17	B	Muscimol	0.5	4.0	Right	RT	-16.2	9.7
18	B	Muscimol	0.5	4.0	Right	RT	-19.9	6.5
19	S	Muscimol	1.0	2.0	Left	RT	16.0	-10.0
20	S	Saline	1.0	9.0	Left	RT	15.1	-9.6
21	S	Saline	1.0	9.0	Left	RT	12.0	-5.0
22	S	Muscimol	1.0	1.5	Left	RT	12.6	-9.6
23	S	Muscimol	1.0	1.5	Left	RT	10.0	-6.0
24	S	Muscimol	1.0	1.5	Left	RT	11.0	-15.0
25	S	Muscimol	1.0	1.5	Left	RT	11.0	-3.5
26	S	Saline	1.0	9.0	Right	RT	-11.5	12.0
27	S	Muscimol	1.0	1.5	Left	RT	14.0	-4.0
28	S	Saline	1.0	9.0	Right	RT	-9.0	11.0
29	S	Muscimol	1.0	1.5	Right	RT	-7.0	10.5

Supplementary Table 1. Unilateral Muscimol Injections into the monkey SC. (Associated with Figs. 1 and 2 main text) Details of the injections in two monkeys. Each row shows an individual injection (n=29) and each column shows the specifics of one particular injection. The column labeled monkey indicates that the experiment was performed in monkey S or monkey B and the treatment column indicates whether the injection was saline or muscimol. The next column indicates the total injection volume in µl. The column labeled Conc. (µg/µl) lists the concentration of muscimol or saline and the side of injection column lists the side of the brain in which the injection was made. The decision task type column indicates the version of the decision task used (RT or delay). The RF_x column shows the horizontal position of the RF for each injection, and the RF_y column shows the vertical position of the RF.

Muscimol psychometric function parameter comparisons						
Variables	Data sessions	Test	n	p	Test stat.	CI
Alpha (α)	Pre - rec	Wilcoxon signed rank test with Bonferroni correction	23	5.139×10^{-3}	230	—
Beta (β)	Pre - rec	Paired t-test with Bonferroni correction	23	0.203	-1.313	(-0.019, 0.006)
Saline psychometric function parameter comparisons						
Variables	Data sessions	Test	n	p	Test stat.	CI
Alpha (α)	Pre - rec	Paired t-test with Bonferroni correction	6	0.368	0.988	(-2.754, 5.257)
Beta (β)	Pre - rec	Paired t-test with Bonferroni correction	6	0.128	-1.822	(-0.067, 0.018)
Muscimol chronometric function parameter comparisons						
Variables	Data sessions	Test	n	p	Test stat.	CI
Slope awayIF	Pre - rec	Paired t-test with Bonferroni correction	9	0.209	1.366	(-0.873, 2.593)
Slope toIF	Pre - rec	Paired t-test with Bonferroni correction	9	0.409	-0.872	(-3.101, 1.609)
Intercept awayIF	Pre - rec	Paired t-test with Bonferroni correction	9	0.030	2.630	(-4.886, 216.526)
Intercept toIF	Pre - rec	Paired t-test with Bonferroni correction	9	0.007	3.613	(32.553, 240.499)
Saline chronometric function parameter comparisons						
Variables	Data sessions	Test	n	p	Test stat.	CI
Slope awayIF	Pre - rec	Paired t-test with Bonferroni correction	4	0.983	-0.023	(-6.257, 6.189)
Slope toIF	Pre - rec	Paired t-test with Bonferroni correction	4	0.579	0.621	(-0.879, 1.186)
Intercept awayIF	Pre - rec	Paired t-test with Bonferroni correction	4	0.736	0.370	(-71.196, 85.050)
Intercept toIF	Pre - rec	Wilcoxon signed rank test with Bonferroni correction	4	0.250	9	—

Supplementary Table 2. Statistics for recovery psychometric and chronometric function fits

(Associated with Fig. 2 of the main text) All statistical tests that were performed with the recovery data for the muscimol and saline injections for the α, β, RT slope, and RT intercept parameters associated with the fits in Extended Data Fig. 4 are listed. The first column indicates on which parameter the statistics were performed, the second column indicates which paired groups of data sessions were compared, the third column indicates which paired test was performed, the n indicates number of observations, followed by the p-value, test-statistic, and the confidence interval if a t-test was performed. The statistical results are sectioned by the injection (saline or muscimol) data that the statistics were performed on, marked by horizontal lines. We performed pairwise comparisons between parameters of pre-muscimol to recovery (the 2nd pairwise comparison of pre-muscimol to post-muscimol are shown in the main text) with Bonferroni corrections (two pairwise comparison tests performed, $\alpha = 0.05/2 = 0.025$).

Muscimol accuracy comparisons

Multiple comparisons of differences in accuracies between pre- and post-muscimol sessions for monkey S

Task / side	Sessions	Test	n	p	Test stat.	CI
Decision toIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	12	0.002	0	—
Decision awayIF	Pre, post	Paired samples t-test with Bonferroni correction	12	7.000E-06	7.967	(0.078, 0.139)
Selection toIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	12	0.656	28	—
Selection awayIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	12	0.779	20	—

Multiple comparisons of differences in accuracies between pre- and post-muscimol sessions for monkey B

Task / side	Sessions	Test	n	p	Test stat.	CI
Decision toIF	Pre, post	Paired samples t-test with Bonferroni correction	11	2.280E-04	-5.599	(-0.263, -0.113)
Decision awayIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	11	0.003	66	—
Selection toIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	11	0.249	5	—
Selection awayIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	11	0.050	4	—

Saline accuracy comparisons

Multiple comparisons of accuracies between sessions for monkey S

Task / side	Sessions	Test	n	p	Test stat.	CI
Selection toIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	4	0.317	0	—
Selection awayIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	4	0.180	0	—

Multiple comparisons of accuracies between sessions for monkey B

Task/side	Sessions	Test	n	p	Test stat.	CI
Selection toIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	2	0.655	1	—
Selection awayIF	Pre, post	Wilcoxon signed rank test with Bonferroni correction	2	0.655	2	—

Muscimol peak saccadic velocity comparisons

Multiple comparisons between toIF peak velocities between the decision and selection task post-muscimol for monkey S

Exp. #	Task	Test	n (Decision)	n (Selection)	p	Test stat.	CI
1	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	344	123	1.58×10^{-7}	-5.243	—
2	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	401	133	0.081	-1.744	—
3	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	218	162	0.292	1.054	—
11	Decision, selection	Independent t-test with Bonferroni correction	246	73	0.335	0.966	(-22.670, 44.665)
19	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	210	170	7.109×10^{-22}	9.162	—
22	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	261	169	5.664×10^{-23}	9.869	—
23	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	237	166	0.015	2.438	—
27	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	290	172	2.915×10^{-16}	8.177	—
29	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	221	156	1.210×10^{-11}	6.779	—

Multiple comparisons between toIF peak velocities between the decision and selection task post-muscimol for monkey B

Exp. #	Task	Test	n (Decision)	n (Selection)	p	Test stat.	CI
4	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	80	161	0.816	-0.233	—
5	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	283	172	2.237×10^{-29}	11.253	—
6	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	173	126	0.051	1.955	—
7	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	212	167	9.141×10^{-5}	3.912	—
8	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	253	170	2.800×10^{-17}	8.455	—
9	Decision, selection	Independent t-test with Bonferroni correction	321	85	7.059×10^{-9}	5.915	(25.677, 70.254)
12	Decision, selection	Independent t-test with Bonferroni correction	299	75	1.277×10^{-30}	12.615	(94.601, 147.342)
15	Decision, selection	Wilcoxon rank-sum test with Bonferroni correction	174	174	6.028×10^{-37}	12.698	—

Supplementary Table 3. Statistics for choice accuracy and peak saccadic velocity analysis

(Associated with Fig. 3 of the main text) This table reports the statistical results of comparing differences in performance accuracy for the decision and selection tasks, for toIF and awayIF trials, pre- and post-muscimol. Also listed are the statistics for assessing the differences in post-muscimol peak saccadic velocity for toIF saccades made in the decision and the selection tasks, for each muscimol injection shown in Fig. 3d and e. N=9 muscimol injections for monkey S and n=8 muscimol injections for monkey B. The table is divided by horizontal black lines into three main sections - muscimol accuracy comparisons, saline accuracy comparisons, and muscimol peak saccadic velocity comparisons. Each of these main sections is further separated by monkey. The muscimol accuracy comparisons section and the saline accuracy comparisons section have the same ordering of columns. The first column, "task/side" describes the task and side of accuracy from pre to post, i.e. "decision toIF" means the test assessed differences in accuracy for toIF trials in the decision task between pre- and post-muscimol. The second column describes the data sessions compared and the next column describes the statistical test that was used to compare the two data sets. We used a paired samples t-test when both samples (data from pre and post) consisted of normally distributed data or the Wilcoxon signed rank test when either of the two samples consisted of non-normally distributed data, implementing a Bonferroni correction (four accuracy multiple comparison tests per monkey, $\alpha = 0.05/4 = 0.0127$). The next column indicates the sample size (number of injections). The column after shows the p-value of the statistical comparison analysis followed by the test statistic. The last column indicates the confidence interval if a t-test was performed. For the section assessing differences in post-muscimol peak saccadic velocity for toIF saccades made in the decision task compared to the selection task, on an injection-by-injection basis, the first column describes the injection

experiment number, the second column shows the tasks from which the data are compared, and the next column describes the specific statistical test that was used to compare the two samples of peak velocities of saccades made. A two independent sample t-test (for normally distributed data sessions) or Wilcoxon rank-sum test (for non-normally distributed data sessions) was used to compare the post-muscimol toIF peak saccadic velocities between the decision and selection tasks with Bonferroni corrections of $\alpha = 0.05/9 = 0.0056$ for the tests for monkey S, since nine t-tests were performed on data from each injection, and $\alpha = 0.05/8 = 0.0063$ for monkey B, since eight t-tests were performed on data from each injection, totaling 17 injections. Six injections were excluded due to technical issues with the eye tracker that impacted measurement of eye speed but not assessment of choice or RT. The next column indicates the sample size, the number of trials where the peak saccadic velocities were recorded, from the decision task and the following column indicates the sample size of the peak saccadic velocities from the selection task. The next column relates the p-value of the test followed by the test statistic in the next column for the specific statistical test performed. The last column indicates the confidence interval if a t-test was performed.

Monkey S in-sample predictions

Model	Choice data	RT mean	RT 25th percentile	RT median	RT 75th percentile	AIC	BIC
HDDM	98.3%	55.3%	59.8%	79.0%	54.9%	N/A	N/A
HDDM-Δ	97.6%	29.2%	26.9%	54.8%	31.0%	N/A	N/A
HDDM-*w*	94.5%	15.2%	28.6%	49.2%	15.6%	N/A	N/A
HDDM-τ	89.9%	56.2%	60.7%	69.2%	50.2%	N/A	N/A
HDDM-*a,w*	94.6%	36.4%	57.8%	74.8%	31.7%	N/A	N/A
HDDM-*G*	94.9%	46.0%	56.9%	74.3%	44.0%	N/A	N/A
HDDM Pre **	98.9%	19.2%	34.7%	69.8%	16.5%	N/A	N/A
DDM Pre **	98.0%	61.2%	34.0%	63.0%	63.3%	N/A	N/A
UGM Pre **	93.5%	78.6%	-11.0%	52.7%	65.8%	N/A	N/A
DDM Pre	92.8%	-7.3%	-24.0%	-3.4%	0.2%	20183.22	20269.41
UGM Pre	92.5%	6.3%	-54.6%	-10.4%	4.9%	21499.91	21586.10
DDM Post	90.0%	1.5%	-14.3%	5.0%	4.1%	19842.78	19929.02
UGM Post	89.1%	11.2%	-93.9%	-8.9%	11.2%	21121.18	21207.41
UGM-δ Post	87.8%	-6.6%	10.9%	-24.2%	-7.3%	21520.48	21588.23
UGM-*m* Post	20.0%	-14.7%	-28.8%	-144.3%	-3.3%	22464.49	22470.65

Monkey B in-sample predictions

Model	Choice data	RT mean	RT 25th percentile	RT median	RT 75th percentile	AIC	BIC
HDDM	99.3%	65.9%	76.2%	88.4%	69.2%	N/A	N/A
HDDM-Δ	98.3%	21.0%	-2.3%	49.0%	28.5%	N/A	N/A
HDDM-w	93.5%	5.1%	-15.2%	33.9%	17.3%	N/A	N/A
HDDM-τ	91.3%	62.5%	65.1%	74.4%	66.6%	N/A	N/A
HDDM-*a,w*	94.0%	19.0%	63.6%	81.8%	4.2%	N/A	N/A
HDDM-*G*	97.2%	64.4%	74.5%	86.5%	68.5%	N/A	N/A
HDDM Pre**	99.2%	-277.9%	-99.4%	36.9%	-17.9%	N/A	N/A
DDM Pre **	98.6%	-30.2%	-166.6%	-6.0%	35.2%	N/A	N/A
UGM Pre **	97.6%	61.8%	11.9%	51.4%	61.9%	N/A	N/A
DDM Pre	98.2%	-147.2%	-94.7%	43.0%	-44.8%	6658.55	6728.62
UGM Pre	98.1%	53.6%	-9.2%	52.0%	65.2%	6392.01	6462.07
DDM Post	96.7%	63.9%	16.6%	45.6%	64.4%	5788.23	5857.04
UGM Post	97.5%	74.0%	11.4%	53.1%	62.0%	5851.70	5920.51
UGM-δ Post	83.1%	-165.6%	-149.5%	-56.9%	-77.0%	6524.93	6578.99
UGM-*m* Post	26.2%	22.5%	-1.3%	26.1%	27.8%	6187.52	6192.43

Monkey S out-of-sample predictions

Model	Choice data	RT mean	RT 25th percentile	RT median	RT 75th percentile
HDDM	96.7%	-55.5%	-59.5%	46.0%	-60.9%
HDDM-Δ	96.3%	-21.7%	-146.0%	32.2%	-41.0%
HDDM-*w*	95.1%	-50.3%	-125.4%	24.3%	-73.9%
HDDM-τ	95.6%	15.7%	-32.0%	44.6%	-16.6%
HDDM-*a,w*	95.7%	-129.6%	-54.4%	34.9%	-130.1%
HDDM-*G*	97.9%	-78.2%	-62.6%	36.4%	-91.1%
HDDM Pre **	73.3%	8.7%	-14.1%	18.7%	4.2%
DDM Pre **	55.5%	18.6%	4.0%	34.1%	-11.4%
UGM Pre **	52.7%	45.9%	3.6%	20.6%	26.9%
DDM Pre	76.4%	-8.3%	-21.7%	-2.7%	-10.5%
UGM Pre	74.5%	7.6%	-7.8%	-4.3%	4.0%
DDM Post	66.9%	8.9%	-16.4%	2.6%	8.2%
UGM Post	63.5%	13.4%	-15.8%	0.0%	8.2%
UGM-δ Post	61.9%	-1.0%	2.9%	13.1%	1.4%
UGM-*m* Post	38.5%	-22.2%	-29.1%	-68.2%	-4.2%

Monkey B out-of-sample predictions

Model	Choice data	RT mean	RT 25th percentile	RT median	RT 75th percentile
HDDM	93.0%	-132.5%	10.7%	27.3%	-79.1%
HDDM-Δ	91.9%	-273.5%	-4.2%	9.6%	-262.3%
HDDM-w	94.9%	-285.2%	14.5%	38.5%	-271.2%
HDDM-τ	94.9%	-178.2%	-23.8%	25.7%	-132.0%
HDDM-*a,w*	95.0%	-218.4%	15.31%	38.2%	-185.3%
HDDM-*G*	93.4%	-135.9%	15.7%	38.8%	-93.6%
HDDM Pre **	60.0%	-189.9%	13.1%	-9.6%	-129.7%
DDM Pre **	83.4%	12.5%	-94.0%	-58.9%	28.1%
UGM Pre **	81.4%	19.1%	1.3%	20.6%	22.3%
DDM Pre	89.7%	-4.6%	-53.2%	6.3%	18.6%
UGM Pre	89.1%	34.5%	-8.7%	38.5%	37.2%
DDM Post	88.0%	-65.0%	-29.0%	-1.0%	-133.8%
UGM Post	84.3%	-8.3%	-27.9%	-31.4%	-34.2%
UGM-δ Post	77.1%	-168.5%	-133.2%	-77.7%	-66.2%
UGM-*m* Post	42.8%	5.1%	0.5%	24.2%	-44.5%

Supplementary Table 4. Table of R^2_{pred} for in-sample and out-of-sample predictions (Associated with Fig. 5 of the main text). For model comparisons, we obtained goodness of fit measures such as R^2_{pred} for choice performance data, RT mean, RT 25th percentile, RT median, RT 75th percentiles, along with the AIC and BIC values (where can be applied) for the in-sample predictions where various models were fit to a random 80% of the data (Supplementary Note) for monkey S (top left) and monkey B (top right). Also shown are the R^2_{pred} values for out-of-sample predictions (bottom left for monkey S and bottom right for monkey B), where the predictions of the fits are compared to the 20% of the data that was not fit to assess if the data was overfit and

see if the same conclusions of model comparison are supported. Each section of the table starts with the model comparisons between the full HDDM, HDDM with free drift rate offset (HDDM-Δ), HDDM with free proportional start point (HDDM-w), the HDDM with free non-decision time (HDDM-τ), and the HDDM with the free proportional start point and bound (HDDM-a,w). For both monkeys, the in-sample R^2_{pred} shows that the HDDM-Δ fit explains the choice data best compared to HDDM variants with specific parameters free to vary. The HDDM-a,w also captures some of the choice data but not to the extent of the HDDM-Δ, which fully captures the magnitude of the decision bias, which is visualized in Extended Data Fig. 7. Like HDDM-a,w, the HDDM-w also captures some of the decision bias, but not sufficiently to capture the magnitude of decision bias seen in the data. The HDDM-τ explains the least of the choice data out of the models. The R^2_{pred} goodness of fits are also shown for HDDM-G, the full HDDM with the same parameters but with drift rates constrained to a linear comparison of two accumulators to test a model with two independent accumulators with gain (Supplementary Note). The HDDM-G fits the data well but not as well as the full HDDM. The predictions from the model fits denoted with ** are the models that were fitted to individual pre-muscimol data (rather than pooled pre-muscimol data) with nine conditions (17%, 10%, 5%, 3% separately for toIF and awayIF trials, with the 0% collapsed across toIF and awayIF trials) for direct model comparison between the HDDM, DDM, and UGM for each monkey (Supplementary Note) to see whether the UGM is also a reasonable decision-making model assumption as the HDDM/DDM in our behavioral data. The next four models are the DDM and UGM fits for pre and post data that were pooled (Supplementary Note) to show the measures of fits that correspond to the parameter estimates seen in Extended Data Fig. 6. Finally, the model comparisons for the full UGM, UGM with free drift rates (UGM-δ), and UGM with free urgency slope (UGM-m) for the post-muscimol data (while keeping the other parameters fixed at pre-

muscimol parameter estimate values, Supplementary Note), are shown. For monkey S, the AIC and BIC values indicate that the UGM-δ is a better fit than the UGM-*m*, whereas for monkey B, the AIC and BIC values favor the UGM-*m* than the UGM-δ. However, the in-sample R^2_{pred} values for choice data show that the UGM-*m* fails to capture the choice data for both monkeys whereas the UGM-δ captures the choice data almost as well as the full UGM (also visualized in Extended Data Fig. 7). The UGM-δ better explains the shift in decision bias than UGM-*m* in both monkeys, while the UGM-*m* better explains the RT distribution change from monkey B, but not for monkey S, likely due to individual strategy differences. The out-of-sample predictions generally show the same results although they are more variable.

	Original	Shifted	Narrow	Wide
Prob. of a drift rate offset decrease in monkey B	98.99	99.22	98.83	99.26
Prob. of a drift rate offset decrease in monkey S	99.72	99.61	99.73	99.82
BF for a drift rate offset not equal to zero in monkey B	17.87	13.47	21.01	6.63
BF for a drift rate offset not equal to zero in monkey S	3194165.63	45818420.04	45451.49	10040625.32
Prob. of a start point decrease in monkey B	29.48	31.29	31.03	30.54
Prob. of a start point decrease in monkey S	95.14	95.17	94.94	95.22
BF for a start point equal to .5 in monkey B	2.82	5.93	1.59	6.00
BF for a start point equal to .5 in monkey S	11.51	18.20	5.72	22.50
Prob. of an increase in non-decision time in monkey B	97.14	96.47	96.99	97.00
Prob. of an increase in non-decision time in monkey S	94.79	93.30	94.65	93.70
Prob. of an increase in symmetric boundary in monkey B	95.19	94.96	93.87	95.14
Prob. of an increase in symmetric boundary in monkey S	78.85	79.42	78.10	80.32
Prob. of an increase in lapse proportion in monkey B	53.98	53.15	54.28	51.48
Prob. of an increase in lapse proportion in monkey S	72.52	69.38	70.25	62.55

Supplementary Table 5. Probability and Bayes Factor (BF) estimates derived from posterior distributions for specific parameter changes with different prior distributions for hierarchical mean parameters (μ). "Original" refers to the HDDM listed in the Main text and Methods. "Shifted" refers to a HDDM with shifted hierarchical mean priors from "Original" HDDM (Supplementary Note). "Narrow" refers to a HDDM with half the width hierarchical mean priors compared to the "Original" HDDM (Supplementary Note). "Wide" refers to a HDDM with double the width hierarchical mean priors compared to the "Original" HDDM.