

# Competing neural representations of choice shape evidence accumulation in humans

Krista Bond,<sup>1,2,3\*</sup>, Javier Rasero,<sup>1</sup>, Raghav Madan<sup>4</sup>,  
Jyotika Bahuguna<sup>1</sup>, Jonathan Rubin<sup>2,5</sup>, Timothy Verstynen<sup>1,2,3,6\*</sup>

<sup>1</sup>Department of Psychology, Carnegie Mellon University, United States

<sup>2</sup>Center for the Neural Basis of Cognition, United States

<sup>3</sup>Carnegie Mellon Neuroscience Institute, United States

<sup>4</sup>Department of Biomedical and Health Informatics,  
University of Washington, United States

<sup>5</sup>Department of Mathematics, University of Pittsburgh, United States

<sup>6</sup>Department of Biomedical Engineering, Carnegie Mellon University, United States

\*To whom correspondence should be addressed;

kbond@andrew.cmu.edu & timothyv@andrew.cmu.edu.

1      **Changing your mind requires shifting the way streams of information lead to**  
2      **a decision. Using *in silico* experiments we show how the cortico-basal ganglia-**  
3      **thalamic (CBGT) circuits can feasibly implement shifts in the evidence accu-**  
4      **mulation process. When action contingencies change, dopaminergic plasticity**  
5      **redirects the balance of power, both within and between action representa-**  
6      **tions, to divert the flow of evidence from one option to another. This finding**  
7      **predicts that when competition between action representations is highest, the**  
8      **rate of evidence accumulation is lowest. We then validate this prediction in**  
9      **a sample of *homo sapiens* as they perform an adaptive decision-making task**  
10     **while whole-brain hemodynamic responses are recorded. These results paint**

11        **a holistic picture of how CBGT circuits manage and adapt the evidence accu-**  
12        **mulation process in mammals.**

13        **One-sentence Summary:** Interactions between cortical and subcortical circuits in the mam-  
14        malian brain flexibly control the flow of information streams that drive decisions by shifting the  
15        balance of power both within and between action representations.

16        **Introduction**

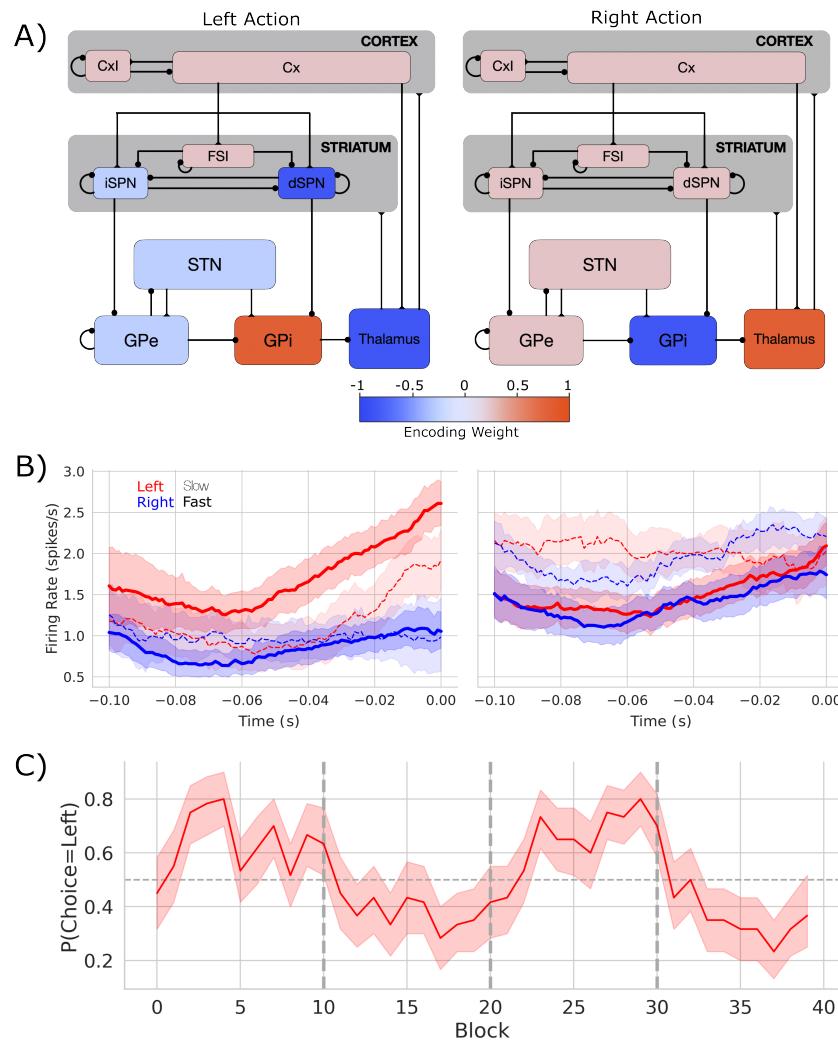
17        Choice is fundamentally driven by information. The process of deciding between available ac-  
18        tions is continually updated using incoming sensory signals, processed at a given accumulation  
19        rate, until sufficient evidence is reached to trigger one action over the other (1, 2). The param-  
20        eters of this evidence accumulation process are highly plastic, adjusting to both the reliability  
21        of sensory signals (3–7) and previous choice history (8–13), to balance the speed of a given  
22        decision with local demands to choose the right action.

23        We recently showed how environmental change influences the decision process by period-  
24        ically switching the reward associated with a given action in a 2-choice task (7). This reward  
25        contingency change induces competition between old and new action values, forcing a change-  
26        of-mind about the most rewarding option. This internal competition prompts humans to dy-  
27        namically reduce the rate at which they accumulate evidence (drift-rate,  $v$ , in a normative drift  
28        diffusion model, DDM (2)) and sometimes also increases the threshold of evidence they need to  
29        trigger an action (boundary height,  $a$ ) (7). The result is a change of the decision policy to a slow,  
30        exploratory state. Over time feedback-learning pushes the system back into an exploitative state  
31        until the environment changes again (see also (11) and (12)).

32        Here we investigate the underlying neural mechanisms that drive dynamic decision policies  
33        in a changing environment. We start with a set of theoretical experiments, using biologically re-  
34        alistic spiking network models, to test how the cortico-basal ganglia-thalamic (CBGT) circuits

35 influence the evidence accumulation process (14–18). These experiments both explain previous  
 36 results (7) and make specific predictions as to how competition between action representations  
 37 drives changes in the decision policy. We then test these predictions in humans using a high-  
 38 powered, within-participant neuroimaging design, collecting data over thousands of trials where  
 39 action-outcome contingencies change on a semi-random basis.

40



41

42 **Fig. 1. Biologically based CBGT network dynamics and behavior.** A) Each CBGT nucleus  
 43 is organized into left and right action channels with the exception of a common population

44 of striatal fast spiking interneurons (FSIs) and another of cortical interneurons (CxI). Values  
45 show encoded weights for a left action. Network schematic adapted from (19). B) Firing rate  
46 profiles for D1-SPNs (left panel) and D2-SPNs (right panel) prior to stimulus onset (t=0) for  
47 a left choice. D1-SPN activity in left and right channels is shown in red and blue, respectively.  
48 Thick solid lines represent fast trials (short RTs) and thin dashed lines represent slow trials (long  
49 RTs). C) Choice probability for the CBGT network model. The reward for left and right actions  
50 changed every 10 trials, marked by vertical dashed lines. The horizontal dashed line represents  
51 chance performance.

## 52 **Results**

### 53 **CBGT circuits can control decision parameters under uncertainty**

54 Both theoretical (9, 12, 14, 19–21) and experimental (18) evidence suggest that the CBGT cir-  
55 cuits play a critical role in the evidence accumulation process (for a review see (22)). The  
56 canonical CBGT circuit (Fig. 1A) includes two dissociable control pathways: the direct (fa-  
57 cilitation) and indirect (suppression) pathways (23, 24). A critical assumption of the canonical  
58 model is that the basal ganglia are organized into multiple "channels" mapped to specific action  
59 representations (25, 26), each containing a direct and indirect pathway. While a strict, segregated  
60 action channel organization may not accurately reflect the true underlying circuitry, striatal neu-  
61 rons have been shown to organize into task-specific spatiotemporal assemblies that qualitatively  
62 reflect independent action representations (27–31). Within these action channels, activation of  
63 the direct pathway, via cortical excitation of D1-expressing spiny projection neurons (SPNs) in  
64 the striatum, releases GABAergic signals that can suppress activity in the CBGT output nucleus  
65 (internal segment of the globus pallidus, GPi, in primates or substantia nigra pars reticulata,  
66 SNr, in rodents) (26, 32–34). This relieves the thalamus from tonic inhibition, thereby exciting  
67 postsynaptic cortical cells and facilitating action execution. Conversely, activation of the indi-

68 rect pathway via D2-expressing SPNs in the striatum controls firing in the external segment of  
69 the globus pallidus (GPe) and the subthalamic nucleus (STN), resulting in strengthened basal  
70 ganglia inhibition of the thalamus. This weakens drive to postsynaptic cortical cells and reduces  
71 the likelihood that an action is selected in cortex.

72 Critically, the direct and indirect pathways converge in the GPi/SNr (35, 36). This suggests  
73 that these pathways compete to control whether each specific action is selected (37). The ap-  
74 parent winner-take-all selection policy and action-channel like coding (27–31) also imply that  
75 action representations themselves compete. Altogether, this neuroanatomical evidence suggests  
76 that competition both between and within CBGT pathways controls the rate of evidence accu-  
77 mulation during decision making (12, 15, 19).

78 To illustrate this process, we designed a spiking neural network model of the CBGT cir-  
79 cuits, shown in Fig. 1A, with dopamine-dependent plasticity occurring at the corticostriatal  
80 synapses (17, 38). The network performed a probabilistic 2-arm bandit task with switching re-  
81 ward contingencies (7; see Supp. Methods). The experimental task followed the same general  
82 structure as our prior work (7). In brief, the network selected one of two targets, each of which  
83 returned a reward according to a specific probability distribution. The relative reward probabili-  
84 ties for each target were held constant at 75% and 25% and the action-outcome contingency  
85 was changed every 10 trials, on average. For the purpose of this study we focus primarily on  
86 the neural and behavioral effects that occur around the switching of the optimal target. We  
87 used four different network instances (see Supp. Methods) as a proxy for simulating individual  
88 differences over human participants.

89 Figure B shows the firing rates of dSPNs and iSPNs in the left action channel, time-locked  
90 to selection onset (when thalamic units exceed 30Hz, t=0), for both fast (<196ms) and slow  
91 (> 314.5ms) decisions. As expected, the dSPNs show a ramping of activity as decision onset  
92 is approached and the slope of this ramp scales with response speed. In contrast, we see that

93 iSPN firing is sustained during slow movements and weakly ramps during fast movements.  
94 However, iSPN firing was relatively insensitive to left versus right decisions. This is consistent  
95 with our previous work showing that differences in direct pathways track primarily with choice  
96 while indirect pathway activity modulates overall response speeds (12, 19) as supported by  
97 experimental studies (39–41).

98 We then modeled the behavior of the CBGT network using a hierarchical version of the  
99 DDM (42), a canonical formalism for the process of evidence accumulation during decision-  
100 making (2) (Fig. 2A). This model returns four key parameters with distinct influences on  
101 evidence accumulation. The drift rate ( $v$ ) represents the rate of evidence accumulation, the  
102 boundary height ( $a$ ) represents the amount of evidence required to cross the decision threshold,  
103 nondecision time ( $t$ ) is the delay in the onset of the accumulation process, and starting bias ( $z$ )  
104 is a bias to begin accumulating evidence for one choice over another (see Methods section).

105 We tracked internal estimates of action-value and environmental change using trial-by-trial  
106 estimates of two ideal observer parameters, the belief in the value of the optimal choice ( $\Delta B$ )  
107 and change point probability ( $\Omega$ ), respectively (see (3, 7) and Methods for details). Using these  
108 estimates, we evaluated how a suspected change in the environment and the belief in optimal  
109 choice value influenced underlying decision parameters. Consistent with prior observations in  
110 humans (7) we found that both  $v$  and  $a$  were the most pliable parameters across experimental  
111 conditions for the network. Specifically, we found that the model mapping  $\Delta B$  to drift rate and  
112  $\Omega$  to boundary height and the model mapping  $\Delta B$  to drift rate provided equivocal best fits to the  
113 data over human participants ( $\Delta DIC_{\text{null}} = -29.85 \pm 12.76$  and  $\Delta DIC_{\text{null}} = -22.60 \pm 7.28$ ,  
114 respectively; see (43) and Methods for guidelines on model fit interpretation). All other models  
115 failed to provide a better fit than the null model (Supp. Table 2). Consistent with prior work  
116 (7), we found that the relationship between  $\Omega$  and the boundary height was unreliable (mean  
117  $\beta_{a \sim \Omega} = 0.069 \pm 0.152$ ; mean  $p = 0.232 \pm 0.366$ ). However, drift rate reliably increased with  $\Delta B$

118 in three of four participants (mean  $\beta_{v \sim \Delta B} = 0.934 \pm 0.386$ ; mean  $p < 0.001$ ; 4/4 participants  
119  $p < 0.001$ ; Supp. Table 3).

120 These effects reflect a stereotyped trajectory around a change point, whereby  $v$  immediately  
121 plummets and  $a$  briefly increases, with  $a$  quickly recovering and  $v$  slowly growing as reward  
122 feedback reinforces the new optimal target (7). Because prior work has shown that the change  
123 in  $v$  is more reliable than changes in  $a$  (7) and because  $v$  determines the direction of choice, we  
124 focus the remainder of our analysis on the control of  $v$ .

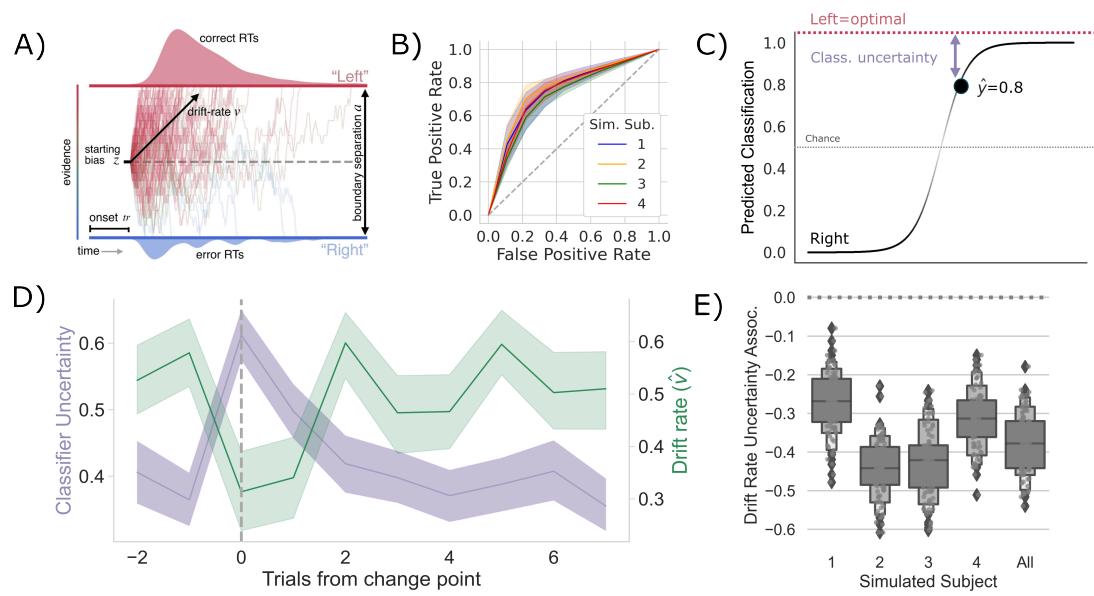
125 To test whether these shifts in  $v$  are driven by competition within and between action chan-  
126 nels, we predicted the network's decision on each trial using a LASSO-PCR classifier trained  
127 on the pre-decision firing rates of the network (see Measuring neural action representations).  
128 The cross-validated accuracy for the four simulated participants is shown in Figure B. This  
129 model was able to predict the chosen action with  $\approx 70\%$  accuracy (72-77%) for each simulated  
130 participant, with an overall accuracy of  $\approx 74\%$ . Examining the encoding pattern in the simu-  
131 lated network, we see lateralized activation over left and right action channels (Fig. 1A), with  
132 opposing weights in GPi and thalamus, and, to a lesser degree, contralateral encoding in STN  
133 and in both indirect and direct SPNs in striatum. We do not observe contralateral encoding in  
134 cortex, which likely reflects the emphasis on basal ganglia structures and lumped representation  
135 of cortex in the model design.

136 To quantify the competition between action channels, we took the unthresholded prediction  
137 from the LASSO-PCR classifier,  $\hat{y}_t$ , and calculated its distance from the optimal target (i.e.,  
138 target with the highest reward probability) on each trial (Supp. Fig. 3; Fig. 2C). This provided  
139 an estimate of the classifier's uncertainty driven by the separability of pre-decision activity  
140 across action channels. In other words, the distance from the optimal target should increase  
141 with increased co-activation of circuits that represent opposing actions.

142 If the competition in action channels is also driving  $v$ , then there should be a negative cor-

143 relation between the classifier's uncertainty and  $v$ , particularly around a change point. Indeed,  
 144 this is exactly what we see (Fig. 2D). In fact, the classifier's uncertainty and  $v$  are consistently  
 145 negatively correlated across all trials in every simulated participant and in aggregate (Fig. 2E).  
 146 Thus, in our model of the CBGT pathways, competition between action representations drives  
 147 changes in  $v$  in response to environmental change.

148



149

150 **Fig. 2. Competition between action plans should drive evidence accumulation.** A) Decision  
 151 parameters were estimated by modeling the joint distribution of reaction times and responses  
 152 within a drift diffusion framework. B) Classification performance for single-trial left and right  
 153 actions shown as an ROC curve. The gray dashed line represents chance performance. C) Pre-  
 154 dicted left and right responses. The distance of the predicted response from the optimal choice  
 155 represents classifier uncertainty for each trial. For example, here the predicted probability of  
 156 a left response on the first trial  $\hat{y}_{t_1}$  is 0.8. The distance from the optimal choice on this trial  
 157 and, thereby, the classifier uncertainty  $u_{t_1}$ , is 0.2. D) Change-point-evoked classifier uncer-  
 158 tainty (lavender) and drift rate (green). The change point is marked by a dashed line. E) The

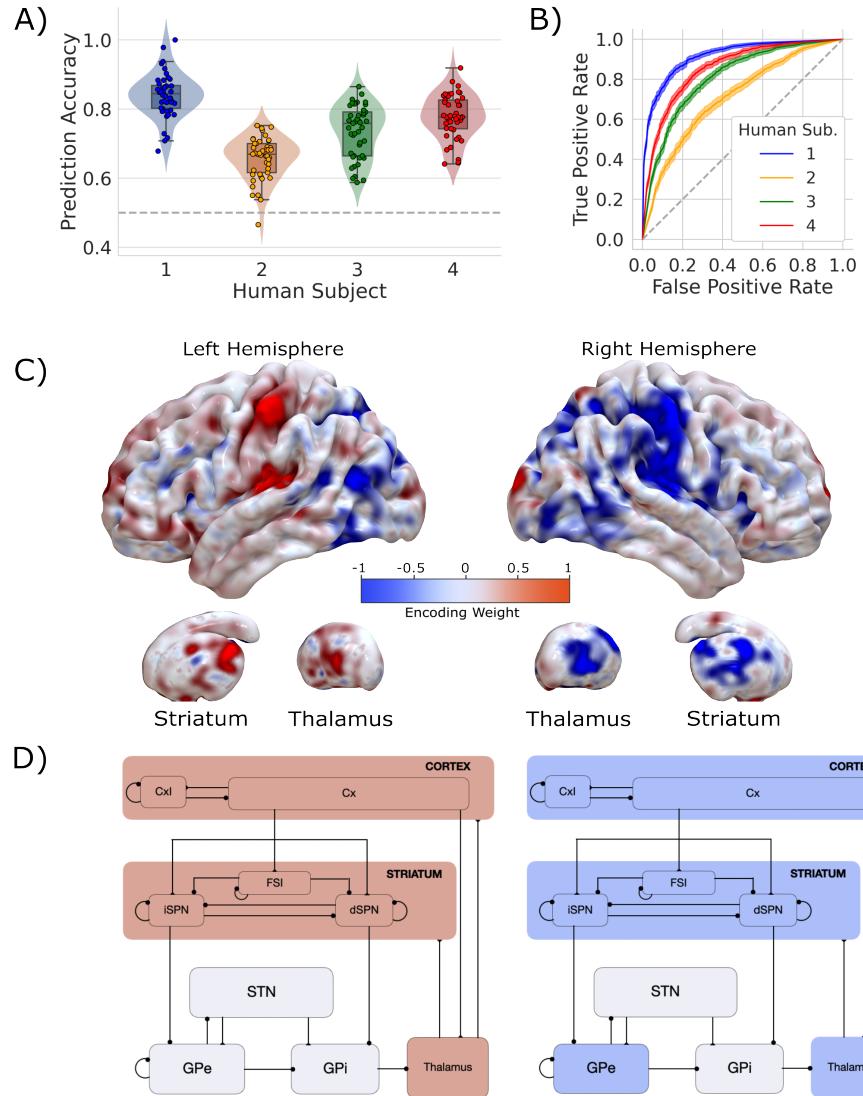
159 association between classifier uncertainty and drift rate. Results for individual participants are  
160 presented along with aggregated results.

161

162 ***Homo sapiens* adapt decision policies in response to change**

163 To test the predictions of our model, a sample of primates (*Homo sapiens*, n=4) played a dy-  
164 namic two-armed bandit task under experimental conditions similar to those used for the sim-  
165 ulated CBGT network and prior behavioral work (7) as whole brain hemophysiological signals  
166 were recorded using functional magnetic resonance imaging (fMRI). On each trial, participants  
167 were presented with a male and female Greeble (44). The goal was to select the Greeble most  
168 likely to give a reward. Selections were made by pressing a button with their left or right hand  
169 to indicate the left or right Greeble on the screen. We collected 2700 trials over 45 runs from  
170 nine separate imaging sessions per participant. Consistent with our within-participant design,  
171 statistical analyses estimated effects on a single-participant basis.

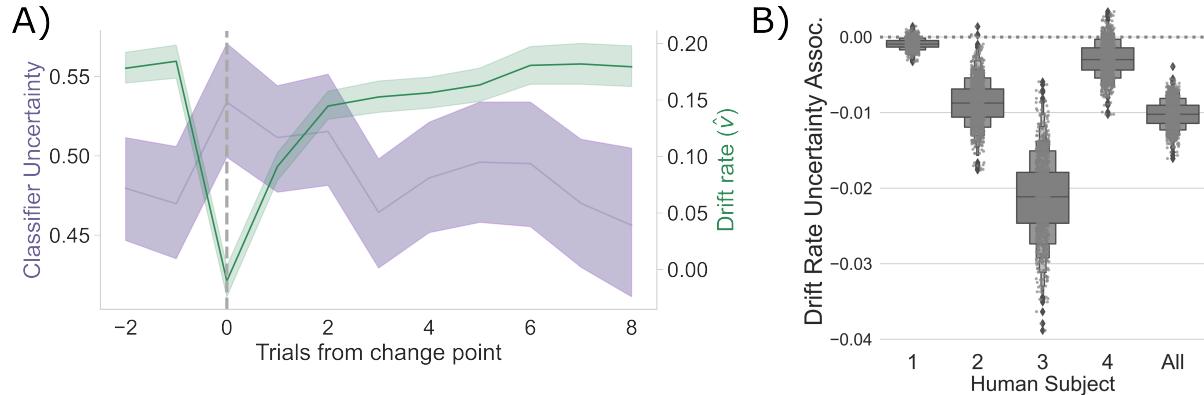
172



173

174 **Fig. 3. Single-trial prediction of action plan competition in humans.** A) Overall classifi-  
175 cation accuracy for single-trial actions for each participant. Each point corresponds to the per-  
176 formance for each cross-validation fold. B) Classification performance for single-trial actions  
177 shown as an ROC curve. The gray dashed line represents chance performance. C) Participant-  
178 averaged encoding weight maps in standard space for both hemispheres. D) The mean encoding  
179 weights within each CGBT node in both hemispheres. See encoding weight scale above for ref-  
180 erence.

181



182

183 **Fig. 4. Competition between action plans drives evidence accumulation in humans.** A)  
184 Classifier uncertainty (lavender) and estimated drift rate ( $\hat{v}$ ; green) dynamics. B) The associa-  
185 tion between classifier uncertainty and drift rate by participant and in aggregate.

186

187 Overall, speed and accuracy across conditions matched what we observed in previous ex-  
188 periments (Experiment 2 in (7)). Specifically, we see a consistent effect of change point on both  
189 RT and accuracy that matches the behavior of our network (Supp. Fig. 2; Supp. Table 1).

190 To address how a change in the environment shifted underlying decision dynamics, we used  
191 a hierarchical DDM modeling approach (42) as we did with the network behavior (see Methods  
192 for details). Given previous empirical work (7) and the results from our CBGT network model  
193 showing that only  $v$  and, less reliably,  $a$  respond to a shift in the environment (7), we focused  
194 our subsequent analysis on these two parameters. Consistent with the predictions from our  
195 CBGT model, we found equivocal fits for the model mapping both  $\Delta B$  to  $v$  and  $\Omega$  to  $a$  and  
196 a simpler model mapping  $\Delta B$  to  $v$  (see Supp. Table 2 for average results). This pattern was  
197 fairly consistent at the participant level, with 3/4 participants showing  $\Delta B$  modulating  $v$  (Supp.  
198 Table 3). These results suggest that as the belief in the value of the optimal choice approaches  
199 the reward value for the optimal choice, the rate of evidence accumulation increases.

200 Taken altogether, we confirm that humans rapidly shift how quickly they accumulate evi-  
201 dence (and, to some degree, how much evidence they need to make a decision) in response to  
202 a change in action-outcome contingencies. This mirrors the decision parameter dynamics pre-  
203 dicted by the CBGT model. We next evaluated how this change in decision policy tracks with  
204 competition in neural action representations.

205 **Measuring action representations in the brain**

206 To measure competition in action representations, we first needed to determine how individual  
207 regions (i.e., voxels) contribute to single decisions. For each participant, trial-wise responses  
208 at every voxel were estimated by means of a general linear model (GLM), with trial modeled  
209 as a separate condition in the design matrix. Therefore, the  $\hat{\beta}_{t,v}$  estimated at voxel  $v$  reflected  
210 the magnitude of the evoked response on trial  $t$ . As in the CBGT model analysis, these whole-  
211 brain, single-trial responses were then submitted to a LASSO-PCR classifier to predict left/right  
212 response choices. The performance of the classifier for each participant was evaluated with a  
213 45-fold cross-validation, iterating through all runs so that each one corresponded to the hold-out  
214 test set for one fold.

215 Our classifier was able to predict single trial responses well above chance for each of the  
216 four participants (Fig. A and B), with mean prediction accuracy ranging from 65% to 83%  
217 (AUCs from 0.72 to 0.92). Thus, as with the CBGT network model, we were able to reliably  
218 predict trial-wise responses for each participant. Fig 3C shows the average encoding map for  
219 our model as an illustration of the influence of each voxel on our model predictions (Supp. Fig.  
220 4 displays individual participant maps). These maps effectively show voxel-tuning towards  
221 rightward (blue) or leftward (red) responses. Qualitatively, we see that cortex, striatum, and  
222 thalamus all exhibit strongly lateralized influences on contralateral response prediction. Indeed,  
223 when we average the encoding weights in terms of principal CBGT nuclei (Fig. 3D), we confirm

224 that these three regions largely predict contralateral responses. Supp. Fig. 5 provides a more  
225 detailed summary of the encoding weights across multiple cortical and subcortical regions.

226 These results show that we can reliably predict single-trial choices from whole-brain hemi-  
227 dynamic responses for individual participants. Further, key regions of the CBGT pathway con-  
228 tribute to these predictions. Next, we set out to determine whether competition between these  
229 representations for left and right actions correlates with changes in the drift rate, as predicted  
230 by the CBGT network model (Fig. 2C).

## 231 **Competition between action representations may drive drift-rate**

232 To evaluate whether competition between action channels correlates with the magnitude of  $v$  on  
233 each trial, as the CBGT network predicts (Fig. 2C), we focused our analysis on trials surround-  
234 ing the change point, following analytical methods identical to those described in the previous  
235 section and shown in Fig. 2C.

236 Consistent with the CBGT network model predictions, following a change point,  $v$  shows  
237 a stereotyped drop and recovery as observed in the CBGT network (Fig. 2C) and prior be-  
238 havioral work (7) (Fig. 4A). This drop in  $v$  tracked with a relative increase in classifier uncer-  
239 tainty, and subsequent recovery, in response to a change in action-outcome contingencies (mean  
240 bootstrapped  $\beta$ :  $-0.021$  to  $-0.001$ ;  $t$  range:  $-3.996$  to  $-1.326$ ;  $p_{S1} = 0.057$ ,  $p_{S2} < 0.001$ ;  
241  $p_{S3} < 0.001$ ;  $p_{S3} = 0.080$ ,  $p_{All} < 0.001$ ). As with the CBGT network simulations (Fig. 2D),  
242 we also observe a consistent negative correlation between  $v$  and classifier uncertainty over all  
243 trials, irrespective of their position to a change point, in each participant and in aggregate (Fig.  
244 4B; Spearman's  $\rho$  range:  $-0.08$  to  $-0.04$ ;  $p$  range:  $< 0.001$  to  $0.043$ ).

245 These results clearly suggest that, as predicted by our CBGT network simulations and prior  
246 work (12, 17, 45), competition between action representations drives changes in the rate of evi-  
247 dence accumulation during decision making in humans.

## 248 Discussion

249 We investigated the underlying mechanisms that drive shifts in decision policies when the rules  
250 of the environment change. We first tested an implementation-level theory of how CBGT net-  
251 works contribute to changes in decision policy parameters. This theory predicted that the rate  
252 of evidence accumulation is driven by competition across action representations. Using a high-  
253 powered within-participants fMRI design conducted with four human primates, wherein each  
254 participant served as an independent replication test, we found evidence consistent with our  
255 CBGT network simulations. Specifically, as action-outcome contingencies change, thereby in-  
256 creasing uncertainty of optimal choice, decision policies shift with a rapid decrease in the rate of  
257 evidence accumulation, followed by a gradual recovery to baseline rates as new contingencies  
258 are learned (see also (7)). These results empirically validate prior theoretical and computa-  
259 tional work predicting that competition between neural populations encoding distinct actions  
260 modulates how information is used to drive a decision (9, 12, 14, 20, 21).

261 Our findings here align with prior work on the role of competition in the regulation of  
262 evidence accumulation. In the decision-making context, the ratio of dSPN to iSPN activation  
263 *within* an action channel has been linked to the drift-rate of single-action decisions (14–16,  
264 37). In the motor control context, this competition manifests as movement vigor (46–48). Yet,  
265 our results show how competition *across* channels drives drift-rate dynamics. So how do we  
266 reconcile these two effects? Mechanistically, the strength of each action channel is defined by  
267 the relative difference between dSPN and iSPN influence. In this way, competition across action  
268 channels is defined by the relative balance of direct and indirect pathway activation within each  
269 channel. Greater direct vs. indirect pathway competition in one action channel, relative to  
270 another, makes that action decision relatively slow and reduces the overall likelihood that it is  
271 selected. This mechanism is consistent with prior theoretical (12, 45) and empirical work (18).

272 While our current work postulates a mechanism by which changes in action-outcome con-  
273 tingencies drive changes in evidence accumulation through plasticity within the CBGT circuits,  
274 the results presented here are far from conclusive. For example, our model of the underlying  
275 neural dynamics predicts that the certainty of individual action representations is encoded by  
276 the competition between direct and indirect pathways (see also (12, 38, 45)). Thus, external  
277 perturbation of dSPN (or iSPN) firing, say with optogenetic methods, during decision-making  
278 should causally impact the evidence accumulation rate and, subsequently, the speed (or slow)  
279 the speed at which the new action-outcome contingencies are learned. Indeed, there is already  
280 some evidence for this outcome (see (18), but also (49) for contrastive evidence).

281 Our model, however, has very specific predictions with regards to disruptions of each path-  
282 way within an action representation. Disrupting the *balance* of dSPN and iSPN efficacy should  
283 selectively impact the drift-rate (and, to a degree, onset bias; see (45)), while non-specific dis-  
284 ruption of global iSPN efficacy across action representations should selectively disrupt bound-  
285 ary height (and, to a degree, accumulation onset time; see again (45)).

286 Thus, increasing the difference between dSPN and iSPN firing in the channel representing  
287 the new optimal-action, say by selective excitation of the relevant dSPNs, should speed up the  
288 time to resolve the credit assignment problem during learning. This would result in faster and  
289 more accurate learning following an environmental change and lead to characteristic signatures  
290 in the distribution of reaction times, as well as choice probabilities, reflective of a shift in evi-  
291 dence accumulation rate. Of course, testing these predictions is left to future work.

## 292 Conclusion

293 As the world changes and certain actions become less optimal, successful behavioral adapta-  
294 tion requires flexibly changing how sensory evidence drives decisions. Our simulations and  
295 hemophysiological experiments in human primates show how this process can occur within the

296 CBGT circuits. Here, a shift in action-outcome contingencies induces competition between en-  
297 coded action plans by modifying the relative balance of direct and indirect pathway activity in  
298 CBGT circuits, both within and between action channels, slowing the rate of evidence accumu-  
299 lation to promote adaptive exploration. If the environment subsequently remains stable, then  
300 this learning process accelerates the rate of evidence accumulation for the optimal decision by  
301 increasing the strength of action representations for the new optimal choice. This highlights  
302 how these macroscopic systems promote flexible, effective decision-making under dynamic en-  
303 vironmental conditions.

304 **References and Notes**

- 305 1. J. I. Gold, M. N. Shadlen, *Annu. Rev. Neurosci.* **30**, 535 (2007).
- 306 2. R. Ratcliff, *Psychological review* **85**, 59 (1978).
- 307 3. M. R. Nassar, R. C. Wilson, B. Heasly, J. I. Gold, *Journal of Neuroscience* **30**, 12366  
308 (2010).
- 309 4. R. C. Wilson, Y. Niv, *Frontiers in human neuroscience* **5**, 189 (2012).
- 310 5. M. R. Nassar, *et al.*, *Nature neuroscience* **15**, 1040 (2012).
- 311 6. T. E. Behrens, M. W. Woolrich, M. E. Walton, M. F. Rushworth, *Nature neuroscience* **10**,  
312 1214 (2007).
- 313 7. K. Bond, K. Dunovan, A. Porter, J. E. Rubin, T. Verstynen, *Elife* **10**, e65540 (2021).
- 314 8. A. E. Urai, J. W. De Gee, T. H. Donner, *BioRxiv* p. 251595 (2018).
- 315 9. R. Ratcliff, M. J. Frank, *Neural computation* **24**, 1186 (2012).

316 10. M. L. Pedersen, M. J. Frank, G. Biele, *Psychonomic bulletin & review* **24**, 1234 (2017).

317 11. K. Dunovan, T. Verstynen, *Journal of Neuroscience* **39**, 2251 (2019).

318 12. K. Dunovan, C. Vich, M. Clapp, T. Verstynen, J. Rubin, *PLoS computational biology* **15**,  
319 e1006998 (2019).

320 13. A. G. Mendonça, *et al.*, *Nature Communications* **11**, 1 (2020).

321 14. K. Dunovan, T. Verstynen, *Frontiers in neuroscience* **10**, 106 (2016).

322 15. S. Bariselli, W. Fobbs, M. Creed, A. Kravitz, *Brain research* (2018).

323 16. J. G. Mikhael, R. Bogacz, *PLoS Comput. Biol.* **12**, e1005062 (2016).

324 17. J. E. Rubin, C. Vich, M. Clapp, K. Noneman, T. Verstynen, *European Journal of Neuro-*  
325 *science* (2020).

326 18. M. M. Yartsev, T. D. Hanks, A. M. Yoon, C. D. Brody, *Elife* **7**, e34929 (2018).

327 19. C. Vich, M. Clapp, T. Verstynen, J. E. Rubin, *bioRxiv* (2022).

328 20. R. Bogacz, K. Gurney, *Neural computation* **19**, 442 (2007).

329 21. R. Bogacz, E.-J. Wagenmakers, B. U. Forstmann, S. Nieuwenhuis, *Trends in neurosciences*  
330 **33**, 10 (2010).

331 22. A. Gupta, *et al.*, *Frontiers in Computational Neuroscience* **15** (2021).

332 23. R. L. Albin, A. B. Young, J. B. Penney, *Trends in Neurosciences* **18**, 63 (1995).

333 24. D. M. Friend, A. V. Kravitz, *Trends in neurosciences* **37**, 301 (2014).

334 25. J. W. Mink, *Prog. Neurobiol.* **50**, 381 (1996).

335 26. G. E. Alexander, M. R. DeLong, P. L. Strick, *Annual review of neuroscience* **9**, 357 (1986).

336 27. A. Adler, I. Finkes, S. Katabi, Y. Prut, H. Bergman, *The Journal of neuroscience : the*  
337 *official journal of the Society for Neuroscience* **33**, 4854 (2013).

338 28. A. Klaus, *et al.*, *Submitted* **95**, 1171 (2017).

339 29. G. Barbera, *et al.*, *Neuron* **92**, 202 (2016).

340 30. L. Carrillo-Reid, S. Hernandez-Lopez, D. Tapia, E. Galarraga, J. Bargas, *Journal of Neu-*  
341 *roscience* **31**, 14972 (2011).

342 31. N. Badreddine, *et al.*, *Cell Reports* **39** (2022).

343 32. A. V. Kravitz, L. D. Tye, A. C. Kreitzer, *Nat. Neurosci.* **15**, 816 (2012).

344 33. K. Gurney, T. J. Prescott, P. Redgrave, *Biological cybernetics* **84**, 411 (2001).

345 34. J. W. Mink, *Progress in neurobiology* **50**, 381 (1996).

346 35. H. Kitano, I. Tanibuchi, K. Jinnai, *Brain Research* **784**, 228 (1998).

347 36. N. N. Foster, *et al.*, *Nature* **598**, 188 (2021).

348 37. K. Dunovan, B. Lynch, T. Molesworth, T. Verstynen, *Elife* **4**, e08723 (2015).

349 38. C. Vich, K. Dunovan, T. Verstynen, J. Rubin, *Communications in Nonlinear Science and*  
350 *Numerical Simulation* **82**, 105048 (2020).

351 39. E. A. Yttri, J. T. Dudman, *Nature* **533**, 1 (2016).

352 40. B. U. Forstmann, *et al.*, *Proceedings of the National Academy of Sciences* **107**, 15916  
353 (2010).

354 41. T. V. Maia, M. J. Frank, *Nature neuroscience* **14**, 154 (2011).

355 42. T. V. Wiecki, I. Sofer, M. J. Frank, *Frontiers in neuroinformatics* **7**, 14 (2013).

356 43. K. P. Burnham, D. R. Anderson, *Model selection and inference* (Springer, 1998), pp. 75–  
357 117.

358 44. I. Gauthier, M. J. Tarr, *Vision research* **37**, 1673 (1997).

359 45. C. Vich, M. Clapp, T. Verstynen, J. Rubin (2021).

360 46. E. A. Yttri, J. T. Dudman, *Nature* **533**, 402 (2016).

361 47. J. T. Dudman, J. W. Krakauer, *Current Opinion in Neurobiology* **37**, 158 (2016).

362 48. R. S. Turner, M. Desmurget, *Current opinion in neurobiology* **20**, 704 (2010).

363 49. L. Ding, J. I. Gold, *Journal of Neuroscience* **30**, 15747 (2010).

364 50. W. Wei, J. E. Rubin, X.-J. Wang, *Journal of Neuroscience* **35**, 4052 (2015).

365 51. L. Hurwicz, *Decision and Organization* p. 320 (1972).

366 52. J. O. Ledyard, *Allocation, Information and Markets* (Springer, 1989), pp. 141–151.

367 53. T. V. Wiecki, I. Sofer, M. J. Frank, *Frontiers in neuroinformatics* **7**, 14 (2013).

368 54. D. J. Spiegelhalter, N. G. Best, B. P. Carlin, A. Van Der Linde, *Journal of the royal statis-  
369 tical society: Series b (statistical methodology)* **64**, 583 (2002).

370 55. O. Esteban, *et al.*, *Nature Methods* **16**, 111 (2018).

371 56. J. Diedrichsen, R. Shadmehr, *NeuroImage* **27**, 624 (2005).

372 57. S. Haufe, *et al.*, *NeuroImage* **87**, 96 (2014).

## 373 **Acknowledgements**

374 We thank all members of the CoAx Lab and collaborators for their feedback on the development  
375 of this work.

## 376 **Funding**

377 Air Force Research Laboratory, Grant Office ID: 180119.

## 378 **Author Contributions**

379 K.B.: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project  
380 administration, Software, Visualization, Writing - original draft, Writing - review and editing;

381 JR: Formal analysis, Software, Visualization, Writing - original draft, Writing - review and edit-  
382 ing;

383 RM: Data curation, Formal analysis, Software, Writing - review and editing;

384 JB: Formal analysis, Software, Writing - review and editing;

385 JR: Conceptualization, Writing - review and editing;

386 TV: Conceptualization, Formal analysis, Funding acquisition, Investigation, Project administra-  
387 tion, Resources, Supervision, Validation, Writing - review and editing

## 388 **Competing Interests**

389 The authors declare no competing interests.

## 390 **Data and Materials Availability**

391 Behavioral data and computational derivatives are publically available here. Raw and prepro-  
392 cessed hemodynamic data, in addition to physiological measurements collected for quality con-

393 trol, are available here.

## 394 **Supplementary Materials:**

395 Materials and Methods Supplementary Text

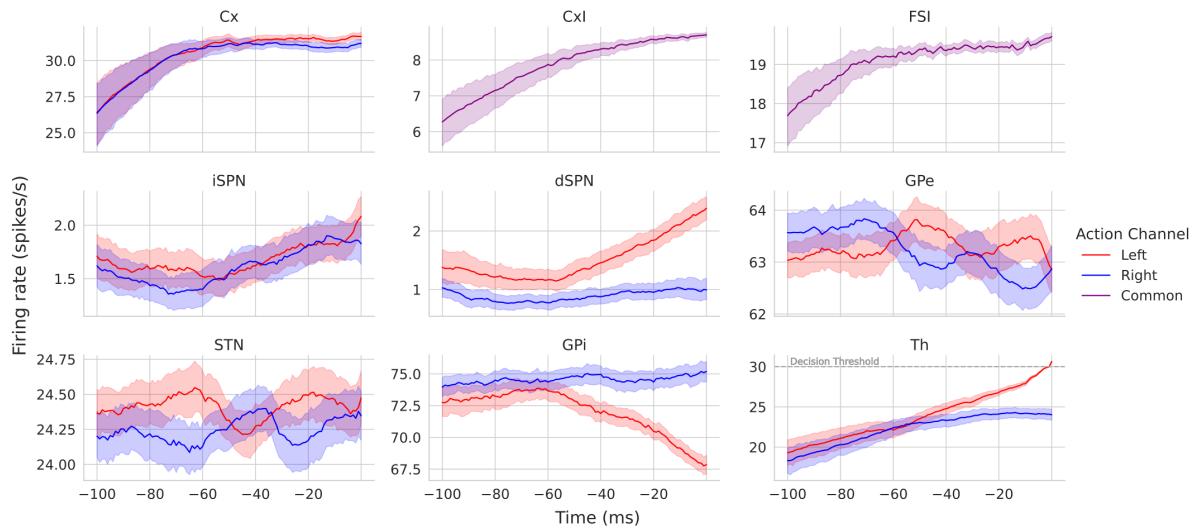
396 Figs. S1 to S5

397 Tables S1 to S8

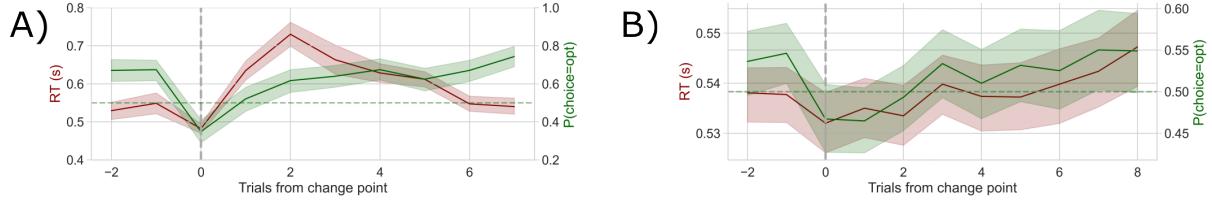
398 Supp. References (0)

399 No supplementary data included (see Data Availability section above)

## 400 **Figures**



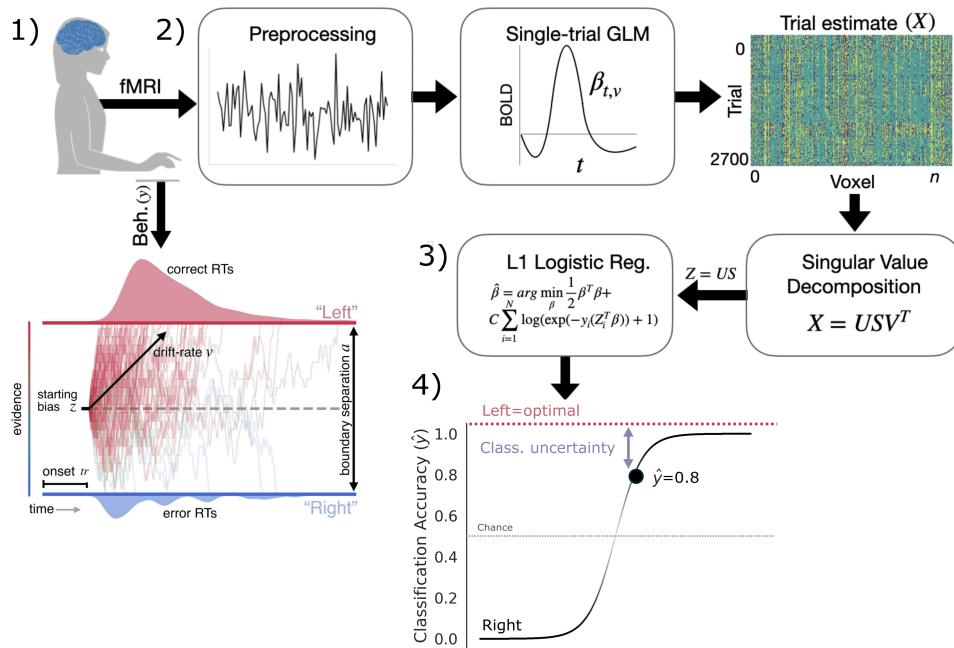
402 **Supp. Fig. 1. Simulated CBGT nuclei firing rates for a left decision.** Each panel shows the  
403 firing rates for each CBGT nucleus 100 ms prior to a left decision. The decision threshold for  
404 thalamus (30 spikes/second) is marked with a horizontal gray line. Note that the y axes have  
405 different limits for each nucleus due to differences of scale in their firing rates.



406

407 **Supp. Fig. 2. Simulated and human behavior.** Change point evoked reaction times are shown  
 408 in red and accuracy, or the probability of selecting the optimally rewarding choice, is shown in  
 409 green. Chance is marked as a green horizontal dashed line. The change point is marked by the  
 410 vertical gray line. A) Simulated behavior. B) Human behavior.

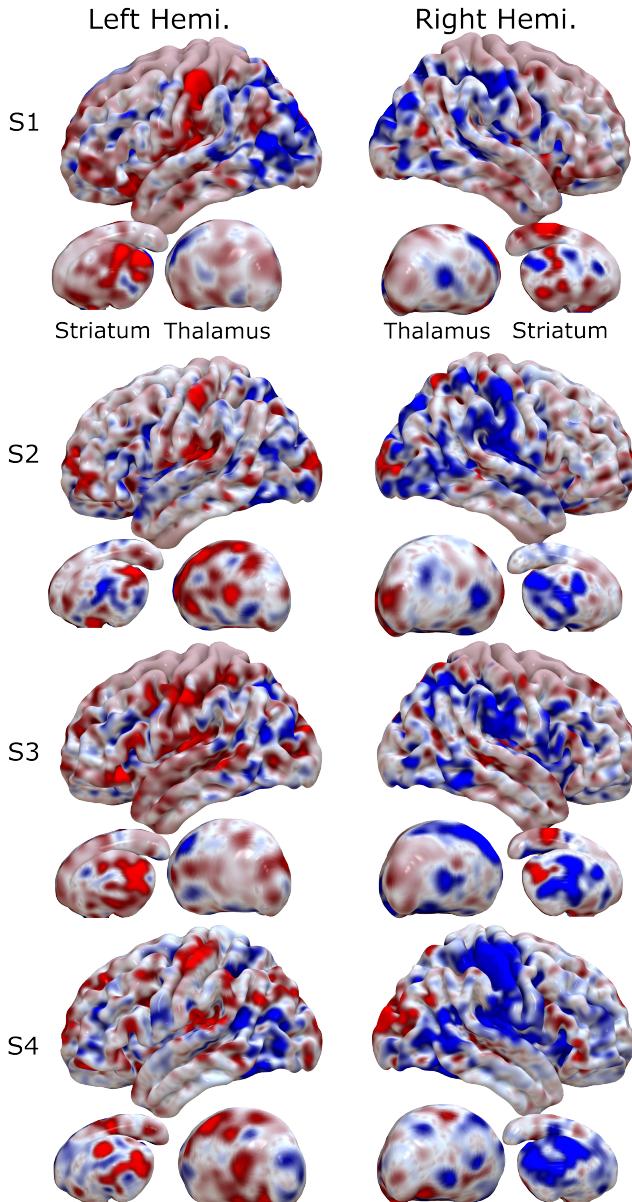
411



412

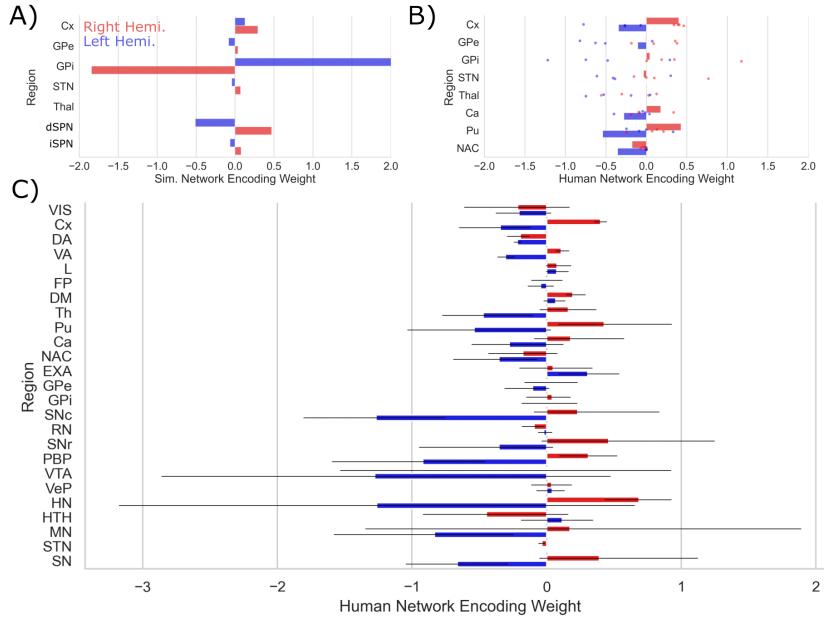
413 **Supp. Fig. 3. Analysis method.** Step 1. Behavioral response collection and DDM (Drift  
 414 Diffusion Model) parameter estimation. In the case of the simulated CBGT network, this step  
 415 involved simulating responses to experimental manipulations. Step 2. Preprocessing and single-  
 416 trial estimates of the hemodynamic response. Step 3. Singular Value Decomposition and Logis-  
 417 tic regression with an L1 penalty. After crossvalidation, this outputs a predicted response (left

418 or right), here coded as 0 or 1. Step 4. Calculating classifier uncertainty from cross-validated  
419 response prediction. The further the predicted response from the inflection point of the logis-  
420 tic function, the more certain the prediction. The distance of this predicted response from the  
421 optimal choice represents classifier uncertainty for each trial. Here, the predicted probability of  
422 a left response  $\hat{y}_{t1}$  is 0.2. The distance from the optimal choice on this trial, and, thereby, the  
423 classifier uncertainty is 0.2.



424

425 **Supp. Fig. 4. Encoding maps in standardized space for each participant.** Rows represent  
426 individual participants. Columns refer to left and right views of the whole brain. Thalamus and  
427 striatum are shown beneath each cortical map. Values are z-scored.



428

429 **Supp. Fig. 5. Encoding patterns by CBGT node.** A) Simulated CBGT encoding weights.  
430 B) Human CBGT encoding weights for comparison with the simulated CBGT network re-  
431 sults. Each point represents the average result for each participant. Bars represent participant-  
432 averaged data. C) The full set of human CBGT encoding weights for all captured nodes from  
433 whole-brain imaging. Gray error bars represent 95% CIs over participants. Left hemisphere  
434 weights are marked in blue and right hemisphere weights are marked in red.

435 **Tables**

**Table S1. Behavior**

<b>Simulated</b>			<b>Human</b>		
Sim. Part.	RT(s)	Accuracy	Human Part.	RT (s)	Accuracy
1	0.604	0.59	1	0.553	0.538
2	0.559	0.624	2	0.537	0.541
3	0.608	0.61	3	0.531	0.553
4	0.596	0.648	4	0.54	0.511
All	0.592 ± 0.176	0.618	All	0.540 ± 0.076	0.536

436

437 **Supp. Table 1. Behavior.** Simulated and human reaction time (in seconds) and accuracy.

**Table S2. Model fits**

**Simulated**

	$\Delta B$	$\Omega$	$\Delta DIC_{\text{null}}$	$\Delta DIC_{\text{best}}$
I	v	a	$-29.85 \pm 12.76$	$-4.49 \pm 5.91$
II	a	v	$-23.94 \pm 22.56$	$-10.40 \pm 11.22$
III	-	v	$-6.16 \pm 4.24$	$-28.19 \pm 13.62$
IV	v	-	$-22.60 \pm 7.28$	$-11.74 \pm 14.80$
V	-	a	$-7.04 \pm 11.06$	$-27.30 \pm 8.16$
VI	a	-	$-17.72 \pm 21.49$	$-16.62 \pm 11.88$
VII	-	-	$0.00 \pm 0.00$	$-34.34 \pm 15.97$

**Human**

	$\Delta B$	$\Omega$	$\Delta DIC_{\text{null}}$	$\Delta DIC_{\text{best}}$
I	v	a	$-14.90 \pm 20.58$	$-1.52 \pm 1.04$
II	a	v	$-0.44 \pm 1.11$	$-15.99 \pm 18.56$
III	-	v	$-1.47 \pm 1.30$	$-14.96 \pm 18.56$
IV	v	-	$-13.80 \pm 16.61$	$-2.63 \pm 3.62$
V	-	a	$-1.03 \pm 4.46$	$-15.40 \pm 15.60$
VI	a	-	$1.00 \pm 0.71$	$-17.42 \pm 19.52$
VII	-	-	$0.00 \pm 0.00$	$-16.43 \pm 19.53$

438

439 **Supp. Table 2. Model fits.** Deviance Information Criterion (DIC) values for regression models

440 tested.

**Table S3. Human model fits by participant**

Part.	$\Delta B$	$\Omega$	$\Delta DIC_{null}$	$\Delta DIC_{best}$	
I	1	v	a	0.61	-2.32
II	1	a	v	0.08	-1.79
III	1	-	v	-1.71	0.00
IV	1	v	-	1.13	-2.84
V	1	-	a	-0.36	-1.35
VI	1	a	-	1.93	-3.64
VII	1	-	-	0.00	-1.71
I	2	v	a	-9.91	-1.73
II	2	a	v	-0.69	-10.95
III	2	-	v	-1.17	-10.47
IV	2	v	-	-11.64	0.00
V	2	-	a	1.89	-13.52
VI	2	a	-	0.46	-12.10
VII	2	-	-	0.00	-11.64
I	3	v	a	-45.08	0.00
II	3	a	v	-1.85	-43.23
III	3	-	v	-3.07	-42.01
IV	3	v	-	-37.41	-7.68
V	3	-	a	-7.53	-37.55
VI	3	a	-	1.16	-46.25
VII	3	-	-	0.00	-45.08
I	4	v	a	-5.23	-2.05
II	4	a	v	0.71	-7.99
III	4	-	v	0.07	-7.35
IV	4	v	-	-7.28	0.00
V	4	-	a	1.90	-9.18
VI	4	a	-	0.43	-7.70
VII	4	-	-	0.00	-7.28

441

442 Supp Table 3. Human model fits by participant.

## 443 Materials & Methods

### 444 Simulations

445 We simulated neural dynamics and behavior using a biologically based, spiking cortico-basal  
446 ganglia-thalamic (CBGT) network model (11, 19). The network representing the CBGT circuit  
447 is composed of 9 neural populations: cortical interneurons (CxI), excitatory cortical neurons  
448 (Cx), striatal D1/D2-spiny projection neurons (dSPNs/iSPNs), striatal fast-spiking interneurons  
449 (FSI), the internal (GPi) and external globus pallidus (GPe), the subthalamic nucleus (STN), and  
450 the thalamus (Th). All the neuronal populations are segregated into two action channels with the  
451 exception of cortical (CxI) and striatal interneurons (FSIs). Each neuron in the population was  
452 modeled with an integrate-fire-or-burst-model (50), and a conductance-based synapse model  
453 was used for NMDA, AMPA and GABA receptors. The neuronal and network parameters  
454 (inter-nuclei connectivity and synaptic strengths) were tuned to obtain realistic baseline firing  
455 rates for all the nuclei. The details of the model are described in our previous work (19) as well  
456 as in the appendix for the sake of completeness.

457 Corticostriatal weights for D1 and D2 neurons in striatum were modulated by phasic dopamine  
458 to model the influence of reinforcement learning on network dynamics. The details of STDP  
459 learning are described in detail in previous work (38), but key details are shown below. As a  
460 result of these features of the CBGT network, it was capable of learning under realistic experi-  
461 mental paradigms with probabilistic reinforcement schemes (i.e. under reward probabilities and  
462 unstable action-outcome values).

### 463 Threshold for CBGT network decisions

464 A decision between the two competing actions (“left” and “right”) was considered to be made  
465 when either of the thalamic subpopulations reached a threshold of 30Hz. This threshold was  
466 set based on the network dynamics for the chosen parameters with a aim to obtain realistic

467 reaction times. The maximum time allowed to reach a decision was 1000ms. If none of the  
468 thalamic subpopulations reach the threshold of 30Hz, no action was considered to be taken.  
469 Such trials were dropped from further analysis. Reaction/decision times were calculated as time  
470 from stimulus onset to decision (either subpopulation reaches the threshold). The "slow" and  
471 "fast" trials were categorized as reaction times  $\geq$  75th percentile (314.5ms) and reactions time  
472  $<$  50th percentile (196.0ms), respectively, of the reaction time distributions. The firing rates  
473 of the CBGT nuclei during the reaction times were used for prediction analysis as discussed in  
474 Section .

#### 475 **Corticostriatal weight plasticity**

476 The corticostriatal weights are modified by a dopamine-mediated STDP rule, where the pha-  
477 sic dopamine is modulated by reward prediction error. The internal estimate of the reward is  
478 calculated at every trial by a Q-learning algorithm which is subtracted from the reward associ-  
479 ated with the experimental paradigm to yield a trial-by-trial estimate of the reward prediction  
480 error. The effect of dopaminergic release is receptor dependent; a rise in dopamine promotes  
481 potentiation for D1-SPNs and depression for D2-SPNs. The degree of change in the weights is  
482 dependent on an eligibility trace which is proportional to the coincidental pre-synaptic (cortical)  
483 and post-synaptic (striatal) firing rates. The STDP rule is described in detail in (38) as well as  
484 in the appendix.

#### 485 **In silico experimental design**

486 We follow the paradigm of a 2 arm bandit task, where the CBGT network learns to consistently  
487 choose the rewarded action until the block changes (i.e the reward contingencies switch), at  
488 which point the CBGT network re-learns the rewarded action (reversal learning). Each session  
489 consists of 40 trials with a block change every 10 trials. The reward probabilities represent a  
490 conflict of (75%, 25%); that is, in a left block, 75% of the left actions are rewarded, whereas

491 25% of the right actions are rewarded. The inter-trial-interval in network time is fixed to 600ms.

492 To maximize the similarity between the CBGT network simulations and our human data, we  
493 randomly varied the initialization of the network such that neurons with a specific connection  
494 probability were randomly chosen for each simulated subject, with the background input to the  
495 nuclei for each simulated subject as a mean-reverting random walk (noise was drawn from the  
496 normal distribution  $N(0,1)$ ). These means are listed in Supp. Table 4.

## 497 **Participants**

498 Four neurologically healthy adult human primates (two female, all right-handed, 29-34 years  
499 old) were recruited and paid \$30 per session, in addition to a performance bonus and a bonus  
500 for completing all nine sessions. These participants were recruited from the local university  
501 population.

502 All procedures were approved by the Carnegie Mellon University Institutional Review Board.  
503 All research participants provided informed consent to participate in the study and consent to  
504 publish any research findings based on their provided data.

## 505 **Experimental design**

506 The experiment used male and female Greebles (44) as selection targets. Participants were  
507 first trained to discriminate between male and female Greebles to prevent errors in perceptual  
508 discrimination from interfering with selection on the basis of value. Using a two-alternative  
509 forced choice task, participants were presented with a male and female Greeble and asked to  
510 select the female, with the male and female Greeble identities resampled on each trial. Partic-  
511 ipants received binary feedback regarding their selection (correct or incorrect). This criterion  
512 task ended after participants reached 95% accuracy. After reaching perceptual discrimination  
513 criterion for each session, each participant was tested under nine reinforcement learning con-  
514 ditions composed of 300 trials each, generating 2700 trials per participant in total. Data were

515 collected from four participants in accordance with a replication-based design, with each partic-  
516 ipant serving as a replication experiment. Participants completed these sessions in randomized  
517 order. Each learning trial presented a male and female Greeble (44), with the goal of selecting  
518 the gender identity of the Greeble that was most rewarding. Because individual Greeble identi-  
519 ties were resampled on each trial, the task of the participant was to choose the gender identity  
520 rather than the individual identity of the Greeble which was most rewarding.

521 Probabilistic reward feedback was given in the form of points drawn from the normal dis-  
522 tribution  $\mathcal{N}(\mu = 3, \sigma = 1)$  and converted to an integer. These points were displayed at the  
523 center of the screen. For each run, participants began with 60 points and lost one point for  
524 each incorrect decision. To promote incentive compatibility (51, 52), participants earned a cent  
525 for every point earned. Reaction time was constrained such that participants were required to  
526 respond within between 0.1 s and 0.75 s from stimulus presentation. If participants responded  
527 in  $\leq 0.1$  s,  $\geq 0.75$  s, or failed to respond altogether, the point total turned red and decreased by  
528 5 points. Each trial lasted 1.5 s and reward feedback for a given trial was displayed from the  
529 time of the participant's response to the end of the trial. To manipulate change point probability,  
530 the gender identity of the most rewarding Greeble was switched probabilistically, with a change  
531 occurring every 10, 20, or 30 trials, on average. To manipulate the belief in the value of the  
532 optimal target, the probability of reward for the optimal target was manipulated, with  $P$  set to  
533 0.65, 0.75, or 0.85. Each session combined one value of  $P$  with one level of volatility, such  
534 that all combinations of change point frequency and reward probability were imposed across  
535 the nine sessions. Finally, the position of the high-value target was pseudo-randomized on each  
536 trial to prevent prepotent response selections on the basis of location.

537 **Behavioral analysis**

538 Statistical analyses and data visualization were conducted using custom scripts written in R (R  
539 Foundation for Statistical Computing, version 3.4.3) and Python (Python Software Foundation,  
540 version 3.5.5). Binary accuracy data were submitted to a mixed effects logistic regression anal-  
541 ysis with either the degree of conflict (the probability of reward for the optimal target) or the  
542 degree of volatility (mean change point frequency) as predictors. The resulting log-likelihood  
543 estimates were transformed to likelihood for interpretability. RT data were log-transformed and  
544 submitted to a mixed effects linear regression analysis with the same predictors as in the previ-  
545 ous analysis. To determine if participants used ideal observer estimates to update their behavior,  
546 two more mixed effects regression analyses were performed. Estimates of change point proba-  
547 bility and the belief in the value of the optimal target served as predictors of reaction time and  
548 accuracy across groups. As before, we used a mixed logistic regression for accuracy data and a  
549 mixed linear regression for reaction time data.

550 **Estimating evidence accumulation using drift diffusion modeling**

551 To assess whether and how much the ideal observer estimates of change point probability ( $\Omega$ )  
552 and the belief in the value of the optimal target ( $\Delta B$ ) (3, 7) updated the rate of evidence accu-  
553 mulation ( $v$ ), we regressed the change-point-evoked ideal observer estimates onto the decision  
554 parameters using hierarchical drift diffusion model (HDDM) regression (53). These ideal ob-  
555 server estimates of environmental uncertainty served as a more direct and continuous measure  
556 of the uncertainty we sought to induce with our experimental manipulations. Using this more  
557 direct approach, we pooled change point probability and belief across all conditions and used  
558 these values as our predictors of drift rate and boundary height. Responses were accuracy-  
559 coded, and the belief in the difference between targets values was transformed to the belief in  
560 the value of the optimal target ( $\Delta B_{\text{optimal}(t)} = B_{\text{optimal}(t)} - B_{\text{suboptimal}(t)}$ ). This approach allowed

561 us to estimate trial-by-trial covariation between the ideal observer estimates and the decision  
562 parameters.

To find the models that best fit the observed data, we conducted a model selection process using Deviance Information Criterion (DIC) scores. A lower DIC score indicates a model that loses less information. Here, a difference of two points from the lowest-scoring model cannot rule out the higher scoring model; a difference of three to seven points suggests that the higher scoring model has considerably less support; and a difference of 10 points suggests essentially no support for the higher scoring model (43, 54). We evaluated the DIC scores for the set of fitted models relative to an intercept-only regression model ( $DIC_{intercept} - DIC_{model_i}$ ).

563 **MRI Data Acquisition**

564 Neurologically healthy human participants (N=4, 2 female) were recruited. Each participant  
565 was tested in nine separate imaging sessions using a 3T Siemens Prisma scanner. Session 1  
566 included a set of anatomical and functional localizer sequences (e.g., visual presentation of  
567 Greeble stimuli with no manual responses, and left vs. right button responses to identify motor  
568 networks). Sessions 2-10 collected five functional runs of the dynamic 2-armed bandit task  
569 (60 trials per run). Male and female "greebles" served as the visual stimuli for the selection  
570 targets (44), with each presented on one side of a central fixation cross. Participants were  
571 trained to respond within 1.5 seconds.

572 To minimize the convolution of the hemodynamic response from trial to trial, inter-trial  
573 intervals were sampled according to a truncated exponential distribution with a minimum of 4  
574 s between trials, a maximum of 16 s, and a rate parameter of 2.8 s. To ensure that head position  
575 was stabilized and stable over sessions, a CaseForge head case was customized and printed for  
576 each participant. The task-evoked hemodynamic response was measured using a high spatial  
577 (2mm<sup>3</sup> voxels) and high temporal (750ms TR) resolution echo planar imaging approach. This

578 design maximized recovery of single-trial evoked BOLD responses in subcortical areas, as well  
579 as cortical areas with higher signal-to-noise ratios. During each functional run, eye-tracking  
580 (EyeLink, SR Research Inc.), physiological signals (ECG, respiration, and pulse-oximetry via  
581 the Siemens PMU system) were also collected for tracking attention and for artifact removal.

## 582 **Preprocessing**

583 fMRI data were preprocessed using the default pipeline of fMRIprep (55), a standard toolbox  
584 for fMRI data preprocessing that provides stability to variations in scan acquisition protocols, a  
585 minimal user manipulation, and easily interpretable, comprehensive output results reporting.

## 586 **Single-trial response estimation**

587 By means of a univariate general linear model (GLM) within participant trial-wise responses at  
588 the voxel-level were estimated. Specifically, for each fMRI run preprocessed BOLD time series  
589 were regressed onto a design matrix, where each task trial corresponded to a different column,  
590 and was modeled using a boxcar function convolved with the default hemodynamic response  
591 function given in SPM12. Thus, each column in the design matrix estimated the average BOLD  
592 activity within each trial. In order to account for head motion, the six realignment parameters (3  
593 rotations, 3 translations) were included as covariates. In addition, a high-pass filter (128 s) was  
594 applied to remove low-frequency artifacts. Parameter and error variance were estimated using  
595 the RobustWLS toolbox, which adjusts for further artifacts in the data by inversely weighting  
596 each observation according to its spatial noise (56).

597 Finally, estimated trial-wise responses were concatenated across runs and sessions and then  
598 stacked across voxels to give a matrix,  $\hat{\beta}_{t,v}$ , of T (trial estimations) x V (voxels) for each partic-  
599 ipant.

600 **Single-trial response prediction**

601 A machine learning approach was applied to predict left/right greeble choices from the trial-  
602 wise responses. First, using the trial-wise hemodynamic responses, we estimated the contrast  
603 in neural activation when the participant made a left versus right selection. A Lasso-PCR clas-  
604 sifier (i.e. an L1-constrained principal component logistic regression) was estimated for each  
605 participant according to the below procedure.

606 First, a singular value decomposition (SVD) was applied to the input matrix  $X$ :

$$X = USV^T, \quad (1)$$

607 where the product matrix  $Z = US$  represents the principal component scores, i.e. the projected  
608 values of  $X$  into the principal component space, and  $V^T$  an orthogonal matrix whose rows are  
609 the principal directions in feature space. Then the binary response variable  $y$  (Left/Right choice)  
610 was regressed onto  $Z$ , where the estimation of the  $\beta$  coefficients is participant to a L1 penalty  
611 term  $C$  in the objective function:

$$\hat{\beta} = \arg \min_{\beta} \frac{1}{2} \beta^T \beta + C \sum_{i=1}^N \log(\exp(-y_i(Z_i^T \beta)) + 1), \quad (2)$$

612 where  $\beta$  and  $Z$  include the intercept term,  $y_i = \{-1, 1\}$  and  $N$  is the number of observations.  
613 Projection of the estimated  $\hat{\beta}$  coefficients back to the original feature (voxel) space was done to  
614 yield a weight map  $\hat{w} = V\hat{\beta}$ , which in turn was used to generate final predictions  $\hat{y}$ :

$$\hat{y} = \frac{1 - e^{-x \cdot \hat{w}}}{1 + e^{-x \cdot \hat{w}}}, \quad (3)$$

615 where  $x$  denotes the vector of voxel-wise responses for a given trial (i.e. a given row in the  
616  $X$  matrix). When visualizing the resulting weight maps, these were further transformed to  
617 encoded brain patterns. This step was performed to aid in correct interpretation in terms of

618 the studied brain process, because doing this directly from the observed weights in multivariate  
619 classification (and regression) models can be problematic (57).

620 Here, the competition between left-right neural responses decreases classifier decoding ac-  
621 curacy, as neural activation associated with these actions becomes less separable. Therefore,  
622 classifier prediction serves as a proxy for response competition. To quantify uncertainty from  
623 this, we calculated the Euclidean distance of these decoded responses  $\hat{y}$  from the statistically  
624 optimal choice on a given trial, *opt\_choice*. This yielded a trial-wise uncertainty metric derived  
625 from the decoded competition between neural responses.

$$\hat{U} = d(\hat{y}, \text{opt\_choice}). \quad (4)$$

626 The same analytical pipeline was used to calculate single trial responses for simulated data  
627 with a difference that trial-wise average firing rates of all nuclei from the simulations were used  
628 instead of fMRI hemodynamic responses.

629 **Neuron model**

630 We used integrate-and-fire-or-burst model that models the membrane potential  $V(t)$  as

$$C \frac{dV}{dt} = -g_L(V(t) - V_L) - g_T h(t) H(V(t) - V_h)(V(t) - V_T) - I_{syn}(t) - I_{ext}(t) \quad (5)$$

$$\frac{dh}{dt} = \begin{cases} \frac{-h(t)}{\tau_h^-} & , \text{when } V(t) \geq V_h \\ \frac{(1-h(t))}{\tau_h^+} & , \text{when } V(t) < V_h \end{cases}$$

631 where  $g_L$  represents the leak conductance,  $V_L$  is the leak reversal potential and the first term  
632  $g_L(V(t) - V_L)$  is the leak current; a low threshold  $Ca^{2+}$  current with maximum conductance as  
633  $g_T$ , gating variable  $h(t)$ , a heaviside function  $H$ , reversal potential  $V_T$ ;  $I_{syn}$  is the synaptic cur-  
634 rent and  $I_{ext}$  is the external current. This neuron model is capable of producing post inhibitory

635 bursts, regulated by the gating variable that decays with the time constant  $\tau_h^-$ , when the mem-  
 636 brane potential reaches a certain threshold  $V_h$  and rises with time constant  $\tau_h^+$ . However, when  
 637  $g_T$  is set to zero, the neuronal dynamics reduce to a leaky integrate and fire neuron. Currently,  
 638 we model GPe and STN neuronal populations with bursty neurons and the remaining neuronal  
 639 populations with leaky integrate-and-fire neurons, with conductance-based synapses.

640

641 The synaptic current  $I_{syn}(t)$  consists of three components, two excitatory currents corre-  
 642 sponding to AMPA and NMDA receptors and one inhibitory current corresponding to GABA  
 643 receptors, and is calculated as below:

$$I_{syn} = g_{\text{AMPA}} s_{\text{AMPA}}(t)(V(t) - V_E) + \frac{g_{\text{NMDA}} s_{\text{NMDA}}(t)(V(t) - V_E)}{1 + e^{-0.062V(t)/3.57}} + g_{\text{GABA}} s_{\text{GABA}}(t)(V(t) - V_I)$$

where  $g_i$  represents the maximum conductance corresponding to the receptor  $i \in (\text{AMPA},$   
 NMDA and GABA),  $V_I$  and  $V_E$  represent the excitatory and inhibitory reversal potentials, and  
 $s_i$  represents the gating variable for the channels, with dynamics given by:

$$\frac{ds_{\text{AMPA}}}{dt} = \sum_j \delta(t - t_j) - \frac{s_{\text{AMPA}}}{\tau_{\text{AMPA}}}$$

$$\frac{ds_{\text{NMDA}}}{dt} = \alpha(1 - s_{\text{NMDA}}) \sum_j \delta(t - t_j) - \frac{s_{\text{NMDA}}}{\tau_{\text{NMDA}}}$$

$$\frac{ds_{\text{GABA}}}{dt} = \sum_j \delta(t - t_j) - \frac{s_{\text{GABA}}}{\tau_{\text{GABA}}}$$

644 The gating variables for AMPA and GABA acts as leaky integrators that are increased by all  
 645 incoming spikes, with an additional constraint for NMDA that ensures that the maximum value  
 646 of  $s_{\text{NMDA}}$  remains below 1.

647 The values of neuronal parameters for all the nuclei are listed in Table S4, external inputs to  
 648 the CBGT nuclei are listed in Table S5, and the synaptic parameter values are listed in Table S6.

649

**Table S4. Neuronal parameters**

Parameter	unit	Cx	CxI	dSPN	iSPN	FSI	GPe	STN	Thalamus
$\tau_m$ (membrane time constant)	ms	20	10	20	20	10	20	20	27.78
$V_{rest}$ (resting membrane potential)	mV	-70	-70	-70	-70	-70	-70	-70	-70
$V_{threshold}$ (threshold potential)	mV	-50	-50	-50	-50	-50	-50	-50	-50
$V_L$ (leak reversal)	mV	-55	-55	-55	-55	-55	-55	-55	-55
$g_T$ (low threshold $\text{Ca}^{2+}$ maximal conductance)	$\text{mS/cm}^2$	0	0	0	0	0	0.06	0.06	0
$V_h$ (threshold potential for burst activation)	mV	-60	-60	-60	-60	-60	-60	-60	-60
$V_T$ ( $\text{Ca}^{2+}$ reversal potential)	mV	120	120	120	120	120	120	120	120
$\tau_h^+$ (burst duration in ms)	ms	20	20	20	20	20	20	20	20
$\tau_h^-$ (hyperpolarization duration)	ms	100	100	100	100	100	100	100	100

650

651 **Supp. Table 4. Neuronal parameters.** Neuronal parameters for each nucleus are listed in the  
 652 left column, with values shown on the right.

653

#### 654 **Spike timing dependent plasticity rule**

655 The plasticity rule we use is a dopamine modulated STDP rule also described in (38). All the  
 656 values of the relevant parameters are listed in Table S8. The weight update of a corticostriatal  
 657 synapse is controlled by three factors: 1) an eligibility trace, 2) the type of the striatal neuron  
 658 (iSPN/dSPN), and 3) the level of dopamine.

To compute these quantities for a given synapse, an activity trace of each neuron in the pre-synaptic and post-synaptic populations is tracked via the equations

$$\tau_{PRE} \frac{dA_{PRE}}{dt} = \Delta_{PRE} X_{PRE}(t) - A_{PRE}(t)$$

$$\tau_{POST} \frac{dA_{POST}}{dt} = \Delta_{POST} X_{POST}(t) - A_{POST}(t)$$

659 where  $X_{PRE}, X_{POST}$  are spike trains, such that  $A_{PRE}$  and  $A_{POST}$  maintain a filtered record of

660 synaptic spiking of the pre/post neuron, respectively, with spike impact parameters  $\Delta_{PRE}, \Delta_{POST}$   
661 and time constants  $\tau_{PRE}, \tau_{POST}$ .

662 If the post-synaptic spike follows the spiking activity of the pre-synaptic population closely  
663 enough in time, then eligibility trace ( $E$ ) increases and allows for plasticity to occur. On the  
664 other hand, if a pre-synaptic spike follows the spiking activity of the post-synaptic population,  
665 then  $E$  decreases. In absence of any activity and spikes, the eligibility trace decays to zero with  
666 a time constant  $\tau_E$ . Putting these effects together, we obtain the equation

$$\tau_E \frac{dE}{dt} = X_{POST}(t)A_{PRE}(t) - X_{PRE}(t)A_{POST}(t) - E.$$

667 The synaptic weight update depends on the dopamine receptor type of the striatal neuron;  
668 that is, if the neuron is a dSPN or iSPN. We assume that a phasic dopamine release promotes  
669 long term potentiation (LTP) in dSPNs and long term depression (LTD) in iSPNs. This factor  
670 is indicated by the learning rate parameter  $\alpha_w$ , which is set to a positive value for dSPNs and a  
671 negative value for iSPNs. The weight update dynamics is given by:

$$\frac{dw}{dt} = [\alpha_{w-X} E(t) f_X(K_{DA})(W_{max}^X - w)]^+ + [\alpha_{w-X} E(t) f_X(K_{DA})(w - W_{min})]^- \quad (6)$$

672 where  $X \in \{ \text{dSPN, iSPN} \}$  with  $\alpha_{w-dSPN} > 0$  and  $\alpha_{w-iSPN} < 0$ . Here, the weights of  
673 the corticostriatal synapses are bounded between the maximal value  $W_{max}^X$ , which depends on  
674 the SPN type, and a minimal value of  $W_{min} = 0.001$ . The precise values used for all relevant  
675 parameters are listed in Table S8.

676 In the weight update rule (6),  $K_{DA}$  represents the dopamine level present. This quantity  
677 changes as a result of phasic release of dopamine (increments of size  $DA_{inc}$ ), which is correlated  
678 to the reward prediction error encountered in the environment. The parameter  $C_{scale}$  defines the  
679 scaling between the reward prediction error and the amount of dopamine released, and  $K_{DA}$   
680 obeys the equation

$$\tau_{DOP} \frac{K_{DA}}{dt} = C_{scale} (DA_{inc}(t) - K_{DA}) \delta(t) - K_{DA},$$

681 where

$$DA_{inc}(t) = r(t) - Q_{chosen}(t)$$

682 for reward  $r(t)$  and expected value  $Q_{chosen}(t)$  of the chosen action. Trial-by-trial estimates of  
 683 the values of the actions (left/right) are maintained by a simple Q-update rule:

$$Q_a(t+1) = Q_a(t) + \alpha_q(r(t) - Q_a(t))$$

684 where  $a \in \{\text{left, right}\}$  and where  $\alpha_q$  represents the learning rate of the Q-values.

685 Finally, the function  $f_X(K_{DA})$  converts the level of dopamine into an impact on plasticity  
 686 in a way that depends on the identity  $X$  of the post-synaptic neuron, as follows:

$$f_X(K_{DA}) = \begin{cases} K_{DA}, & X = dSPN, \\ \frac{K_{DA}}{c + |K_{DA}|}, & X = iSPN, \end{cases}$$

687 where  $c$  sets the dopamine level where  $f_{iSPN}$  reaches half-maximum.

688

**Table S5. External input to the CBGT populations**

Population	Receptor	External input mean frequency	External input efficacy	Number of external connections
CxI	AMPA	3.7	1.2	640
Cx	AMPA	2.3	2.0	800
dSPN	AMPA	1.3	4.0	800
iSPN	AMPA	1.3	4.0	800
FSI	AMPA	3.6	1.55	800
GPi	AMPA	0.8	5.9	800
GPe	AMPA	4	2.0	800
GPe	GABA	2	2.0	2000
STN	AMPA	4.45	1.65	800
Thalamus	AMPA	2.2	2.5	800

689

690 **Supp. Table 5. External inputs to CBGT nuclei.** Each nucleus is listed on the left, with input

691 frequency, efficacy, and number of connections listed by receptor.

692

**Table S6. Synaptic parameters**

Parameter	unit	Value
$\tau_{\text{AMPA}}$	ms	2
$V_E$	mV	0
$\tau_{\text{NMDA}}$	ms	100
$\tau_{\text{GABA}}$	ms	5
$V_I$	mV	-70
$\alpha$	-	0.6332

693

694 **Supp. Table 6. Synaptic parameters.** Parameters for the simulated synapses.

695

**Table S7. CBGT connectivity**

Connection type	Connection probability	g (nS)	Receptor
Cx-dSPN	1.0	0.015	AMPA
Cx-dSPN	1.0	0.02	NMDA
Cx-iSPN	1.0	0.015	AMPA
Cx-iSPN	1.0	0.02	NMDA
Cx-FSI	1.0	0.43	AMPA
Cx-Th	1.0	0.025	AMPA
Cx-Th	1.0	0.035	NMDA
Cx-Cx	0.13	0.0127	AMPA
Cx-Cx	0.13	0.08	NMDA
Cx-CxI	0.0725	0.113	AMPA
Cx-CxI	0.0725	0.525	NMDA
CxI-Cx	0.5	1.05	GABA
CxI-CxI	1.0	1.075	GABA
dSPN-dSPN	0.45	0.28	GABA
dSPN-iSPN	0.45	0.28	GABA
dSPN-GPi	1.0	2.09	GABA
iSPN-iSPN	0.45	0.28	GABA
iSPN-dSPN	0.5	0.28	GABA
iSPN-GPe	1.0	4.07	GABA
FSI-FSI	1.0	3.2583	GABA
FSI-dSPN	1.0	1.77	GABA
FSI-iSPN	1.0	1.66	GABA
GPe-GPe	0.067	1.75	GABA
GPe-STN	0.067	0.35	GABA
GPe-GPi	1.0	0.058	GABA
STN-GPe	0.1617	0.07	AMPA
STN-GPe	0.1617	1.51	NMDA
STN-GPi	1.0	0.038	GABA
GPi-Th	1.0	0.033	GABA
Th-dSPN	1.0	0.38	AMPA
Th-iSPN	1.0	0.38	AMPA
Th-FSI	0.83	0.1	AMPA
Th-Cx	0.83	0.03	NMDA

696

697 **Supp. Table 7. CBGT connectivity.** Connection type and probability by nucleus and receptor.

698

**Table S8. Number of neurons in each CBGT population**

Population	Number of neurons
Cx	204
CxI	186
dSPN	75
iSPN	75
FSI	75
GPe	750
GPi	75
STN	750
Th	75

699

700 **Supp. Table 8. Number of neurons in each CBGT nucleus.**

701

**Table S9. STDP parameters**

Parameter	Value
$\Delta_{PRE}$	0.8
$\Delta_{POST}$	0.04
$\tau_{PRE}$	15
$\tau_{POST}$	6
$\tau_E$	100
$\alpha_{w-dSPN}$	39.5
$\alpha_{w-iSPN}$	-38.2
$W_{max}^{dSPN}$	0.055
$W_{max}^{iSPN}$	0.035
$W_{min}$	0.001
c	2.5
$\tau_{DOP}$	2.0
$\alpha_q$	0.6
$C_{scale}$	85

702

703 **Supp. Table 9. STDP parameters.**

704