

How the dynamic interplay of cortico-basal ganglia-thalamic pathways shapes the time course of deliberation and commitment

Zhuojun Yu¹, Timothy Verstynen^{1,2}, Jonathan E. Rubin^{2,3}

1 Department of Psychology & Neuroscience Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, United States of America

2 Center for the Neural Basis of Cognition, Pittsburgh, Pennsylvania, United States of America

3 Department of Mathematics, University of Pittsburgh, Pittsburgh, Pennsylvania, United States of America

 These authors contributed equally to this work.

* timothyv@andrew.cmu.edu (TV); jonrubin@pitt.edu (JR)

Abstract

Although the cortico-basal ganglia-thalamic (CBGT) network is identified as a central circuit for decision-making, the dynamic interplay of multiple control pathways within this network in shaping decision trajectories remains poorly understood. Here we develop and apply a novel computational framework — CLAW (Circuit Logic Assessed via Walks) — for tracing the instantaneous flow of neural activity as it progresses through CBGT networks engaged in a virtual decision-making task. Our CLAW analysis reveals that the complex dynamics of network activity is functionally dissectible into two critical phases: deliberation and commitment. These two phases are governed by distinct contributions of underlying CBGT pathways, with indirect and pallidostriatal pathways influencing deliberation, while the direct pathway drives action commitment. We translate CBGT dynamics into the evolution of decision-related policies, based on three previously identified control ensembles (responsiveness, pliancy, and choice) that encapsulate the relationship between CBGT activity and the evidence accumulation process. Our results demonstrate two contrasting strategies for decision-making. Fast decisions, with direct pathway dominance, feature an early response in both boundary height and drift rate, leading to a rapid collapse of decision boundaries and a clear directional bias. In contrast, slow decisions, driven by indirect and pallidostriatal pathway dominance, involve delayed changes in both decision policy parameters, allowing for an extended period of deliberation before commitment to an action. These analyses provide important insights into how the CBGT circuitry can be tuned to adopt various decision strategies and how the decision-making process unfolds within each regime.

Author summary

We investigate how the cortico-basal ganglia-thalamic (CBGT) network coordinates decision-making through its interconnected pathways. Using a novel Circuit Logic Assessed via Walks (CLAW) framework, we trace instantaneous neural activity through virtual CBGT networks as they engage in forced choice decisions. This analysis

uncovers two key phases of a decision: deliberation, shaped by the indirect and pallidostriatal pathways, and commitment, driven by the direct pathway. We also demonstrate that CBGT activity supports two distinct decision styles: fast decisions involve an early decision boundary collapse and strong directional preference, while slow decisions feature minimal changes during an extended deliberation phase. These findings reveal the dynamic mechanisms within the CBGT network that underlie the different decision processes and how these can be tuned to adapt decision-making across varying demands and contexts.

Introduction

Theoretical [1–5] and empirical [6–9] studies have established that the basal ganglia, working together with the thalamus and cortex, can play a critical role in the evidence accumulation process during decision-making. The distributed circuits of the cortico-basal ganglia-thalamic (CBGT) network simultaneously integrate external (e.g., sensory signals) and internal (e.g., learned contingencies) factors, until sufficient evidence is accumulated to allow one action to proceed [10, 11]. The topological structure of the interconnected CBGT nuclei [12–14], including their spatiotemporal organization into separable action representations [15], is ideally suited for managing this evidence accumulation process [3, 16, 17] and adapting its implementation based on environmental feedback [12, 18–21].

Prior computational work [5, 22, 23] has shown how the CBGT network can regulate decision policies (e.g., manage the speed-accuracy [24, 25] and exploration-exploitation [26, 27] tradeoffs). Specifically, interactions between distinct subnetworks within the CBGT pathways can be mapped to certain behavioral outcomes described by parameters in a drift-diffusion model (DDM; [10, 16, 28]), with the low-dimensional relationship between specific configurations of CBGT networks and decision policies characterized in terms of three *control ensembles* [22]. It has been further shown that dopaminergic plasticity at the corticostriatal synapses can alter the activity of CBGT control ensembles, so as to modulate the resulting decision policies in response to feedback following post-action outcomes [23]. These studies suggest that CBGT networks can establish the *static* state of decision policies. In typical evidence accumulation models this would be represented by fixed values of parameters such as drift rate (the rate of evidence accumulation toward a decision) and boundary height (the amount of evidence needed to make a decision), with gradual adaptation of these factors over the course of sequences of decisions if feedback is provided. However, it remains to be investigated (1) how information flows through different regions of the CBGT network during the course of an individual decision and (2) how the complex dynamics of CBGT networks give rise to variations in *dynamic* decision policies on this fast timescale.

To address the first of these open questions, we propose in this work a novel computational framework that we call CLAW, short for Circuit Logic Assessed via Walks, to depict the flow of noisy neural activity through a model CBGT network on a moment-by-moment basis (Fig 1A-D). Our CLAW analysis is related to, but extends, the classical framework for conceptually organizing CBGT activity, which assumes that the internal circuitry of the basal ganglia implements action selection through a division of labor between the structurally and functionally dissociable *direct* pathway (Fig 1A, green) and *indirect* pathway (blue) [13, 29, 30], with the former facilitating action selection and the latter suppressing action selection [29]. Discoveries over the past decade have revealed that, rather than functioning as independent and antagonistic mechanisms for facilitating or suppressing action selection, the direct and indirect pathways may engage in a dynamic competition for control over basal ganglia

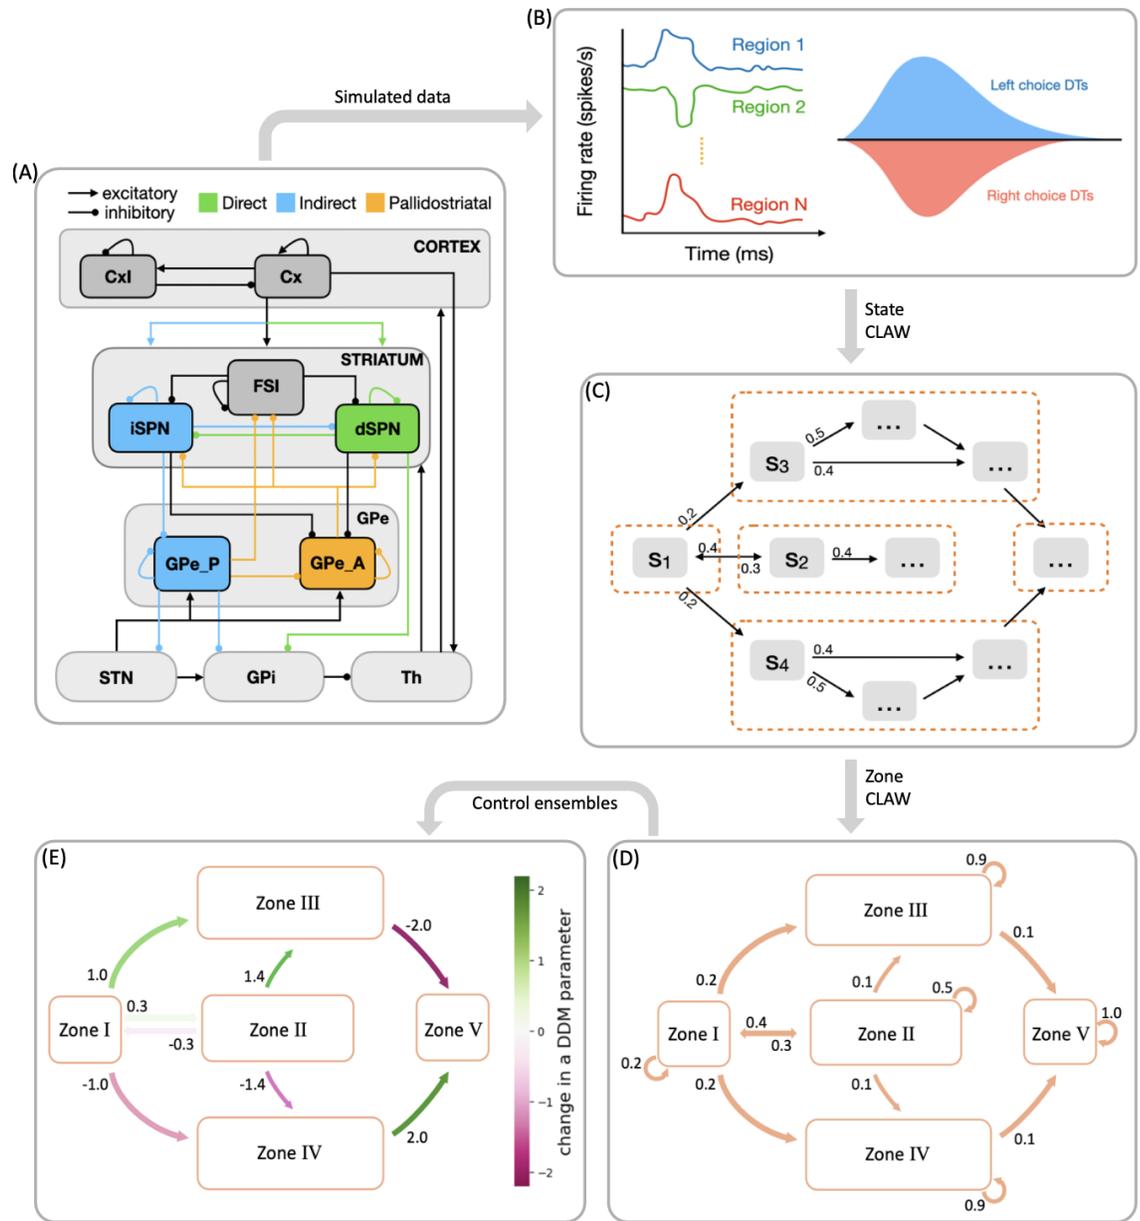


Fig 1. Schematic showing the steps in analyzing decision-making dynamics within the cortical-basal ganglia-thalamic (CBGT) network. (A) CBGT network, with connections color-coded as follows: classical direct pathway in green, classical indirect pathway in blue, and pallidothalamic pathway in gold. Arrows ending in dots indicate postsynaptic sites of inhibitory connections, while arrows ending in triangles indicate excitatory connections. Cx, cortical neurons; CxI, inhibitory cortical interneurons; FSI, fast spiking interneurons; dSPN, direct spiny projection neurons; iSPN, indirect spiny projection neurons; GPe, external globus pallidus; GPe_P, prototypical neurons; GPe_A, arkypallidal neurons; GPi, internal globus pallidus; STN, subthalamic nucleus; Th, thalamus. (B) Simulated behavioral data, including firing rates, choices, and decision times, generated by a spiking model for the CBGT network. (C) A chain of state transitions derived from the processed data. (D) States grouped into zones, with associated transition probabilities. (E) Levels of activity within zones mapped to DDM parameters through control ensemble analysis, showing how dynamic decision policies emerge from flow of network activity between zones.

output [3, 17, 31, 32]. Moreover, this model of CBGT pathways has been further complicated by a recent reappraisal for the cellular complexity of the external segment of the globus pallidus (GPe; see [33]), now seen as being composed of two general classes of neurons: prototypic (GPeP) and arky pallidal (GPeA) cells [34–37]. The arky pallidal neurons are currently thought to play a critical role in regulating striatal signaling, particularly during tasks requiring reactive inhibitory control [38–41]. In light of these findings, here we propose a computational model that incorporates an additional third pathway, the *pallido-striatal* pathway (Fig 1A, gold), that conveys ascending feedback inhibition to the striatum and is modulated by both striatal projection pathways. Then applying the CLAW framework to simulated CBGT dynamics in our model specifically allows us to generate novel insights and predictions about the complex interactions among all CBGT nuclei during decision processes, helping us to extract the details of how specific pathways interact to control decision-making dynamics. Here we use this approach to explore the way the direct, indirect, and pallido-striatal pathways regulate the bidirectional control of information through CBGT circuits and hence the dynamics of evidence accumulation as a decision is being made.

To address the second question, about the dynamic variations of decision policies within the fast timescale of individual decisions, we build on the prior observation of separable control ensembles within the CBGT pathways [22, 23]. These groupings of basal ganglia components, dubbed the *responsiveness*, *pliancy*, and *choice* ensembles, each represent a relation between CBGT firing patterns and DDM parameters that encode a decision policy. Responsiveness modulates how quickly evidence evaluation begins and the standard of evidence needed for decisions, in a positively correlated way, and is largely associated with the overall activity in the corticothalamic loop and direct pathway. Pliancy modulates evidence onset and the amount of evidence required to make a decision in opposite ways, and is strongly dependent on the overall activity of the indirect pathway. Choice modulates the direction and speed of evidence accumulation towards one action. With this ensemble the overall activity across channels has very little impact on the resulting decision policy, but instead the between-channel differences in firing activity show a robust association. This control ensemble framework provides mechanistic predictions about the roles of CBGT pathways in tuning decision-related policies [22], and how dopamine-mediated synaptic plasticity gradually shifts decision properties across multiple decisions [23]. Yet until now, the detailed dynamics of how these control ensembles shape information flow during the decision process has remained unknown.

In this work we set out to replicate the existence of the three independent control ensembles, using the more biologically realistic model of CBGT circuits that includes the pallido-striatal pathway for additional control. We then try to translate firing rate changes associated with transitions between different network states to changes in the DDM parameter space (as sketched in Fig 1D-E), allowing us to predict dynamic fluctuations and adjustments in decision policies over the fast timescale on which individual decisions are made. Our results reveal how instantaneous information flow through different CBGT pathways produces a dynamic decision policy and demonstrate two contrasting decision-making modes: fast decisions feature an early collapse in the evidence accumulation threshold and a clear directional bias, while slow decisions with more extended deliberation are driven by a delayed response in the decision boundary height and drift rate. By comparing the dynamics of control ensemble activation across fast and slow trials, we gain important insights about the contributions of specific aspects of CBGT activity to flexible decision-making, which may be harnessed to adapt decision policies during the evidence accumulation process.

Results

CBGT activity and CLAW

Our main goal here was to extract the temporal dynamics of information flow through CBGT circuits during the evidence accumulation process within individual decisions. To this end, we simulated a spiking model of the CBGT network [42] in the context of a simple two-choice task. Implementing a genetic search algorithm [23] we generated 300 distinct networks, each with a different configuration of synaptic weights, that produced firing rates of all CBGT populations within experimentally observed ranges (see also [22]). The networks were largely unbiased with respect to left and right choices, as the same connectivity parameters were shared by both channels. For each sampled network, we simulated 50 trials (i.e., choices) and gathered the time-dependent firing rates from a collection of $N = 10$ distinct cell populations within the CBGT network, across action channels representing left and right choices (Fig 1B). As a reference, Fig 2 shows an example of the time course of firing rates for all the nuclei in two consecutive trials from an example network. The firing rates were tracked up to the decision time (DT; Fig 2, pink region), defined as the time at which the instantaneous firing rate of the thalamic population for either of the channels first reaches a pre-specified decision threshold (set to 30 Hz). The network firing rates were computed in bins of width $\Delta t = 10$ ms and binarized based on whether activity in each bin was above or below a predefined threshold (Fig 2, green horizontal lines). We then defined a *state* $s_k \in \mathbb{R}^N$ ($k = 1, 2, 3, \dots, 2^N$) to be a unique pattern of activity across the N nuclei of interest in our network. Specifically, each $s_k = [\sigma_{k1}, \sigma_{k2}, \dots, \sigma_{kN}]$, where $\sigma_{kj} \in \{0, 1\}$. Thus, for each trial, we can convert the sequence of binarized firing rates into a sequence of states, each representing the pattern of CBGT activity in a time window. By aggregating all of the state sequences, we determine the transition probabilities between states and then construct a chain that describes the flow of neural states that occurs across the simulated decision processes (Fig 1C). See CBGT network in the Methods section for details of our modeling framework.

High-probability states within the CBGT network during the evidence accumulation process were analyzed using a novel procedure that we call *Circuit Logic Assessed via Walks* (CLAW). The CLAW diagram is depicted in Fig 3A. The name derives from the fact that this analysis can uncover how activity flows through the CBGT circuit by looking at how different cell populations are engaged together over the course of a decision. Transitions between states resemble walks, or chains that can revisit the same state, rather than paths, where each state is visited only once [45]. These transition probabilities from the current state to a subsequent state of the CLAW are shown on the directed edges in Fig 3A. We also show how each state relates to the speed characteristics of a decision by color-coding each state with the mean DT of the trials that visited that state. For purposes we elaborate on later, we further divided the decision time distribution of the full set of simulated networks into tertiles of equal mass, respectively defining fast (short DT), intermediate (medium DT), and slow (long DT) networks (Fig 3B). Finally, the table in Fig 3C shows the details of the activity associated with each CLAW state, including the binarized firing rates of the populations of interest, the activation probability of subthalamic nucleus (STN) and arkyvallid populations, the mean DT, the ratio of trials in which the left/right option was chosen, and the Kullback–Leibler (KL) divergence between the left and right trials’ DT distributions.

Note that for our CLAW analysis we used an arbitrary numbering system that indicates the unique state out of all 2^{10} possible states of the network, focusing only on those that occurred with relatively high frequencies. Starting from the decision onset state of the CLAW, i.e., the pre-stimulated state (PS) 60, either the left thalamus, the

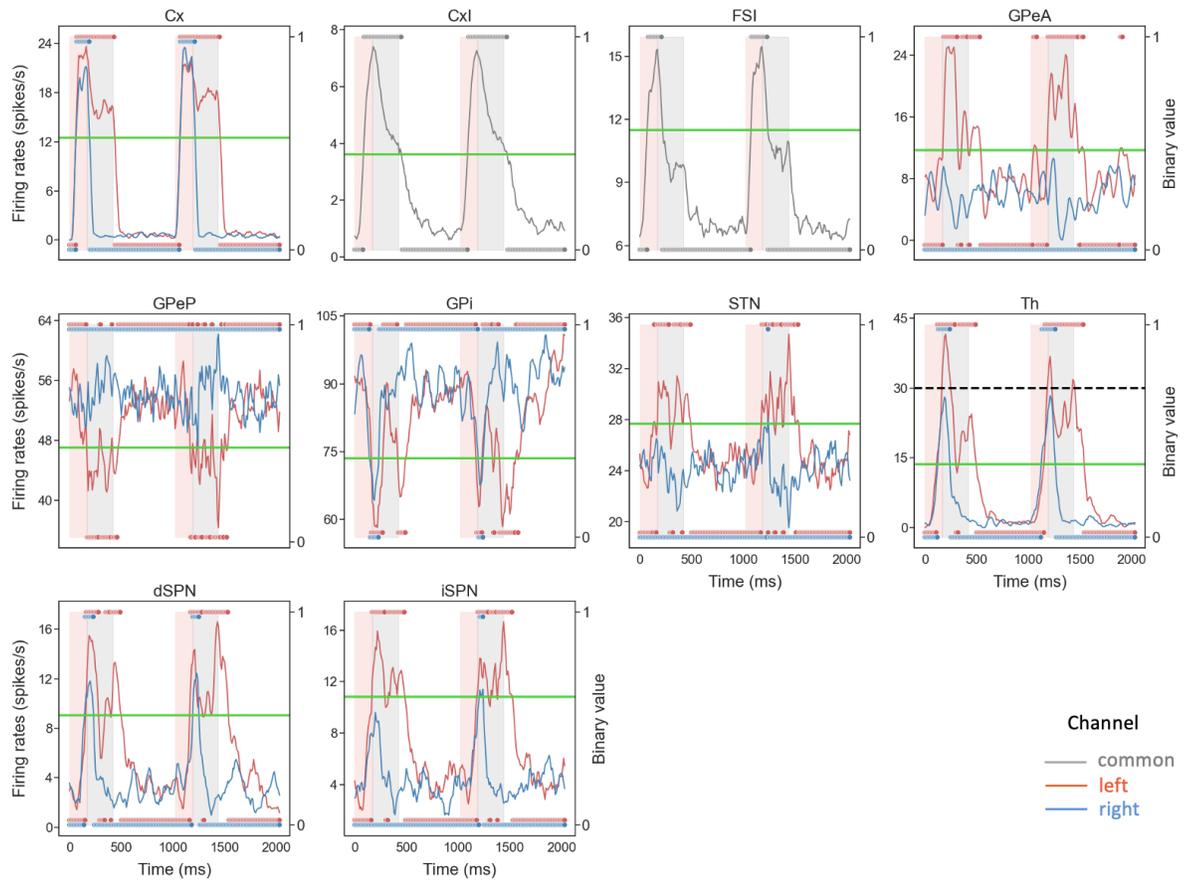


Fig 2. Example of the time course and binarization of firing rates for all nuclei in a two-choice task. Each red (blue) trace corresponds to activity in the left (right) action channel in a CBGT region, and the grey traces correspond to the populations (i.e., CxI and FSI) common to both action channels. Pink regions represent the decision-making phase, occurring before the thalamus (Th) of one of the action channels reaches the decision threshold of 30 Hz (dashed black line in Th panel). Grey regions represent the consolidation phase, where partial cortical input to the selected channel is sustained [43,44]. The unshaded regions represent the inter-trial interval. In each panel the right *y*-axis corresponds to binarized firing rates (dots at 0 or 1), where the horizontal green line indicates the binarization threshold (see Fig 10 for details on how thresholds were determined).

right thalamus, or both ramp up and cross the binarization threshold, progressing through states 61, 62, and 63, while the binarized activation states of the other nuclei remain unchanged. Note that at this point, thalamic firing typically remained below the decision threshold of 30Hz, as this threshold is higher than the thalamic binarization threshold (see the Th panel in Fig 2), which was based on the probability distribution of firing rates for this nucleus. Consequently, evidence accumulation continues through changes in firing of the other CBGT components until an action is selected. Following the flow of transition probabilities in the CLAW diagram revealed two distinct paths of decision-making: the inner CLAW (states colored blue in Fig 3C), with low probabilities of transitioning directly into a decision, and the outer CLAW (states colored green), with high probabilities of transitioning directly into a decision.

We found two distinct loops along the inner CLAW, associated with deliberation: an initial deliberation phase represented by the loop $63 \rightarrow 55 \rightarrow 183 \rightarrow 191 \rightarrow 63$ (and symmetrically $63 \rightarrow 59 \rightarrow 123 \rightarrow 127 \rightarrow 63$), and a second deliberation phase

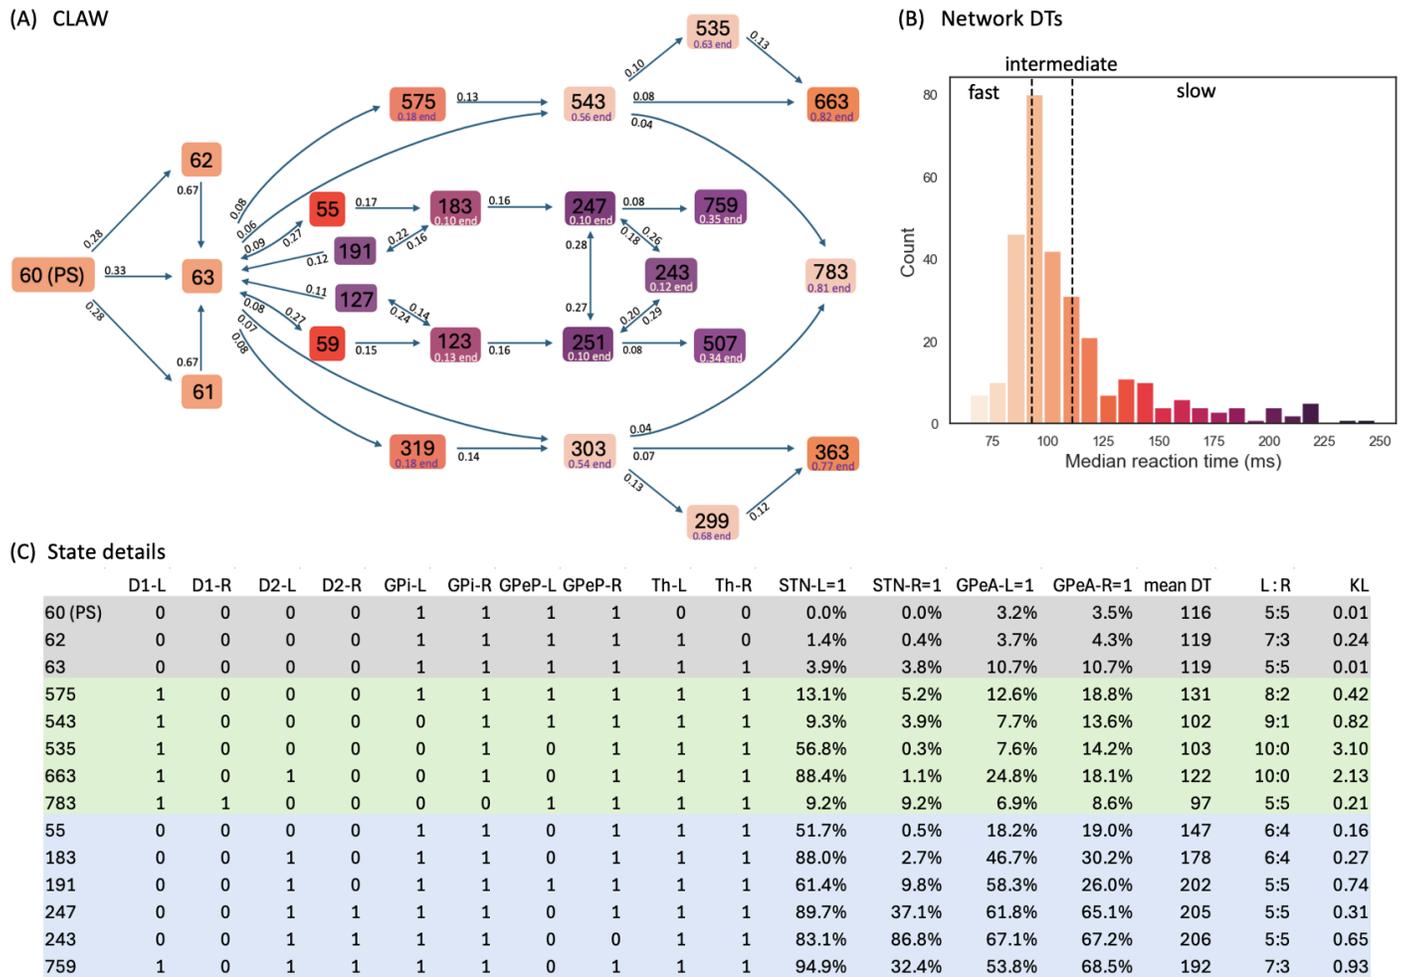


Fig 3. CLAW (Circuit Logic Assessed via Walks) diagram for CBGT network dynamics. (A) CLAW diagram. Black numbers in boxes indicate the network states, and the transition probability from a current state to a subsequent state is indicated by the number near the arrow pointing from the current state. The numbers below certain states (e.g., “0.82 end” below state 663) represent the probability that if these states are reached, then they are the final states prior to decisions. (B) Overall decision time distribution across 300 networks, categorized into three equal-count tertiles defining fast (left), intermediate (middle), and slow (right) networks, demarcated by vertical dashed black lines. The coloring of each CLAW state in panel A corresponds to the mean DT for all trials that visit this state, following the same color-coding scheme as in panel B. (C) Details of the states. A complete explanation of the full set of state properties, including those related to the right choice, is presented in the Supporting Information S1 Table. From left to right, after state labels: binarized firing rates of dSPN (D1), iSPN (D2), GPi, GPeP, and Th for left (-L) and right (-R) channels; probability of activation (binarized firing rate = 1) for STN and GPeA for left and right channels; mean DT over the trials that visit each state; the ratio of trials that chose left/right for those that visit each state; Kullback–Leibler (KL) divergence between left and right trials’ DT distributions. The grey rows correspond to the initial states that occur early in each trial and never lead directly to a decision, and the green and blue rows correspond to outer CLAW and inner CLAW states, respectively. The states in the lower half of the CLAW are not shown; these are symmetric – up to the swap of certain L and R channel binary values – with the states that lie in corresponding positions in the upper half.

represented by the loop 247 ↔ 243 ↔ 251 ↔ 247. The initial deliberation phase provides the network with the flexibility to reconsider its current trajectory of evidence, allowing a trial to return to state 63 for a reassessment of the evidence and then reset the direction of information flow. The inherent uncertainty in this phase, as reflected by the balanced ratio of choosing left versus right at these states, indicates that the options are still being weighed at this stage. During the second deliberation phase, activity had a high probability of switching between three different states via changes in the activation of key nodes like the prototypical pallidal neurons, thereby potentially influencing the subsequent decision outcome (discussed below). This switching reflects a second-tier deliberative process in which alternative actions were actively considered before committing to a final choice.

In contrast to the flexible, evaluative paths along the inner CLAW, the outer CLAW paths were characterized by predominantly one-way transitions, leading unambiguously toward the final choice. In further support of the idea of distinct deliberative and commitment phases, we noted that the states within the inner CLAW displayed no clear preference between left and right choices, which suggests that the network was still exploring different possibilities without any strong bias. Conversely, trials that progressed along the outer CLAW paths exhibited increasing certainty and stronger commitment to the choice. For instance, the left-to-right (L:R) ratio increased along the one-way arrows from state 63 to states 575, 543, 535, and 663, ultimately reaching an absolute 10:0 at states 535 and 663, which signifies a firm commitment to the chosen alternative once the trial reached these states. Therefore, based on comparisons of the transition routes and outcome ratios between the inner and outer paths, the CLAW demonstrates that the complex dynamics of decision-making is separable, reflecting distinct states associated with deliberation and commitment. Different CBGT subnetworks have distinct activity patterns tied to these state transitions and hence distinct circuit logic underlying the decision process.

The above analysis raises questions about how specific pathways in the CBGT network control the decision dynamics. Our model incorporated three distinct pathways (sub-networks): the direct pathway (colored green in Fig 1A), the indirect pathway (blue), and the pallidostriatal pathway (gold). We observed that when the direct pathway, via D1-expressing direct spiny projection neurons (dSPNs) in the striatum, became dominant in an action channel (state 575 for left decisions), the process bypassed the two deliberative phases and rapidly committed to the corresponding action, leading to a short decision time characteristic of fast or intermediate networks. In the case of left decisions, the activation of the left dSPN inhibited the firing of the ipsilateral internal globus pallidus (GPI; 575 to 543), thus increasing left thalamic firing rate to the decision threshold (at states 543, 535, or 663). Notably, this route features either no, or relatively late, engagement of the indirect (including suppression of the GPeP) and pallidostriatal (including activation of the GPeA) pathways, indicating that the direct pathway was the primary driver of action commitment. However, if the indirect pathway, via D2-expressing indirect spiny projection neurons (iSPNs), turned on more quickly than the direct pathway (state 183), the trial spent more time in deliberation and had a longer decision time, as indicated by the darker coloring of the inner CLAW, corresponding to slow decision trials (compare colors in Fig 3A-B). Here the STN also became active (see “STN-L=1” at state 183 in Fig 3C) due to the inhibitory influence of D2 striatal inputs onto the GPeP, which effectively released STN from inhibition. At the same time, enhanced excitation from the STN and reduced inhibition from the GPeP gradually enhanced the firing of arky pallidal neurons, especially when the trial was involved in the second deliberation loop (see “GPeA-L=1” and “GPeA-R=1” at states 243, 247, and 251). Thus, in these cases, the dominant indirect pathway actively induced pallidostriatal pathway activation, further contributing to a slowdown in

decision-making. Note that in this scenario, the direct pathway was always suppressed by the dominant activity of the indirect SPNs and arkyallidal neurons, so that the trial was prevented from quickly executing a decision. Comparing the activation states of the CBGT nuclei and the DTs between the inner CLAW and the outer CLAW shows that the relative dominance of the CBGT pathways at any given time determines how quickly the system converges to a decision: the direct pathway accelerates decision-making while the indirect and pallidostriatal pathways decelerate it. Overall, although the CLAW approach extends beyond the classical CBGT pathways, we see that their distinct roles became evident within the framework of CLAW analysis.

We next explored the mechanistic details underlying the two deliberation phases. Early in the deliberation process, the decision trajectory could turn around from state 183 to 191, in which left prototypical GPe neurons returned to supra-threshold activity. This possibility resulted from the feedback loop between prototypical neurons and the STN: the suppression of activity in GPeP via its indirect striatal input disinhibited the STN, and then via the excitatory connections from STN to GPeP, the prototypical neurons in turn received excitatory input, which allowed them to overcome the inhibition from iSPNs and resume their supra-threshold firing. This state transition highlights the possible role of the reciprocal STN-GPe prototypical neuron loop in promoting exploration by lengthening the deliberation time for evidence accumulation. From state 191, the trajectory could bounce back to state 183 or evolve to state 63. The latter occurred when the enhanced activity of arkyallidal neurons, which arose in state 191 due to the drive signals from STN, reduced the activity of indirect SPNs via the ascending pallidostriatal connections. This input halted the surge of decision-related striatal activity, including the potential activation of dSPNs, suppressed the planned action, and effectively reset the evidence evaluation process. From this process, we see the central role of the GPe, involving both prototypic and arkyallidal cell types, in regulating bidirectional information flow in order to implement the blocking of an incipient response.

On the other hand, if the trajectory did not turn around at state 183 but instead proceeded to state 247, a second deliberation phase could occur. Here, the battle between cortical drive and arkyallidal neurons' inhibition of direct striatal activity determined whether the trajectory was stuck in deliberation or committed to a decision. Importantly, no trial in our simulations failed to select an action, meaning that while the trial could spend time within the second deliberation phase, it did not remain perpetually in a state of deliberation. If the cortical drive became more intense, the ramping activity of the direct pathway (state 759 for the left channel and 507 for the right channel) led to greater commitment to the corresponding decision. In contrast, if dSPNs were inhibited by the stronger signals from arkyallidal neurons, then the indirect pathway could become dominant (states 247, 243, 251) and temporarily suppress the selection of an action, with a descending influence on information flow to the GPeP, STN, GPi, and finally to the thalamus. Specifically, GPi output increased, providing an enhanced suppression of the thalamic responses and leading to a temporarily indecisive state. Note that the STN-GPeP subcircuit was also engaged during this deliberation phase, as demonstrated by the three-state reciprocal transitions between states 247, 243, and 251, where the activity of prototypical neurons in the two channels switched on and off. The competition between the left and right prototypical neurons determined the direction of the subsequent possible commitment, and if they were both sufficiently suppressed at state 243, then the deliberation time would be lengthened. From this phase, it is clear that the interactive effect of the indirect and pallidostriatal pathways played a critical role in prolonging the evidence accumulation process through the combination of a pause process at the striatal level along with downstream effects that allowed for switching between the two possible options.

To reduce the impact of firing rate variability on individual states and to categorize key classes of decision trajectories, we finally partitioned all CLAW states into six zones, as shown in Fig 4. The partition was based on the probabilities and activity changes associated with transitions across different states. Specifically, zone I contained the pre-stimulated state 60 as well as states 61, 62, and 63, where only thalamic populations showed a response in their binarized firing rates to the stimulus at cortex. This zone serves as a pre-decision or launching zone, from which all trials originate, marking the initiation of the decision-making process. Zone II contained the group of states that could have a chance of returning to the pre-decision zone through the initial deliberation process. This inner CLAW zone played a critical role in early evaluation, where the system could explore alternatives rather than immediately converging on a specific action. In contrast, zones III and IV respectively corresponded to the left and right arms of the outer CLAW, each representing a strong commitment to a given action. Once the trial progressed into these zones, there was no return to deliberation and the system was effectively locked into a specific decision. The only exception was that both outer CLAW zones could, with a low probability, transition to the more neutral state 783, which we defined as zone VI. In this state there was an equal likelihood of choosing left or right actions (L:R = 5:5). Such transitions occurred when enough competing evidence arrived to engage the contralateral direct SPNs before a decision was made. Lastly, zone V included the states that comprised the second deliberation process. We identified zones III, IV, V, and VI as *absorption zones*, in which the probability of staying until the decision threshold was reached was approximately one (Fig 4A, loop arrows). Based on the pathway dominance in each zone (Fig 4B), we confirmed our previous observations that commitment to a choice occurs as soon as the direct pathway becomes sufficiently dominant in an action channel (zones III and IV), while increased activation in the indirect and pallidostriatal pathways contributes to ongoing deliberation (zones II and V). This zone partition served as a succinct representation of the state CLAW by capturing the activity of the main pathway components at different phases of decision-making. As a summary of the above analysis, Fig 5 shows the temporal sequence of CBGT pathway activation patterns associated with different classes of decision trajectories along the zone CLAW.

This zone version of the CLAW provides some additional insights into the dynamics of decisions through CBGT networks. First, the dSPNs in both channels were fully activated in zone VI, suggesting a strong competition between the two options. The existence of this neutral zone indicates that the two outer zones were actually not absolute in their commitment: while most trials that reached these zones would remain committed to their initial choice, a few “indecisive trials” ultimately switched to the opponent action. This switch occurred due to the late ramping of the competing action channel, where delayed activation in its direct pathway became sufficiently influential to rapidly override the initial commitment and could lead to a decision reversal. Second, there were direct paths from zone II to the outer zones, representing occasional transitions from the inner CLAW to the outer CLAW that were only evident with grouping of states into zones (hence they did not appear in Fig 3A). These connections highlight the possibility of direct shifts from the early deliberation phase to action commitment, without a need to return to zone I for a fresh re-evaluation or to continue into a prolonged second deliberation phase. Along this shortcut, the dSPN firing rate rapidly increased above the threshold, suppressing the original ramping activity of the iSPN and GPeA, thereby preventing the trial from fully completing the deliberative process and correspondingly accelerating the decision-making. Third, the irreversible transition between the two deliberation zones (II and V) suggests that once the trial moved into a deeper, more fully activated phase of deliberation, the likelihood of resetting deliberation reduced, moving the process toward a concluding outcome and

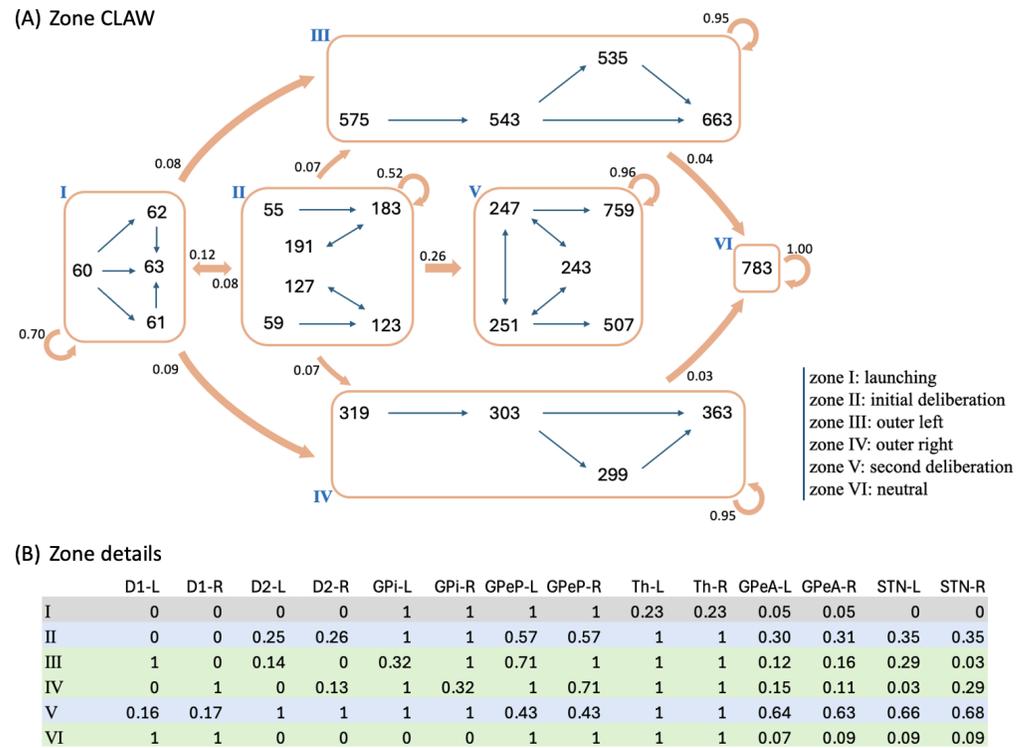


Fig 4. CLAW partitioned into zones. (A) Zone CLAW diagram. Zone I contains the pre-stimulated state 60 as well as states 61, 62, and 63, where only the binarized firing rates of thalamic populations may cross threshold. Zones II and V correspond to the initial and second deliberation phases within the inner CLAW, respectively. The left and right arms of the outer CLAW are represented by zones III and IV, respectively, each of which has a low probability of transitioning into zone VI, which consists of a single neutral state 783. The transition probability from one zone to another is indicated by the number near the arrow pointing from the source zone. The loop arrows represent the probability of staying in a zone or reaching the decision threshold from that zone. (B) Details of the zones. From left to right, after zone labels: probability of activation for dSPN (D1), iSPN (D2), GPI, GPeP, Th, GPeA, and STN for left (-L) and right (-R) channels, when networks are in each zone. The grey rows correspond to the launching zone, and the green and blue rows correspond to the zones containing the outer CLAW and inner CLAW states, respectively.

preventing indefinite uncertainty, albeit with a long decision time. Comparing zone V with zone VI, we see that both dSPN populations in zone VI were above threshold, while the activation pattern in zone V features an analogous balance across channels but for the indirect and pallidostriatal pathways (both iSPN, STN and GPeA populations were above threshold). Hence, the trials reaching in zone V could not turn back from this all-out battle, but rather waited to see which GPI firing rate would fall first. In this sense, although the dominant activity of the indirect and pallidostriatal pathways may temporarily delay the decision, the growing activation of the direct pathway in zone V ultimately ensured convergence to a choice. This interplay among pathways points out the network's ability to balance deliberation and commitment, maintaining functionality even in the face of competing influences. Overall, by categorizing CLAW states into key zones, we gained a clearer understanding of the critical transitions between deliberation and commitment, including certain shifts that were infrequent, but nonetheless enhanced flexibility in complex decision-making dynamics. The zone CLAW also helps us to appreciate the overall stability of decision-making processes despite the inherent

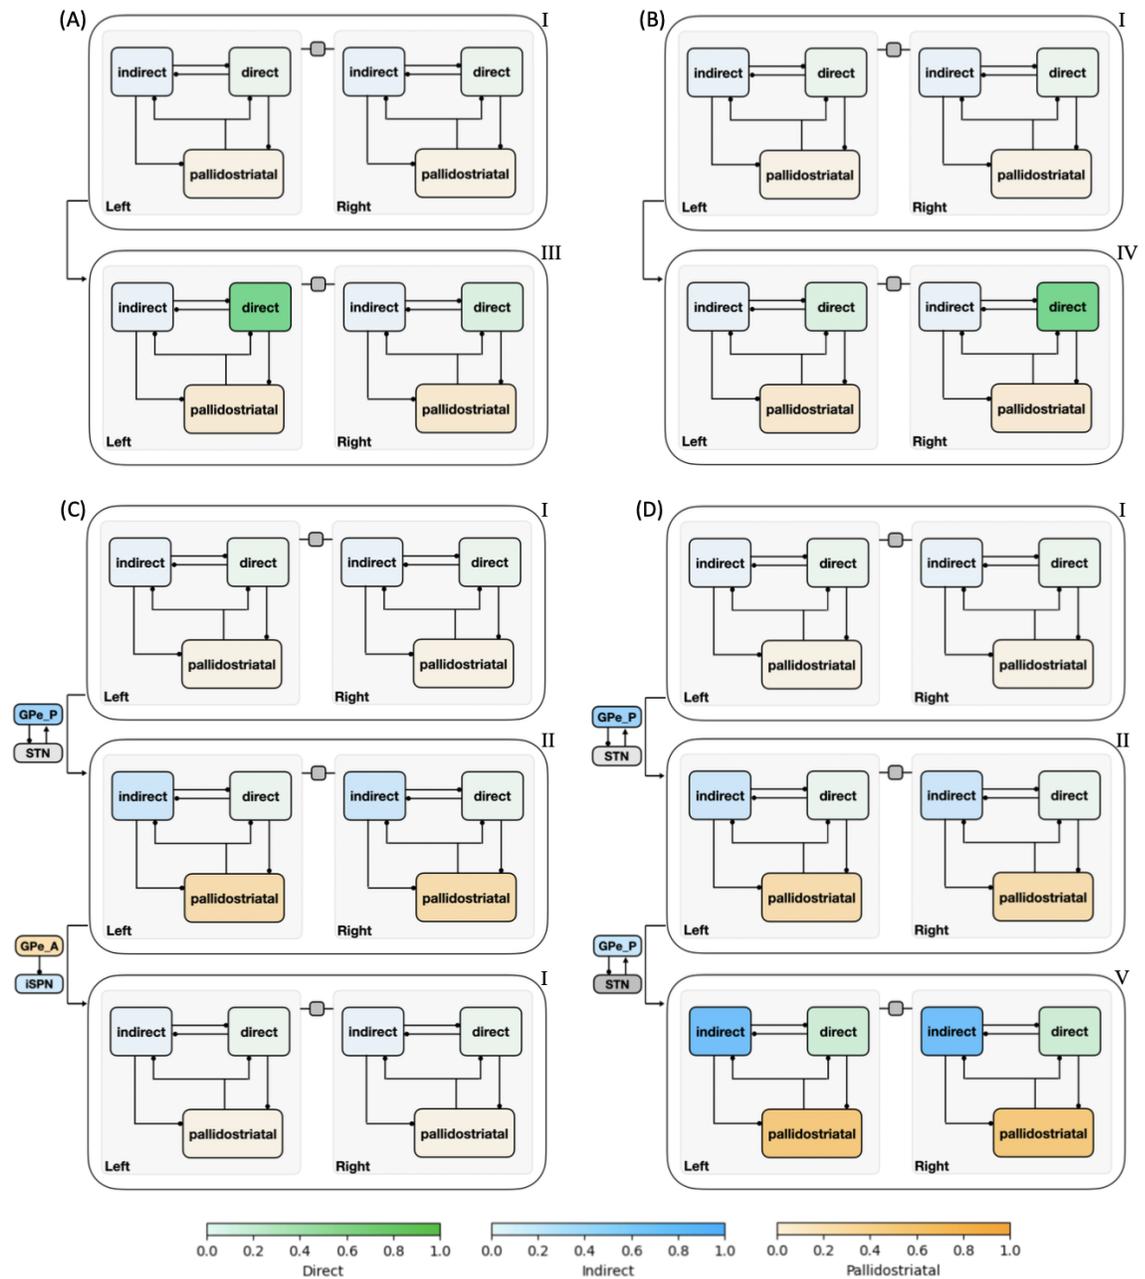


Fig 5. Temporal activation of CBGT pathways for different decision trajectories along the CLAW. The direct, indirect, and pallidostratial pathways across both left and right channels are enclosed in boxes, with their activation strength colored in blue, green, and gold, respectively (color bars at bottom). Darker shades indicate stronger activation within each pathway. Nuclei that play a key role for some phase transitions are indicated near the corresponding arrows. (A) Launching region → Left commitment. (B) Launching region → Right commitment. (C) Launching region → Initial deliberation → Launching region. (D) Launching region → Initial deliberation → Second deliberation. The Roman numeral on the upper right corner of each pattern corresponds to the zone labels in Fig 4.

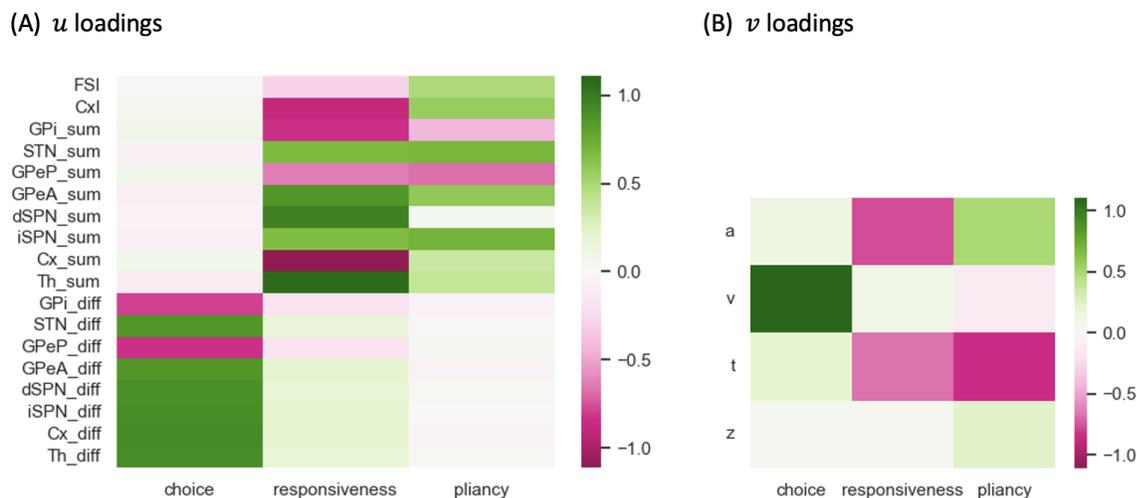


Fig 6. Loadings of the canonical vectors obtained from canonical correlation analysis (CCA) between CBGT activity and DDM parameters. (A) Loadings of u , corresponding to CBGT firing rates. The subscript ‘sum’ refers to the sum of rates in a region across both action channels, while the subscript ‘diff’ refers to the rate in the left channel minus that in the right channel. (B) Loading of v , corresponding to DDM parameters. These three components are referred to as choice, responsiveness, and pliancy control ensembles, respectively.

variability present in neural firing activity, highlighting the robustness of the competition between CBGT subnetworks that effectively controls how incoming information leads to deliberation and decisions.

CLAW and decision policies

Previous modeling studies, using two-choice decision tasks, have identified three low-dimensional control ensembles within the CBGT network. Each is associated with specific changes to the decision policy, defined in terms of parameter configurations of evidence accumulation models [22, 23]. We replicated this analysis using our CBGT network architecture, with the aim of examining the relationship between different CLAW zones and decision policies. First, we gathered the behavioral features (decision times and choices) produced by the simulated CBGT networks and fit their distributions to the DDM using the Hierarchical Sequential Sampling Modeling (HSSM) toolbox.¹ This step allowed us to derive sets of four key DDM parameters associated with the decision process for each network: boundary height a , drift rate v , onset time t , and starting bias z . We then applied canonical correlation analysis (CCA) to compute a low-dimensional mapping between correlated patterns in the space of CBGT network activity and the space of DDM decision parameters. For CBGT network activity, we considered two aspects of activity within each CBGT population: (1) global firing rates across the left and right action channels (i.e., sum of the firing rates in each region across both channels), and (2) bias in activity towards one action (i.e., difference in the firing rates of each cell population, between the left and right channels). From this analysis, we captured the first three pairs of component vectors that maximized the correlation between CBGT activity and DDM parameters. Figure 6 shows color-coded representations of the loadings from the vector u consisting of the firing rate components and from the vector v comprising the DDM parameter components from the CCA. See Control ensembles in the Methods section for more details.

¹For more information on the HSSM toolbox, see <https://lncbrown.github.io/HSSM/>.

Our analysis recovered three components that are nearly identical to the components found in prior work [22, 23]. Comparing across columns, the first component is strongly associated with the drift rate v and the between-channel differences in firing rates, consistent with the previously identified *choice* ensemble. Greater activity of cortical, striatal, STN, GPeA, and thalamic neurons in the left channel relative to the right channel, coupled with weaker activity in GPi and GPeP, corresponds to more positive drift rates toward the left decision. The second component, in contrast, loads heavily on the overall activity of corticothalamic systems as well as the direct pathway, with negative loading on both boundary height a and onset time t . This pattern agrees with what had previously been called the *responsiveness* ensemble, because it modulates, in the same direction, both how quickly evidence evaluation begins and the level of evidence needed for a commitment to a decision and hence the overall response speed. The final component loads heavily on both the indirect and pallidothalamic pathways and shows a positive loading on a but a negative loading on t . This component matches well with the *pliancy* ensemble, reflecting the ease with which the evidence collected can be translated into a commitment to a decision. In our case, the choice ensemble corresponded to the strongest component, describing the most covariance between CBGT activity and DDM components, whereas in prior work it was the weakest of the three components [22, 23]. Other than this ordering of components, our CCA results are highly consistent with prior work using different variants of the CBGT spiking network model.

With these control ensembles in hand, we next investigated how the dynamics of CBGT network activity during decisions tunes the resulting decision policy parameters. To this end, we translated the time series of CBGT firing rates to the time course of individual control ensemble engagement. Specifically, for each simulated trial we computed the change in the 18 CBGT firing rate measurements (see column labels in Fig 6A) between consecutive time bins, denoted by ΔF_k from the $(k - 1)$ -th to the k -th bin. Next, we projected $\Delta F_k \in \mathbb{R}^{18}$ onto the control ensemble space by mapping $W_k = \Delta F_k^T U$, where the components of the u loadings comprise the columns of U , such that $W_k \in \mathbb{R}^3$. Each component of W_k represents how the instantaneous firing rate change from the $(k - 1)$ -th to the k -th time bin corresponds to a change in the activation, or *drive*, of one of the three control ensembles. Finally, we grouped all trials into four types, categorized by their decision times and selected choices: fast left, fast right, slow left, and slow right. Given that decision times vary across trials, we aligned the trials in each group to the response time, as shown in Fig 7. For a full description of this approach, see Control ensembles in the Methods section.

We observed that the main difference between fast choices (Fig 7AC) and slow choices (Fig 7BD) is the timing at which the drives began to change. For the fast trials, the choice drive showed a modest deviation from baseline, becoming positive for the left action and negative for the right action. This happened near the end of the decision time, marking the point at which the system began to show a directional commitment. However, the drive of the responsiveness and pliancy components rose earlier and more dramatically, around the midpoint of the decision process, indicating that the system acted early to modulate the response speed and the degree of evidence necessary for selection. For slow trials, the choice ensemble still changed a bit later than the other two ensembles. Yet, considering the much longer decision times on slow trials relative to fast trials (note the differences in time axis scales in Fig 7AC vs. Fig 7BD), the slow cases exhibited more similar timing across the three ensembles, with all exhibiting a gradual acceleration as the long decision process neared completion. In addition, the increase in the magnitude of the responsiveness and pliancy ensembles was notably smaller in the slow cases than that in the fast cases, while the changes in the choice ensemble closely overlapped (see the right insets in Fig 7). Thus, the comparison across

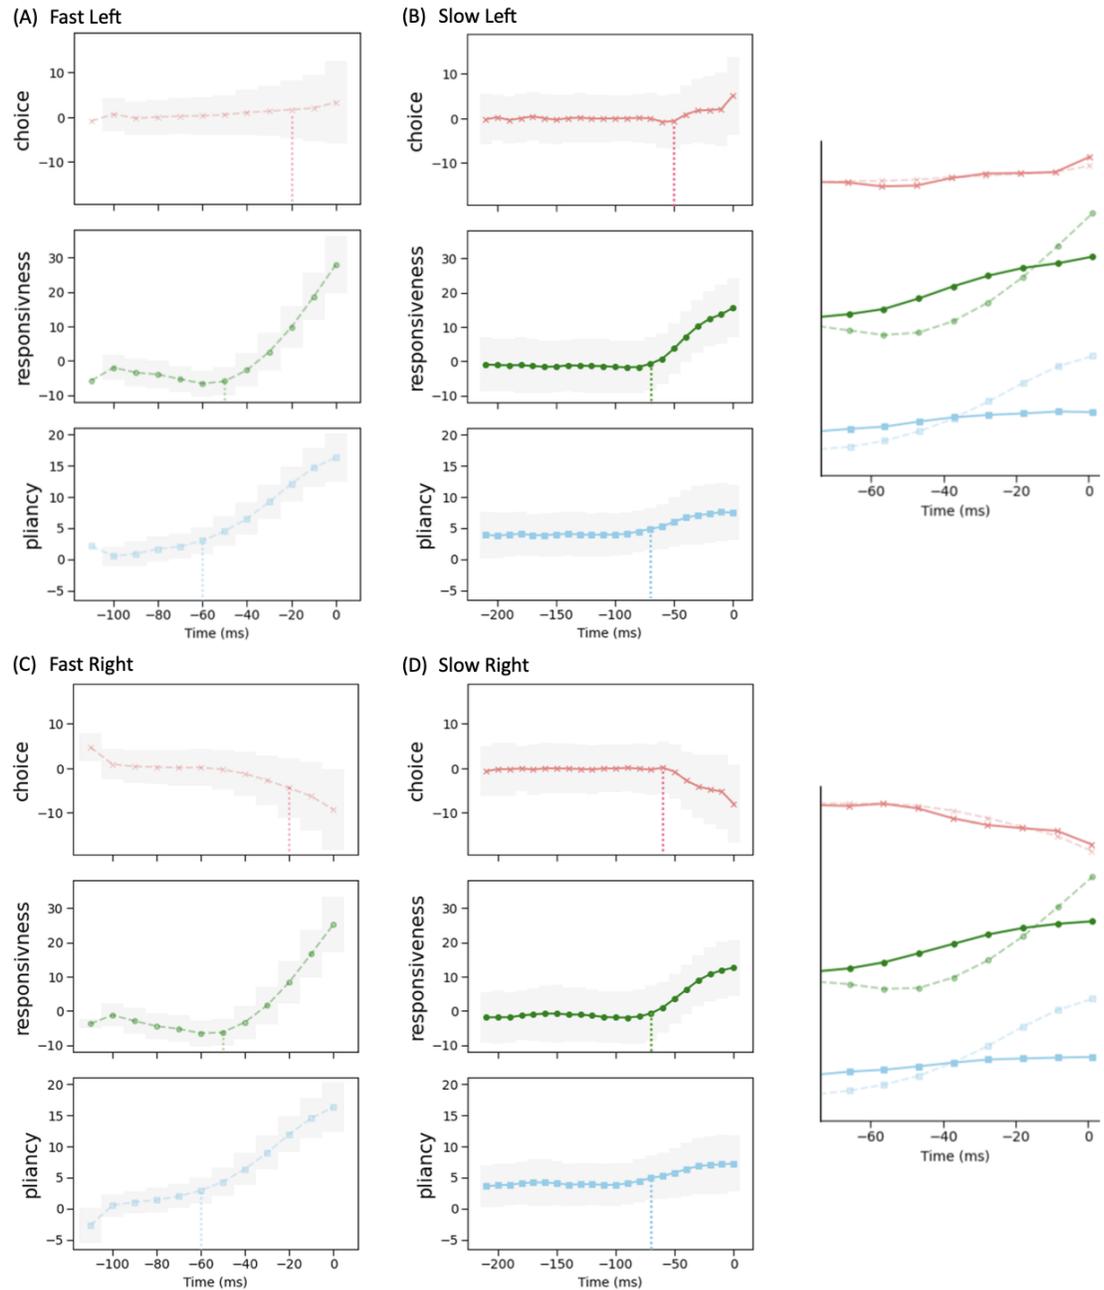


Fig 7. Time evolution of control ensemble drives. The percentage changes in each drive are averaged across all trials within each of four decision groups — (A) fast left, (B) slow left, (C) fast right, and (D) slow right, from the stimulus onset time to the decision time (set to be 0). Since each trial had a different decision time, we performed the averaging by aligning trials on their decision times, with the averages represented by dots connected by lines (light dashed for fast and dark solid for slow). Shaded areas represent the standard deviation for each group. The vertical dotted lines with colors matching the plotted curves mark the timing at which each control ensemble's drive shows a notable response. For fast and slow trials to the same choice, each inset at the right aligns the drives, zoomed in on the time range with pronounced changes (approximately 75 ms), on the same scale for comparison.

decision types shows that the critical components regulating overall decision speed are responsiveness and pliancy. 408

We next considered how these control ensemble changes map to specific activity in the CBGT cellular populations. Consider the distinctions in the temporal patterns of control ensemble activation between fast and slow decisions in terms of the activity of specific CBGT pathway components. First, as shown in Fig 6A, the increase in responsiveness was associated with the overall activity in corticothalamic and direct pathways, and specifically, greater activity in thalamic and dSPN neurons along with weaker activity in cortical and GPi neurons. Second, the increase in pliancy corresponded to the overall activity in components of both the indirect and pallidostriatal pathways, reflected by the amplified activity in iSPN, STN, and GPeA neurons and attenuated activity in GPeP neurons. Third, the change in the choice ensemble was driven by differential activity between the two channels in all pathways. From the timing difference observed from Fig 7, we inferred that a fast decision featured an early surge of overall CBGT activity before the emergence of between-channel differences. This seems counterintuitive, as one might expect that faster decisions would involve earlier differentiation between action channels, reflecting a more immediate commitment to one choice over the other. However, the delay in the emergence of between-channel differences aligns with the notion that fast decisions involve a stereotyped process, where the overall network activity assumes a state that is conducive to rapid implementation of whichever option subsequently becomes favored. As a result, only after the general decision approach had been established by the necessary components of the CBGT network, did the divergence of activity between channels become pronounced. This allowed for a smoother transition and potentially minimized unnecessary conflict or indecision between competing action channels. In contrast, the slow decision process was more deliberative, with prolonged competition and evaluation between options, including additional processing associated with transient pause, action switches, or suppression. These processes required the coordinated efforts of many CBGT nuclei, as we discussed in CBGT activity and CLAW in the context of the initial and second deliberation phases. Consequently, the turning point of the overall network activity for slow decisions occurred much later than in the fast cases and was associated with the emergence of a bias or preference towards one decision option. 409

In big picture terms, the temporal patterns that we observed in the control ensembles are consistent with our previous CLAW analysis showing that, at least in CBGT-driven decisions, there exist two stages — deliberation and commitment — to every choice. Note that, in contrast to the two extended deliberation phases that most fast trials bypassed (cf. Fig 4A, zones II and V), the deliberative stages of the fast trials occurred during the relatively brief time that the fast trials spent in the launching region (Fig 4A, zone I). During this time the rapid responsiveness and pliancy ensemble responses occurred (Fig 7AC) as the CBGT network quickly organized itself for action. This analysis highlights a relationship between the contribution of the three control ensembles and the distinct phases of a decision, achieved through the corresponding patterns of dynamic interactions among CBGT pathways. 410

By exploiting the relation between control ensembles and DDM parameters, we next investigated how the dynamics of CBGT network activity during decisions can be formulated in terms of serially tuning the evidence accumulation process itself during the decision. For this step, and for each pair of zones (i, j) we considered the change in the CBGT firing rate measurements associated with the direct transition from zone i to zone j across the subset of $p \leq 300$ networks that undergo this transition. Denote the firing rate change from zone i to zone j by $\Delta F_{ij} \in \mathbb{R}^{p \times 18}$. We projected ΔF_{ij} onto the control ensemble space using the u loadings derived from the CCA and obtained $W_{ij} \in \mathbb{R}^{p \times 3}$. For each column of W_{ij} , a positive (or negative) element indicates that the 411

sample aligns with the corresponding control ensemble in a positive (or negative) direction, and thus the parameters of DDM that have positive loadings in that control ensemble will increase (or decrease) as the CBGT activity progresses from zone i to zone j . As both responsiveness and pliancy ensembles influence the boundary height and onset time, we finally projected W_{ij} to the DDM space. Using the components of the v loadings to form the columns of V , we computed $P_{ij} = W_{ij}V^T \in \mathbb{R}^{p \times 4}$ which relates the firing rate changes to changes in the DDM parameters. Hence, the median of all p projections in each column of P_{ij} is proportional to the overall change in the direction and magnitude of the corresponding decision policy parameter in the transition from zone i to zone j .

Figure 8 shows the influence of the CBGT network activity flow on each parameter of the DDM. The coloring and number of each arrow indicates the percentage change in the parameter values as the decision trajectory traveled between CLAW zones. Note that since all control ensembles had weak loadings on the bias factor z , we did not consider its associated changes. First, we observed that the modulation of the drift rate v along the outer CLAW was strongly positive when committing to the left action (zones I/II to III) and strongly negative when committing to the right action (zones I/II to IV). These shifts in the drift rate resulted in a clear directional bias and a rapid approach toward the selection of the appropriate action. In contrast, v was less affected when information flowed along the inner CLAW, leading to a much slower accumulation process regardless of the action choice. Reversion from the outer CLAW to zone VI, with equal likelihoods of the two decisions, was associated with changes in v that countered the initial, commitment-related changes. Here we note that the pre-decision states appear to have been somewhat biased towards the left choice, as the magnitude of the modulation of v for I \rightarrow IV was twice as large as that for I \rightarrow III. This indicates that the trajectories may not have needed as much of an increase in drift rate before reaching the left decision threshold. Correspondingly, the drift rate modulation for the initial deliberation transition I \rightarrow II was small and negative, such that this progression counteracted the leftward bias and led to a more neutral exploratory state. Our results on the v dynamics align well with the change in the activation pattern of CBGT regions across zones (Fig 4B), as the drift rate parameter was strongly associated with differential activity between left and right channels (see choice ensemble loadings in Fig 6) — these between-channel differences became significant as the network activity traveled along the outer CLAW, leading to dramatic changes in the drift rate, while they vanished almost completely along the inner CLAW and resulted in minimal drift rate changes.

Second, we observed a “boundary collapse” phenomenon, as each trial proceeded towards a commitment to a choice, as indicated by the pink arrows throughout Fig 8B. This collapse reflected a reduction in the amount of evidence needed for a choice to be made, which facilitated action commitment. An exception occurred at the transitions between the launching and initial deliberation phases (zones I and II), during which the boundary height remained almost unchanged, staying as high as when the evidence accumulation initially started. This invariance suggests a more cautious approach at the onset of decision-making, with the engagement of indirect and pallidostriatal pathway neurons blocking any imminent action. In contrast, progression to zone V lowered the decision boundary substantially, consistent with our previous analysis that the second deliberation ultimately ensured convergence to a choice. A closer inspection of the binarized firing rates in the two deliberation zones (Fig 4B) reveals that the neural populations in both channels were much more activated in zone V compared to zone II. This implies a dramatic alteration in both the responsiveness and pliancy ensembles and therefore the boundary height when transitioning to zone V.

Lastly, the onset time t consistently shortened no matter which decision path was

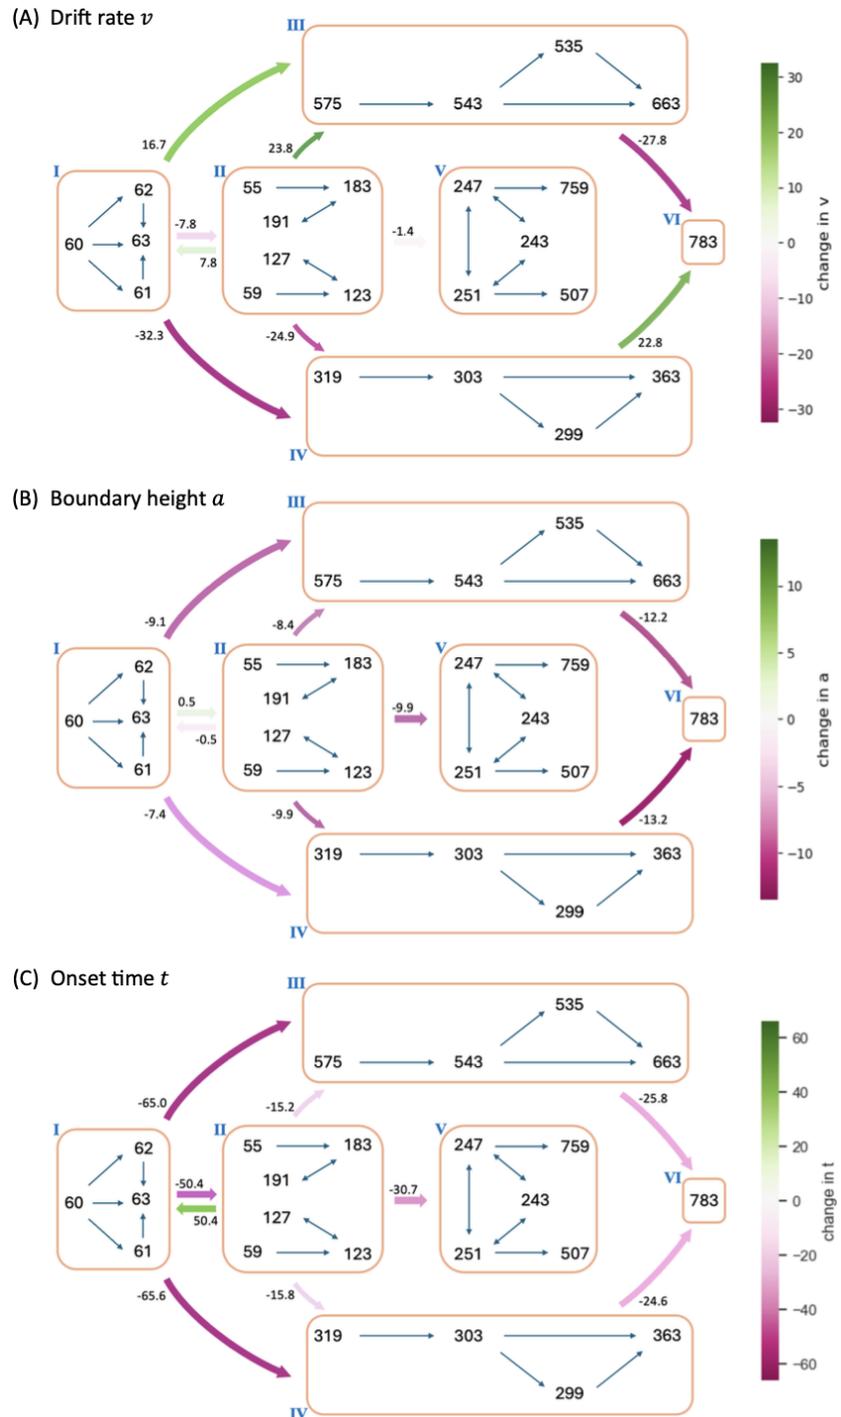


Fig 8. Modulation of decision policy parameters by CBGT activity along the CLAW. The arrows indicate variations of the individual DDM parameters (drift rate v , panel A; boundary height a , panel B; onset time t , panel C) associated with transitions between zones of the CLAW, computed based on the activity changes from each zone to the one immediately downstream from it. The magnitude of each variation is indicated by the coloring (see color bar) and the number at the start of each arrow, representing the percentage change in parameter values.

taken. The effect of varying t is to translate the decision time distribution. The drop in t associated with the progression along the inner CLAW may seem to be at odds with the increase in decision times associated with these pathways (see state colors in Fig 3AB). However, the drop in t reflected the fact that the states in zone II and zone V corresponded to being farther along in the decision process than zones I and II, respectively. When the decision trajectory reached a later stage, e.g., zone V, the process “started” faster because it had already moved past the earlier stages, i.e., zones I and II, and became more prepared. As a result, the onset of the decision-making from this later stage began earlier. The decision time, on the other hand, is measured as the total time for the decision trajectory that visits each zone, starting from the time of trial initiation in zone I and ending in a terminal zone where a decision is made.

Note that when we fit the DDM parameters to each CBGT network’s outputs, we derived a single DDM parameter set that produced responses that most closely matched those of the CBGT network for the whole decision process. These parameters represent a static solution that captures the overall dynamics of decision-making across all zones and transitions for all trials progressing through the network. Then, to link the zone-specific activity patterns to the DDM parameters, we applied the same set of canonical vectors obtained by the CCA, as performed on the static network, to the firing rate changes between zones. Figure 9 provides a schematic illustration of this approach for the transitions from zone I to zones III and IV, respectively. The direct interpretation of this mapping is as follows: if the CBGT firing rates remained static throughout the decision process – resulting in fixed DDM parameters – and we then shifted to a different static set of firing rates (as if a new group of trials occupied a different zone for the entire decision), the canonical vectors would predict how the DDM parameters would need to change to stay consistent with the new activity pattern (Fig 9A-B). In reality, however, CBGT decision output arises from moment-to-moment changes in firing rates throughout each trial, rather than from a fixed set of DDM parameters applied to the entire trial. In other words, trials do not stay within a single zone for the whole decision process but instead transition dynamically between zones as the decision unfolds. Therefore, an accurate interpretation of our “dynamic DDM” analysis is to consider it as the movement of probability mass in the space of possible DDM parameters and DT distributions.

For instance, suppose that the averaged CBGT firing rates are those for zone I. If those stayed fixed, then we would obtain a specific set of DTs and choice probabilities consistent with a specific DDM (i.e., specific set of fit evidence accumulation parameters). Similarly, consider another DDM parameter set that corresponds to the averaged CBGT firing rates for zone III. The difference in the static firing rates between zone I and zone III means that, relative to the first set, the DDM parameters and the probability mass of the DT distribution for the second set differ in a specific way (as depicted in Fig 9B). In the actual case of dynamic decision making, however, firing rate changes occur dynamically as the decision process evolves. Thus, for a decision involving a transition from zone I to zone III, the relevant dynamic decision policy must be intermediate between the two static cases, as illustrated by the magenta color in Fig 9C. For visualization, the figure only shows a possible variation in the drift rate, but the same idea applies to the boundary height and onset time. Although our current analysis does not deliver a specific trajectory of DDM parameters over time (i.e., $v(\tau)$, $a(\tau)$, $t(\tau)$, if time is parameterized by τ), it allows us to estimate the impact of zone transitions on DDM parameters and hence to make a reasonable prediction about how passage through different zones will constrain the DT and choice distribution. Specifically, commitment-related zones induced significant changes in drift rate and boundary height, thus driving fast decisions, while deliberative zones resulted in minimal changes in these parameters and gave rise to slow decisions. Overall, our

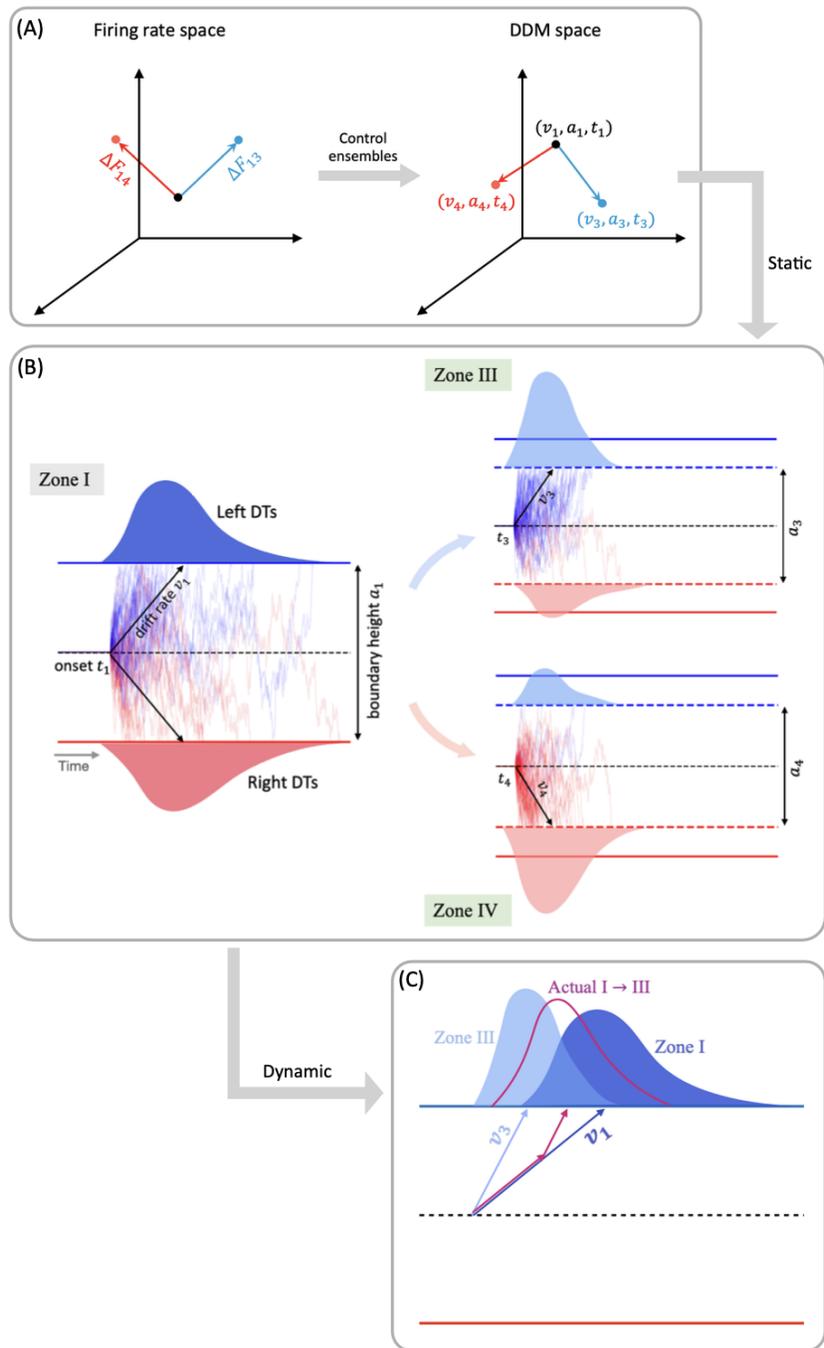


Fig 9. Mapping changes in CBGT activity across CLAW zones to variations in DDM parameters. (A) Firing rate changes during zone transitions are mapped to the DDM space through control ensembles. The vectors ΔF_{13} and ΔF_{14} represent the difference in the firing rate measurements from zone I to zone III (blue) and zone IV (red), respectively. These changes update the *static* DDM parameter set from (v_1, a_1, t_1) to (v_3, a_3, t_3) or (v_4, a_4, t_4) . (B) The corresponding DDM behavior that fits the decision outcomes, assuming that CBGT activity stays in one zone (here, either zone I, zone III, or zone IV) over the entire decision process. (C) Prediction about the actual *dynamic* DDM behavior (magenta) as trials travel from zone I to zone III, the DT distribution of which is constrained by those corresponding to the static cases.

results demonstrate the dynamic nature of decision policies and their evolution across different decision-making scenarios.

Discussion

One of the multifaceted functions of the CBGT circuit is regulation of the evidence accumulation process [3, 16, 46]. Yet, given that the synaptic architecture of the CBGT circuit is intricate, featuring complex interconnected feedforward, reciprocal, and feedback pathways [6], we know very little about how it regulates the flow of information on a moment-by-moment basis during an individual decision. Here we used an integration of multiple computational approaches, including generative spiking neural network models of the CBGT circuit and multiple data analytic techniques, to characterize the evidence accumulation process during individual decisions. Most critical to this was our CLAW framework, using which we were able to capture the flow of activity through CBGT pathways (Fig 3). We then grouped these states into zones to categorize key classes of decision trajectories (Figs 4, 5). Further, by converting the CBGT dynamics into the time evolution of control ensemble drives, we identified that responsiveness and pliancy are the critical ensembles that determine the deliberation period of a decision (Fig 7). Finally, we recast our control ensembles results in terms of DDM parameters and investigated how CBGT activity tunes different aspects of decision policies as the decision process proceeds through different zones of the CLAW (Figs 8, 9). Using this framework, we found that the noisy evidence accumulation process is functionally dissectible into two distinct phases, deliberation and commitment, each governed by competition and cooperation among distinct CBGT subnetworks. In particular, we found that the balance of overall activity across CBGT regions regulates the degree of evidence needed to make a decision early during the process, reflecting control of the deliberation stage. Eventually, a symmetry-breaking allows action-specific direct pathway components to kick in to trigger a transition from deliberation to commitment to a single choice relatively shortly before selection occurs. These findings suggest that the evidence accumulation process during a single decision is highly dynamic and is regulated by a complex interplay of multiple CBGT pathways that play distinct roles in controlling the flow of information during decisions.

Our results broadly confirm the growing perspective [3, 4, 17, 32] that the direct and indirect CBGT pathways engage in a dynamic competition during decision-making, rather than merely functioning in isolation. Instead of a simple switch between “go” and “no-go” signals, the direct–indirect competition implements a decision by continuously balancing the drive for action with the need for deliberation. Our results add to this picture by including the pallidostriatal pathway and highlighting the ways that competition between pathways at key phases of the decision process determines its progression. Specifically, when the system adopts a responsive, pliant state early in the decision (Fig 7AC), the direct pathway associated with one option can gain strength and accelerate the network to action selection (outer CLAW; Figs 3, 4). In other instances, the influence of the indirect pathway becomes dominant, pushing the system toward motor-suppressing states associated with prolonged deliberation (Fig 7BD). During this process, competition between the indirect and pallidostriatal pathways, as well as across action channels, determines the extent of deliberation (inner CLAW; Figs 3, 4). This interactive nature of the CBGT pathways is also critical for encoding uncertainty and for adapting behavior in changing conditions. Our findings indicate that understanding the role of the CBGT circuit in shaping decision-making requires studying the intact dynamical system as a whole.

Previous recent work [22, 23] introduced an “upward mapping approach” used to fit CBGT-driven decision outcomes to a normative DDM model, which led to the

identification of three CBGT control ensembles (choice, responsiveness, and pliancy).
Our work here shows that the control ensemble structure is preserved with the inclusion
of the arky pallidal GPe neurons and demonstrates that the transition from deliberation
and commitment emerges from the evolution of control ensemble activity (cf. Fig 7).
Based on this control ensemble analysis, we translated the changes of CBGT network
activity associated with transitions through CLAW zones into the evolution of
individual DDM parameters, providing a time-dependent algorithmic interpretation of
different decision strategies (see Figs 8, 9). In particular, we observed that fast trials
show an early and rapid reduction in decision boundary height while slow trials
maintain a high decision boundary for a longer time. Unlike traditional DDMs, which
assume that information is accumulated until a fixed threshold is crossed, our results
align with the recently proposed notion of a “boundary collapse” [47–50], where the
decision threshold dynamically adjusts in response to changing task conditions or
sensory evidence. The distinction between fast and slow decisions in our analysis,
however, emphasizes two contrasting decision strategies. On the one hand, the system
may rely on an early and rapid collapse of the decision boundary to push the trial
toward a timely action. This strategy would presumably be critical in situations that
require swift decision-making under time pressure, allowing the system to optimize
speed at the expense of deliberation. On the other hand, the system may devote an
extended time to gradually lowering the evidence accumulation threshold, allowing for
more evaluation of evidence before making a final commitment, thereby prioritizing
accuracy over speed. Hence, the time-varying framework that we have derived serves as
a more complete representation of how decision-making dynamics can be flexibly
tuned to prioritize either speed or accuracy across diverse scenarios.

Moreover, our results reveal that decision policies are highly dynamic during a
decision, extending beyond boundary collapse to include other critical parameters such
as drift rate and onset time, especially given the mixed opinions on the use of dynamic
drift diffusion models [51]. We found that, as evidence accumulates, decisions involving
outer CLAW paths manifest as a strong directional preference (zones III, V, and VI;
Figs 4–8), driving the accumulation process rapidly towards a concluding outcome. In
contrast, decisions that travel through inner CLAW (zones II and V) sustain a slower
and more neutral information flow before committing to a choice. This difference in the
drift rate dynamics between fast and slow decisions highlights the flexibility of the
decision-making system in adjusting its approach to support varying strategies for
making choices. Meanwhile, we observed similar time-evolving patterns of the choice
ensemble drive, which modulated the drift rate parameter, as both fast and slow
decisions neared completion (insets in Fig 8, red curves). This reflects the functional
robustness of the system in terms of maintaining a stable stereotyped process for
decision-making under different conditions.

The results obtained in this work also support the view that prototypic GPe neurons’
connections to STN and GPi form an integral component of the classical indirect
pathway, whereas arky pallidal GPe neurons [33] play a critical role in regulating upward
information processing in the CBGT circuit [41, 52–54]. Complementing experimental
findings on the two subpopulations’ distinct firing rate characteristics [36, 55], our
simulations and analysis suggest new insights into how the two GPe subpopulations may
make divergent contributions to controlling the dynamics of individual decisions. As
shown in Fig 2, the prototypic neurons fire at high rates during resting states, with
many neurons decreasing their activity in response to synaptic drive. In contrast, the
arkypallidal neurons maintain low baseline firing rates during rest, while they rapidly
and robustly increase their activity upon stimulus onset. This divergence in their firing
properties has distinct impacts on the associated basal ganglia targets for each cell type.
For the prototypic neurons, a decrease in firing rate as a decision unfolds fits well with

the classical understanding of the GPe’s role within the indirect pathway, providing a disinhibition of STN and the basal ganglia output nuclei [13, 19, 29]. In our CLAW framework, the second deliberation phase (the loop among states 247, 243, and 251 in Fig 3A), characterized by zero binarized firing rates in one or both prototypic populations, essentially functions as a “brake” on action selection, allowing the system to terminate unwanted actions. This phase closely resembles the dynamics of the classical stop signal task, where the goal is to inhibit a planned action in response to a “stop” signal [56, 57]. This similarity holds even though no trial in our simulations completely ceased or refrained from making a choice, with STN-GPeP subnetwork dynamics evolving such that the trial effectively paused, re-evaluated, and potentially switched between alternative options. Hence, this phase suggests that the prototypic GPe’s brake-like function operates not only during explicit stop signals, but as a more general mechanism for temporal regulation during decision-making.

In contrast to prototypical GPe neurons, the arky pallidal GPe neurons do not project to the STN or other downstream targets of prototypic neurons. Rather they appear to exclusively innervate the striatum via ascending projections [58]. Due to their robust activation upon trial onset, they complement the function of the STN to cancel an action through their direct suppressive effect on firing at the level of striatum. Our analysis shows that in the early deliberation phase (the loop among states 63, 55, 183, and 191 in Fig 3A), the STN acts in the first-tier stage to suppress the basal ganglia output nuclei and pause an incipient response (from 183 to 191). The arky pallidal neurons then become engaged during the second-tier stage to complete the blocking of a response by inhibiting the striatal regions (from 191 to 63). Note that the arky pallidal neurons receive comparatively weak inhibitory striatal inputs compared to their STN excitatory inputs [35, 59], which allows for the recruitment of arky pallidal neurons into the process of resisting commitment when activated by STN. Our results support the emerging standpoint [36, 52, 58, 60] that the GPe, with its heterogeneous collection of cell types, acts not merely as a passive way-station in the indirect pathway, but as an active control mechanism in regulating bidirectional information flow to facilitate flexible execution or blocking of responses [52].

Whether arky pallidal neurons predominantly target SPNs of the indirect or direct pathway remains unclear. It was observed in [61] that the inhibitory signals from arky pallidal neurons have a larger amplitude to iSPNs than dSPNs, with an approximate ratio of 2:1. This suggests a preferential influence over the indirect pathway. Other experimental studies suggested less pronounced bias and more balanced innervation of the two types of striatal populations [58, 62]. Prior modeling work, on the other hand, suggested that a preferential influence of arky pallidal neurons onto dSPNs is crucial for regulating reactive stopping decisions [41]. In our model, in the early deliberation phase, the arky pallidal neurons effectively inhibit both pathways, suppressing the overall progression of information flow in the circuit. This manifests as non-preferential innervation of both types of SPNs by the arky pallidal neurons, subserving an activity pattern akin to an “all stop” signal to the striatum. In the transition from deliberation to commitment, we also found potential roles for inhibition from arky pallidal neurons to both iSPNs and dSPNs. Specifically, the inhibition from arky pallidal neurons to iSPNs in both channels disinhibits the associated dSPNs, facilitating their competition and ultimately leading to a specific response selection (e.g., $247 \rightarrow 759$ and $251 \rightarrow 507$ in Fig 3). Moreover, when a trial becomes stuck in a cycle of indecision – where arky pallidal neurons inhibit both iSPNs and dSPNs (as seen in the three-state reciprocal transitions during the second deliberation phase) – a reduction in dSPN inhibition in either channel can be crucial for breaking the deadlock and allowing the trial to progress toward a decision. Overall, this analysis reveals that the inhibition of striatal populations by arky pallidal neurons plays different roles at

different stages of the decision-making process. Although the relative strengths of the inhibitory pathways from arky pallidal neurons to iSPNs and dSPNs may affect which aspects are dominant, our results suggest that both components contribute to the ongoing competition within the circuit and hence offer a potential reconciliation of prior experimental and theoretical work.

Methodologically, our CLAW analysis relies on binarizing the firing rate within each CBGT nucleus at each time step, using the binarized rates to define activity states, and tracing transitions between states over time. Such a set of states and transitions between them can be called a Boolean network [63]. The Boolean network framework has been heavily used in computational biology, especially to understand dynamics of complex cell signaling pathways [64–66]. The use of this framework has been far less prevalent in neuroscience [67]. Although a recent paper took this approach for the analysis of fMRI data [68], to our knowledge the use of Boolean frameworks for studying the dynamics and evolution of activity in a neural circuit is novel. Our work here highlights the utility of this approach. Harnessing this methodology in future studies could provide an alternative to mean field approaches and other forms of coarse-graining, with a distinct emphasis on discrete states and state transitions that are meaningful for cognitive functions.

Our work suggests several interesting directions for future explorations. We considered a pre-defined action selection task and did not include dopamine-driven corticostriatal plasticity. Given the critical role of dopamine in modulating both the direct and indirect pathway neurons of the striatum [69–73], incorporating plasticity would be expected to significantly alter the pattern observed in our CLAW diagram and will be a key future step. Moreover, we simply considered unbiased evidence for the two available choices, so a natural direction for the extension of this work is to introduce asymmetric evidence in plasticity studies. Another interesting avenue for future inquiry would be to consider the impact of perturbations, consisting of changes in reward contingencies (global) and optogenetic stimulation of specific cell populations (local), on network trajectories and decision outcomes. These extensions could deepen our understanding of how the CBGT circuitry shapes decision processes in response to varying conditions, offering insights into potential therapeutic strategies for disorders related to basal ganglia dysfunction.

Since the CLAW states that we derived from the binarized firing rate data were defined in terms of the activity levels of CBGT populations, our simulation results delivered immediate biological interpretability. The insights from our results suggest several experimentally testable predictions about how, in unfamiliar decision-making tasks that have not yet been learned, distributed control of decisions flows through the CBGT network over time. At the most global level, we predict that especially fast decisions should feature a distinctive early surge in dSPN activity with an associated drop in GPi firing. On the other hand, enhanced iSPN activity and reduced GPeP activity without this dSPN surge should associate with shifts into slow deliberative behavioral states, and should be followed by increased STN and GPeA activity, with an increased level of variability in STN-GPeP circuit activity and intensified GPeA activity together manifesting as a prolonged deliberative phase. This STN-GPeP aspect of our results is consistent with previous modeling work suggesting roles for the STN in action delay [74] and for variable STN-GPeP dynamics in promoting exploration [19, 75]. While the initial increase in iSPN activity and decrease in GPeP activity may predominantly arise in one action channel, the emergence of these changes widely across the relevant nuclei is predicted to signify an entry into a more prolonged, second deliberation phase with variability in the ensuing choices selected.

Optogenetics allows for targeted activation or inhibition of specific neuron populations or types. Our results predict that brief GPeP stimulation or iSPN

inhibition early in the decision-making process will block entry into a deliberative state, allowing dSPN activity to take over and produce a decision. Consistent with the above discussion, we predict that dSPN stimulation and GPeP stimulation should both accelerate decisions; stimulation of these whole populations, at least in the absence of significant experience, should lead to balanced likelihoods of selecting available options. Our results also predict that stimulation of GPeA neurons should yield prolonged deliberation. Finally, a novel prediction of our control ensemble analysis (Fig 7), in light of the composition of the responsiveness and pliancy ensembles (Fig 6), is that even with fast decisions, there should be a discernible rise in activity throughout much of the CBGT network, with the exception of GPeP and GPi, before the channel-specific dSPN surge that induces commitment to a choice. The distinction between fast and slow decisions is predicted to relate not to the presence or absence of this global rise, but rather to its time course and slope (starting earlier but with a more gradual rise for slow decisions) and to whether it is quickly followed by channel-specific dSPN engagement (fast decisions) or not (slow decisions). Overall, we have introduced a new framework, binarization and CLAW, for analyzing the dynamic flow of activity through a neuronal network and have used it to extract novel insights and predictions about CBGT dynamics, including GPeA neurons, during individual decisions. We hope that these advances spur additional experiments and modeling related to the detailed time course of neuronal dynamics during the performance of cognitive functions.

Methods

CBGT network

In this work we simulated behavioral data using a cortico-basal ganglia-thalamic (CBGT) network model. The network consists of 10 distinct populations of spiking model neurons: the cortical interneurons (CxI), the excitatory cortical neurons (Cx), the striatum which includes the D1- and D2-expressing spiny projection neurons (dSPN and iSPN, respectively) and the fast-spiking interneurons (FSI), the globus pallidus external segment, which contains prototypical and arky pallidal neuron subtypes (GPeP and GPeA, respectively), the subthalamic nucleus (STN), the globus pallidus internal segment (GPi), and the thalamus (Th). We considered two groups of neurons for each population, corresponding to the left and right action channels, except for the CxI and FSI, which are shared across both channels. For each decision trial, a choice is selected when the firing rate of the thalamic population within one action channel reaches 30 Hz before the firing rate of the other thalamic population. Figure 1A illustrates the connectivity within the network. Each neuron evolves according to an integrate-and-fire-or-burst model, and the complete model details can be found in [42].

We implemented genetic algorithms (see [23] for details) based on values of 13 selected synaptic connection weights to generate 300 different network configurations that each feature firing rates of all CBGT populations within experimentally observed ranges [22]. In addition, each network exhibited trial timeouts (when no action is selected within 1000 ms) on fewer than 1% of trials, and we checked to ensure that each network was cortico-basal-ganglia driven, with positive correlation between cortical and striatal activity. With each obtained network, we simulated 50 decision trials and gathered temporal firing rate data for all populations across the two action channels.

To unfold the noisy neural firing activity during the dynamic decision-making process, we time-binned the firing rates with $\Delta t = 10$ ms, averaged over neurons in each population, and binarized them based on whether activity was above or below a certain threshold in each time bin. To determine the binarization threshold, we generated firing rate histograms for all trials up to decision times, which showed either unimodal or

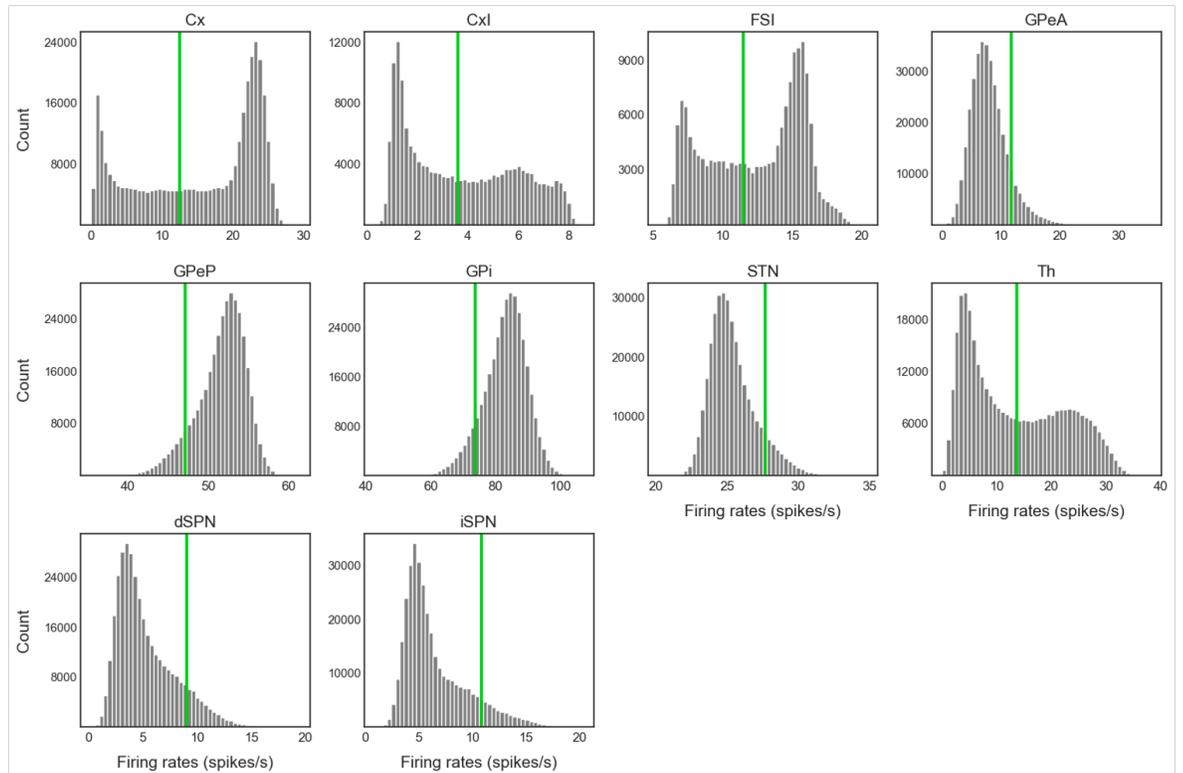


Fig 10. Firing rate histograms for all trials up to decision times. In each panel the vertical green line represents the binarization threshold for the firing rates. In unimodal histograms, the binarization threshold for GPeP and GPi was defined as the firing rate at 10% of the counts, with values above this threshold binarized to 1 and values below it binarized to 0, while for GPeA, dSPN, and iSPN the threshold was set at 90% of the counts. In bimodal histograms, the midpoint between the two peaks was defined as the binarization threshold.

bimodal distributions (Fig 10). For unimodal histograms, the binarization threshold was set at the firing rate corresponding to either 10% (GPeP and GPi) or 90% (GPeA, dSPN, and iSPN) of the total counts, depending on whether the non-stimulated firing rate of each population was above or below its firing rate close to decision times, with values above this threshold binarized to 1 and values below it binarized to 0. For bimodal histograms, the binarization threshold was set at the midpoint between the two peaks. Next we considered $N = 10$ populations of interest — dSPN, iSPN, GPi, GPeP, and Th, for both left and right action channels — and we defined the set $\{s_k\}_{k=1}^{2^N}$ of vectors in \mathbb{R}^N as the base-2 representations of all integers from 1 to 2^N . Each $s_k = [\sigma_{k1}, \sigma_{k2}, \dots, \sigma_{kN}]$, where $\sigma_{kj} \in \{0, 1\}$ for all pairs (k, j) , denotes the k -th unique state of the network. Consequently, we derived a sequence of states from the time-binned, binarized firing rates of every single trial. Using all of these state sequences, we computed the transition probabilities between all possible states. See Fig 11 for a sketch of the above procedures used to process the firing rate data. Finally, we built a chain for the decision dynamics (as depicted in Fig 3) based on the state transitions. Specifically, for each state s_k , we sorted the subsequent n ($n < 2^N$) paths from largest to smallest and retained the largest p ($p < n$) paths until the $(p + 1)$ -th path exhibited a gap of 25% or more relative to the p -th path. This whole approach for the binarization and CLAW reduces firing rate variability and provides a more refined model for the decision dynamics.

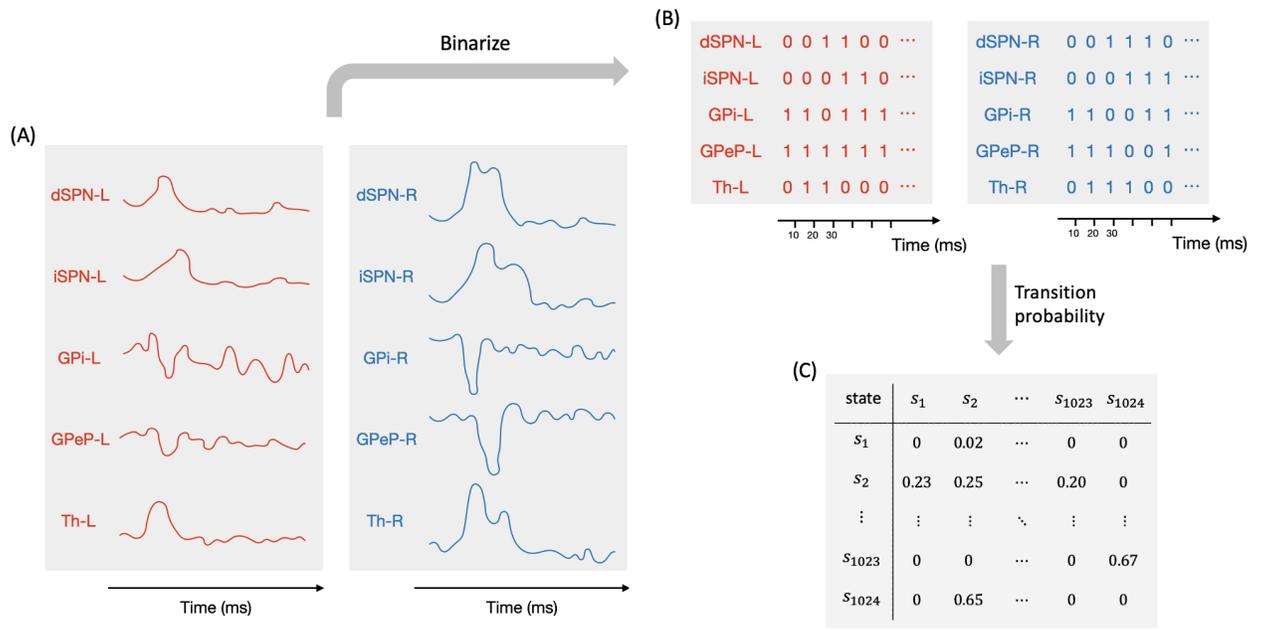


Fig 11. Procedures for processing firing rate data. (A) Populations of interest are specified, with the simulated temporal activity. (B) Firing rates within each population in each time bin (10 ms) are binarized into 1 (above threshold) or 0 (below threshold). (C) Transition probabilities are computed using the full set of binarized activity sequences.

Control ensembles

We replicated the steps of identifying three control ensembles, as introduced previously [22]. First, we mapped the behavior of our CBGT network to the drift-diffusion model (DDM) by fitting the distributions of decision times and choices of the simulated trials, using the Hierarchical Sequential Sampling Modeling (HSSM) toolbox implemented in Python, and obtained the configuration of DDM parameters (boundary height a , drift rate v , onset time t , and starting bias z) corresponding to each of the 300 network configurations. Then, through the application of canonical correlation analysis (CCA) on trial-averaged firing rates $F_{\text{all}} \in \mathbb{R}^{300 \times 18}$ and DDM parameters $D_{\text{all}} \in \mathbb{R}^{300 \times 4}$, we captured pairs of linear combinations within each data set, given by loadings $(u, v) \in \mathbb{R}^{18} \times \mathbb{R}^4$, that yielded the maximal correlation across data sets. This analysis identified three such components associated with distinct decision factors, referred to as choice, responsiveness, and pliancy. The canonical loadings that we obtained are shown in Fig 6; see [22, 23] for additional details on the method.

We converted the time-binned series of CBGT firing rates to a discretized time series of individual control ensemble drives. To do so, for each trial we computed $\Delta F_k \in \mathbb{R}^{18}$, the difference of the averaged firing rates in the 18 model populations between time bins k and $k - 1$. Then, ΔF_k was projected onto the three control ensemble components obtained from the above CCA, via $W_k = \Delta F_k^T U$ where the u loadings form the columns of U , such that W_k contains three elements, each representing the change of the corresponding control ensemble associated with evolving from the $(k - 1)$ -st to the k -th bin. Note that ΔF_0 was set to zero as a baseline, representing the initial state of firing rates before any changes in the decision process occurred. As a result, the firing rates at subsequent time bins ($k > 0$) were adjusted relative to this baseline. The time series of each element of W_k across the trials in each of the four decision groups (fast left, fast right, slow left, and slow right) is displayed in Fig 7.

Following the same approach, we used the trial-averaged firing rate in each CLAW

zone to predict the potential changes in DDM parameters as the CBGT activity evolves across zones. Here we considered the changes in firing rates between each pair of connected zones across all $p \leq 300$ networks that had trials visiting both zones, denoted by $\Delta F_{ij} \in \mathbb{R}^{p \times 18}$ from zone i to zone j , and used these changes to compute $W_{ij} = \Delta F_{ij} U$ where W_{ij} is a p by 3 matrix. The median of each column of W_{ij} indicates the modulation of the associated control ensemble driven by the transition in CBGT activity from zone i to zone j . Finally, we projected W_{ij} to the DDM parameter space via $P_{ij} = W_{ij} V^T$ where the v loadings comprise the columns of V , yielding a p by 4 matrix. The median of each column of P_{ij} reflects the overall change in the corresponding static DDM parameter from zone i to zone j , as shown in Fig 8.

Supporting information

S1 Table. CLAW state details. The table in Fig 3 includes the details of the left-related and neutral states in the upper half of the CLAW, while this table provides the information for all CLAW states, with each pair of left- and right-related states showing symmetry up to the swap of certain L and R channel binary values.

Acknowledgments

We thank Jyotika Bahuguna, Cati Vich Llompart, and Eric Yttri for helpful comments on this work. This work was supported by NIH award R01DA059993 as part of the CRCNS program. JER is also partly supported by NIH award R01NS125814, also part of the CRCNS program.

References

1. Redgrave P, Prescott TJ, Gurney K. The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*. 1999;89(4):1009–1023.
2. Wei W, Wang XJ. Inhibitory control in the cortico-basal ganglia-thalamocortical loop: complex regulation and interplay with memory and decision processes. *Neuron*. 2016;92(5):1093–1105.
3. Dunovan K, Verstynen T. Believer-skeptic meets actor-critic: rethinking the role of basal ganglia pathways during decision-making and reinforcement learning. *Frontiers in neuroscience*. 2016;10:106.
4. Mikhael JG, Bogacz R. Learning reward uncertainty in the basal ganglia. *PLoS computational biology*. 2016;12(9):e1005062.
5. Dunovan K, Vich C, Clapp M, Verstynen T, Rubin J. Reward-driven changes in striatal pathway competition shape evidence evaluation in decision-making. *PLoS computational biology*. 2019;15(5):e1006998.
6. Smith Y, Bevan M, Shink E, Bolam JP. Microcircuitry of the direct and indirect pathways of the basal ganglia. *Neuroscience*. 1998;86(2):353–387.
7. Hollerman JR, Tremblay L, Schultz W. Influence of reward expectation on behavior-related neuronal activity in primate striatum. *Journal of neurophysiology*. 1998;80(2):947–963.

8. Cromwell HC, Schultz W. Effects of expectations for different reward magnitudes on neuronal activity in primate striatum. *Journal of neurophysiology*. 2003;89(5):2823–2838. 905
906
907
9. Yartsev MM, Hanks TD, Yoon AM, Brody CD. Causal contribution and dynamical encoding in the striatum during evidence accumulation. *Elife*. 2018;7:e34929. 908
909
910
10. Ratcliff R. A theory of memory retrieval. *Psychological review*. 1978;85(2):59. 911
11. Gold JI, Shadlen MN. The neural basis of decision making. *Annu Rev Neurosci*. 2007;30(1):535–574. 912
913
12. Alexander GE, DeLong MR, Strick PL. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual review of neuroscience*. 1986;9(1):357–381. 914
915
916
13. Alexander GE, Crutcher MD. Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in neurosciences*. 1990;13(7):266–271. 917
918
919
14. Nakano K, Kayahara T, Tsutsumi T, Ushiro H. Neural circuits and functional organization of the striatum. *Journal of neurology*. 2000;247:V1–V15. 920
921
15. Klaus A, Martins GJ, Paixao VB, Zhou P, Paninski L, Costa RM. The spatiotemporal organization of the striatum encodes action space. *Neuron*. 2017;95(5):1171–1180. 922
923
924
16. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*. 2006;113(4):700. 925
926
927
17. Bariselli S, Fobbs W, Creed M, Kravitz A. A competitive model for striatal action selection. *Brain research*. 2019;1713:70–79. 928
929
18. Gurney K, Prescott TJ, Redgrave P. A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological cybernetics*. 2001;84:401–410. 930
931
932
19. Chakravarthy VS, Joseph D, Bapi RS. What do the basal ganglia do? A modeling perspective. *Biological cybernetics*. 2010;103:237–253. 933
934
20. Tecuapetla F, Jin X, Lima SQ, Costa RM. Complementary contributions of striatal projection pathways to action initiation and execution. *Cell*. 2016;166(3):703–715. 935
936
937
21. Herz DM, Bange M, Gonzalez-Escamilla G, Auer M, Ashkan K, Fischer P, et al. Dynamic control of decision and movement speed in the human basal ganglia. *Nature Communications*. 2022;13(1):7530. 938
939
940
22. Vich C, Clapp M, Rubin JE, Verstynen T. Identifying control ensembles for information processing within the cortico-basal ganglia-thalamic circuit. *PLOS Computational Biology*. 2022;18(6):e1010255. 941
942
943
23. Bahuguna J, Verstynen T, Rubin JE. How cortico-basal ganglia-thalamic subnetworks can shift decision policies to maximize reward rate. *bioRxiv*. 2024;. 944
945
24. Schouten J, Bekker J. Reaction time and accuracy. *Acta psychologica*. 1967;27:143–153. 946
947

25. Chittka L, Skorupski P, Raine NE. Speed–accuracy tradeoffs in animal decision making. *Trends in ecology & evolution*. 2009;24(7):400–407. 948 949
26. Cohen JD, McClure SM, Yu AJ. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2007;362(1481):933–942. 950 951 952
27. Mehlhorn K, Newell BR, Todd PM, Lee MD, Morgan K, Braithwaite VA, et al. Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*. 2015;2(3):191. 953 954 955
28. Ratcliff R, McKoon G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*. 2008;20(4):873–922. 956 957
29. Albin RL, Young AB, Penney JB. The functional anatomy of basal ganglia disorders. *Trends in neurosciences*. 1989;12(10):366–375. 958 959
30. DeLong MR. Primate models of movement disorders of basal ganglia origin. *Trends in neurosciences*. 1990;13(7):281–285. 960 961
31. Wei W, Rubin JE, Wang XJ. Role of the indirect pathway of the basal ganglia in perceptual decision making. *Journal of Neuroscience*. 2015;35(9):4052–4064. 962 963
32. Dunovan K, Lynch B, Molesworth T, Verstynen T. Competing basal ganglia pathways determine the difference between stopping and deciding not to go. *Elife*. 2015;4:e08723. 964 965 966
33. DeLong MR. Activity of pallidal neurons during movement. *Journal of neurophysiology*. 1971;34(3):414–427. 967 968
34. Mallet N, Pogosyan A, Márton LF, Bolam JP, Brown P, Magill PJ. Parkinsonian beta oscillations in the external globus pallidus and their relationship with subthalamic nucleus activity. *Journal of neuroscience*. 2008;28(52):14245–14258. 969 970 971
35. Nevado-Holgado AJ, Mallet N, Magill PJ, Bogacz R. Effective connectivity of the subthalamic nucleus–globus pallidus network during Parkinsonian oscillations. *The Journal of physiology*. 2014;592(7):1429–1455. 972 973 974
36. Dodson PD, Larvin JT, Duffell JM, Garas FN, Doig NM, Kessar N, et al. Distinct developmental origins manifest in the specialized encoding of movement by adult neurons of the external globus pallidus. *Neuron*. 2015;86(2):501–513. 975 976 977
37. Ketzef M, Silberberg G. Differential synaptic input to external globus pallidus neuronal subpopulations in vivo. *Neuron*. 2021;109(3):516–529. 978 979
38. Mallet N, Schmidt R, Leventhal D, Chen F, Amer N, Boraud T, et al. Arkypallidal cells send a stop signal to striatum. *Neuron*. 2016;89(2):308–316. 980 981
39. Schmidt R, Berke JD. A Pause-then-Cancel model of stopping: evidence from basal ganglia neurophysiology. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2017;372(1718):20160202. 982 983 984
40. Diesburg DA, Wessel JR. The Pause-then-Cancel model of human action-stopping: Theoretical considerations and empirical evidence. *Neuroscience & Biobehavioral Reviews*. 2021;129:17–34. 985 986 987

41. Giossi C, Bahuguna J, Rubin JE, Verstynen T, Vich C. Arkypallidal neurons in the external globus pallidus can mediate inhibitory control by altering competition in the striatum. *Proceedings of the National Academy of Sciences*. 2024;121(47):e2408505121. 988-991
42. Clapp M, Bahuguna J, Giossi C, Rubin JE, Verstynen T, Vich C. CBGTPy: An extensible cortico-basal ganglia-thalamic framework for modeling biological decision making. *bioRxiv*. 2024; p. 2023–09. 992-994
43. Cisek P, Kalaska JF. Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron*. 2005;45(5):801–814. 995-997
44. Rubin JE, Vich C, Clapp M, Noneman K, Verstynen T. The credit assignment problem in cortico-basal ganglia-thalamic networks: A review, a problem and a possible solution. *European Journal of Neuroscience*. 2021;53(7):2234–2253. 998-1000
45. Bondy JA, Murty USR. *Graph theory with applications*. vol. 290. Macmillan London; 1976. 1001-1002
46. Bogacz R, Larsen T. Integration of reinforcement learning and optimal decision-making theories of the basal ganglia. *Neural computation*. 2011;23(4):817–851. 1003-1005
47. Milosavljevic M, Malmaud J, Huth A, Koch C, Rangel A. The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision making*. 2010;5(6):437–449. 1006-1009
48. Bowman NE, Kording KP, Gottfried JA. Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. *Neuron*. 2012;75(5):916–927. 1010-1011
49. Hawkins GE, Forstmann BU, Wagenmakers EJ, Ratcliff R, Brown SD. Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*. 2015;35(6):2476–2484. 1012-1014
50. Palestro JJ, Weichart E, Sederberg PB, Turner BM. Some task demands induce collapsing bounds: Evidence from a behavioral analysis. *Psychonomic bulletin & review*. 2018;25:1225–1248. 1015-1017
51. Ratcliff R, Smith PL, Brown SD, McKoon G. Diffusion decision model: Current issues and history. *Trends in cognitive sciences*. 2016;20(4):260–281. 1018-1019
52. Giossi C, Rubin JE, Gittis A, Verstynen T, Vich C. Rethinking the external globus pallidus and information flow in cortico-basal ganglia-thalamic circuits. *European Journal of Neuroscience*. 2024;. 1020-1022
53. Courtney CD, Pamukcu A, Chan CS. Cell and circuit complexity of the external globus pallidus. *Nature Neuroscience*. 2023; p. 1–13. 1023-1024
54. Nambu A, Chiken S. External segment of the globus pallidus in health and disease: Its interactions with the striatum and subthalamic nucleus. *Neurobiology of Disease*. 2024;190:106362. 1025-1027
55. Abdi A, Mallet N, Mohamed FY, Sharott A, Dodson PD, Nakamura KC, et al. Prototypic and arkypallidal neurons in the dopamine-intact external globus pallidus. *Journal of Neuroscience*. 2015;35(17):6667–6688. 1028-1030

56. Nambu A, Tokuno H, Takada M. Functional significance of the cortico-subthalamo-pallidal ‘hyperdirect’ pathway. *Neuroscience research*. 2002;43(2):111–117. 1031
1032
1033
57. Aron AR, Poldrack RA. Cortical and subcortical contributions to stop signal response inhibition: role of the subthalamic nucleus. *Journal of Neuroscience*. 2006;26(9):2424–2433. 1034
1035
1036
58. Mallet N, Mickleman BR, Henny P, Brown MT, Williams C, Bolam JP, et al. Dichotomous organization of the external globus pallidus. *Neuron*. 2012;74(6):1075–1086. 1037
1038
1039
59. Chuhma N, Tanaka KF, Hen R, Rayport S. Functional connectome of the striatal medium spiny neuron. *Journal of Neuroscience*. 2011;31(4):1183–1192. 1040
1041
60. Aristieta A, Barresi M, Lindi SA, Barriere G, Courtand G, de La Crompe B, et al. A disynaptic circuit in the globus pallidus controls locomotion inhibition. *Current Biology*. 2021;31(4):707–721. 1042
1043
1044
61. Glajch KE, Kolver DA, Hegeman DJ, Cui Q, Xenias HS, Augustine EC, et al. Npas1+ pallidal neurons target striatal projection neurons. *Journal of Neuroscience*. 2016;36(20):5472–5488. 1045
1046
1047
62. Mastro KJ, Zitelli KT, Willard AM, Leblanc KH, Kravitz AV, Gittis AH. Cell-specific pallidal intervention induces long-lasting motor recovery in dopamine-depleted mice. *Nature neuroscience*. 2017;20(6):815–823. 1048
1049
1050
63. Albert R, Barabási AL. Dynamics of complex systems: Scaling laws for the period of Boolean networks. *Physical Review Letters*. 2000;84(24):5660. 1051
1052
64. Bornholdt S. Boolean network models of cellular regulation: prospects and limitations. *Journal of the Royal Society Interface*. 2008;5(suppl_1):S85–S94. 1053
1054
65. Wang RS, Saadatpour A, Albert R. Boolean modeling in systems biology: an overview of methodology and applications. *Physical biology*. 2012;9(5):055001. 1055
1056
66. Saadatpour A, Albert R. Boolean modeling of biological regulatory networks: a methodology tutorial. *Methods*. 2013;62(1):3–12. 1057
1058
67. Hopfield JJ. Brain, neural networks, and computation. *Reviews of modern physics*. 1999;71(2):S431. 1059
1060
68. Bertacchini F, Scuro C, Pantano P, Bilotta E. Modelling brain dynamics by boolean networks. *Scientific Reports*. 2022;12(1):16543. 1061
1062
69. Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997;275(5306):1593–1599. 1063
1064
70. Schultz W. Multiple reward signals in the brain. *Nature reviews neuroscience*. 2000;1(3):199–207. 1065
1066
71. Reynolds JN, Wickens JR. Dopamine-dependent plasticity of corticostriatal synapses. *Neural networks*. 2002;15(4-6):507–521. 1067
1068
72. Rubin JE, McIntyre CC, Turner RS, Wichmann T. Basal ganglia activity patterns in parkinsonism and computational modeling of their downstream effects. *European Journal of Neuroscience*. 2012;36(2):2213–2228. 1069
1070
1071

73. Dobbs LK, Kaplan AR, Lemos JC, Matsui A, Rubinstein M, Alvarez VA. Dopamine regulation of lateral inhibition between striatal neurons gates the stimulant actions of cocaine. *Neuron*. 2016;90(5):1100–1113. 1072
1073
1074
74. Frank MJ. Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*. 2006;19(8):1120–1136. 1075
1076
75. Kalva SK, Rengaswamy M, Chakravarthy VS, Gupte N. On the neural substrates for exploratory dynamics in basal ganglia: a model. *Neural networks*. 2012;32:65–73. 1077
1078
1079