bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Dual learning mechanisms drive adaptive inhibitory control

Kyle Dunovan[1,2], Timothy Verstynen[1,2]

1. Dept. Psychology, Carnegie Mellon University
2. Center for the Neural Basis of Cognition, Carnegie Mellon University

Corresponding Author:
Timothy Verstynen
Dept. of Psychology & Center for the Neural Basis of Cognition
436D Baker Hall
Carnegie Mellon University
Pittsburgh, PA 15213
www.cognitiveaxon.com

Keywords: inhibitory control; reinforcement learning; decision-making; basal ganglia; computational modeling;

Acknowledgments: The authors would like to thank Jonathan Rubin, Julie Fiez, and Natasha Tockowicz for their helpful comments on early drafts of this manuscript and Leo Scholl and Tara Molesworth for assistance with data collection. This research was funded by the Pennsylvania State Department of Health Formula Award SAP4100062201 and National Science Foundation CAREER Award #1351748.

Competing Interests: The authors report no competing interests.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Abstract

Goal-directed behavior requires integrating action selection processes with learning systems that adapt control using environmental feedback. These functions intersect in the basal ganglia (BG), which has at least two targets of plasticity: a dopaminergic modulation of striatal pathways and cortical modulation of the subthalamic nucleus (STN). Dual learning mechanisms suggests that feedback signals have a multifaceted impact on BG-dependent decisions. Using a hybrid of accumulation-to-bound decision models and reinforcement learning, we modeled the performance of humans in a stop-signal task where participants (N=75) learned the prior distribution of the timing of a stop signal through trial-and-error feedback. Changes in the drift-rate of the action execution process were driven by errors in action timing, whereas adaptation in the boundary height served to increase caution following failed stops. These findings highlight two interactive learning mechanisms for adapting the control of goal-directed actions based on dissociable dimensions of feedback error.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Introduction

Environmental uncertainty demands that goal-directed actions be executed with a certain degree of caution, requiring agents to strike the appropriate balance between speed and control based on internal goals and contextual constraints. Because of the pervasive and dynamic nature of uncertainty in the real world, the precise amount of behavioral control is rarely a known quantity, and must therefore be learned through trial-and-error. Despite the underlying intuition that intelligent behavior requires both the ability to control actions as well as to learn from consequent feedback, it is not well understood how these learning and control systems interact. One reason for this is that computational models of learning (Frank & Badre, 2012; Sutton & Barto, 1998) and control (Brown & Heathcote, 2008; Schall, Palmeri, & Logan, 2017; Verbruggen & Logan, 2009) have historically emerged from disparate lines of empirical research (see Bogacz & Larsen (2011) and Pedersen, Frank, & Biele (2016) for exceptions), adding difficulty to the already challenging task of inferring cognitive phenomena from gross behavioral measures. Recently, however, insights from cognitive and computational neuroscience have begun to shed light on the interaction of cognitive processes in neural circuits, providing additional empirical anchors for grounding theoretical assumptions (Frank et al., 2015; Heitz & Schall, 2013; Purcell & Palmeri, 2016; Turner, van Maanen, & Forstmann, 2015).

The basal ganglia (BG), a subcortical network that shares dense, reciprocal connections with much of cortex, as well as many subcortical neuromodulators, is known to play a critical role in both learning and control processes - acting as a central integration hub where cortically-distributed control commands (Dunovan, Lynch, Molesworth, & Verstynen, 2015; Majid, Cai, Corey-Bloom, & Aron, 2013; Mallet et al., 2016), sensory evidence (Ding & Gold, 2010, 2012), and decision variables (Forstmann et al., 2008; Herz, Zavala, Bogacz, & Brown, 2016; Keuken et al., 2015) can be weighed and synthesized with feedback-dependent learning signals to facilitate goal-directed behavior (Bogacz & Gurney, 2007; Bogacz & Larsen, 2011; Gurney, Humphries, & Redgrave, 2015). Cortical

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

information enters the BG through three primary pathways: the first two of these pathways, entering via the striatum, are the direct (i.e., facilitating; Figure 1A, green) and indirect pathways (i.e., suppressing; Fig 1A, blue). The third, hyperdirect (i.e., braking; Fig. 1A, red) pathway, enters via the subthalamic nucleus (STN). Given the various sites of structural overlap between the direct and indirect pathways, including their convergence in the output of the BG (Smith, Bevan, Shink, & Bolam, 1998), we recently proposed that, rather than acting as independent "go" and "no-go" levers, these pathways engage in dynamic competition for control over BG-output (Dunovan & Verstynen, 2016). From this perspective, the direct and indirect pathway competition over action execution can be thought of as debate between a "believer" and a "skeptic" (Fig. 1B): activation of the direct pathway increases with evidence supporting action execution and gradually suppresses skepticism of the indirect pathway until a threshold is reached and the action can be gated (Fig. 1B). In contrast, increasing activation of the indirect pathway can proactively suppress or delay action execution in the context of high uncertainty, simultaneously affording more time to reactively cancel action cancellation in response to external stop cues that are mediated by signals from the hyperdirect pathway (Dunovan et al., 2015; Wei & Wang, 2016).

A reasonable corollary of the observed convergence of the direct, indirect, and hyperdirect pathways in the output of the BG is that proactive speeding of execution decisions should also influence the efficacy at which the hyperdirect pathway is able to brake the planned execution. Likewise, proactively slowing the accumulation of "go" activity should not only slow the timing of the execution, but also support easier reactive inhibition by reducing the buildup of "go" activity that must be overcome by the hyperdirect pathway in order to brake the action. To test this intuition, we proposed a variant of the classic independent race model in which reactive braking signals, delivered through the hyperdirect pathway, are assumed to be functionally dependent on the current state of the competition between the direct and indirect pathways. In this nested process model (Fig. 1C) the drift-rate (v_E) of the execution process captures the relative activation of

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

"believer" and "skeptic" populations, leading to a "go" response when the execution process crosses the upper threshold (*a*), reflecting combined strength of motor-inhibiting forces on basal-ganglia output (e.g., hyperdirect pathway, basal firing rate of BG output nucleus). In the event of a stop-cue, a second braking process is instantiated at the current state of the execution process and must reach the bottom boundary before the execution threshold is reached in order to cancel motor output. This nested dependency of the braking process on the accumulating execution process imposes a speed-control tradeoff, such that increasing the execution drift-rate has direct negative consequences on control, making it more difficult for the braking process to override and cancel action output. While this nested process model was able to account for the relationship between proactive control and reactive stopping ability, as well as predict BG output during proactive control (Dunovan et al., 2015), it fails to account for situations in which the control demands are unknown and thus, must be learned from environmental feedback.

Converging lines of evidence from experimental optogenetics (Yttri & Dudman, 2016) and advances in computational modeling (Wei & Wang, 2016) suggest that both primary input structures of the BG, the striatum and STN, are critical for guiding adaptive behavior. For example, similar to classic theories of action-value learning in the striatum, Yttri et al. (2016) found that optogenetic reinforcement of cortical input to direct and indirect pathways led to incremental changes in movement velocity, and that this effect disappeared in the presence of dopamine antagonist. Outside the striatum, multiple lines of evidence now point to the STN as a major source of behavioral adaptation in the BG (Brittain et al., 2012; Cavanagh, Sanguinetti, Allen, Sherman, & Frank, 2014; Frank et al., 2015; Herz et al., 2016). In line with the notion of parallel learning systems, Wei and Wang (2016) showed in a spiking model of the cortico-BG circuitry that plasticity in the striatal and subthalamic pathways gives rise to feedback dependent changes in inhibitory control. Consistent with the nested process model proposed in Dunovan et al., (2015), increasing the connectivity strength of striatal projections

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

in this network model both increased the speed of the execution decision and delayed the stop-signal reaction time. Indeed, this observation lends biological support to the assumptions displayed in Figure 1C about the nested nature of proactive and reactive control signals within the BG. More importantly, however, it suggests distinct functional roles for the two plasticity systems in the BG. First, plasticity in the striatum contributes to adaptive inhibitory control by modulating the drift-rate of the execution process. The STN, on the other hand, should exert control on the height of the execution boundary, delaying action execution in contexts of high uncertainty (Herz et al., 2016). Indeed, Cavanaugh et al., (2014) found that activity in the STN tracked the degree to which subjects slowed responding after committing an error and that this behavioral phenomenon was described by a diffusion model in which errors led to an increase in threshold on subsequent trials. Taken together, these two BG-dependent learning mechanisms predict two means of adapting action control to environmental feedback: a mechanism for optimizing the timing, or vigor, of action execution as well as a post-error slowing mechanism to impose caution on future action decisions.

In order to test the predictions of this dual mechanism hypothesis, we modeled performance of human participants in an adaptive version of the stop-signal task. Specifically, we show how a straightforward hybridization of the nested process model (Dunovan et al., 2015; Dunovan & Verstynen, 2016) with principles of reinforcement learning (Sutton & Barto, 1998), captures feedback-dependent changes in RT and stop accuracy through adaptive changes in the execution drift-rate and boundary height. Drawing on wide ranging evidence from previous studies (Brody & Hanks, 2016; Cavanagh et al., 2014; Ding & Gold, 2010, 2012; Dunovan et al., 2015; Dunovan & Verstynen, 2016; Frank et al., 2015; Herz et al., 2017, 2016; Wei & Wang, 2016; 2016), we argue that these behaviorally derived signatures of learning are consistent with dopaminergic tuning of corticostriatal pathways (e.g., drift-rate adaptation) and error-related changes in corticosubthalamic pathways (e.g., boundary height adaptation).

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Results

Inhibitory control adapts to contextual statistics

Subjects performed an anticipatory version of the stop-signal task (see Figure 2A) similar to that reported previously (Dunovan et al., 2015). On Go trials (n=600), subjects were instructed to make a button press as soon as a vertically rising bar intersected a line near the top of the screen (always occurring at 520 ms) and received a feedback score telling them how early or late their response was (max 100+ pts). On Stop trials, the bar stopped and turned red before reaching the line, signaling the subject to withhold their response. On Context Stop trials (n=200) the stop-signal delay (SSD), defined as the time between trial onset and initiation of the stop-signal, was sampled from one of three probability distributions (Fig. 2B) depending on which group the subject was assigned to. In the Uniform context group, the SSD was sampled uniformly between 10 and 510 ms. In the Early and Late context groups, the SSD was sampled from a normal distribution with means at 250 ms and 350 ms, respectively, both with a standard deviation of 35 ms. To evaluate the effects of Context SSDs on stopping accuracy, Probe Stop trials (n=80) were randomly interspersed throughout all trial blocks with SSDs fixed at 200, 250, 300, 350, or 450 ms (16 trials each). In this paradigm, the Context SSDs establish a prior distribution on the probability of the stop-signal timing, which is known to influence the dynamics of inhibitory control (Shenoy & Yu, 2011). Shifting the mean of the prior distribution of SSDs later in the trial introduces a greater degree of conflict between Go and Stop trial goals. For instance, the majority of stop-signals in Late context occur too late in the trial to effectively cancel a movement intended to occur at 520 ms, requiring subjects to delay their responding on Go trials to achieve the same level of accuracy as subjects in the Early context.

To assess behavioral differences across contexts, we compared accuracy on stop-signal trials at each Probe SSD across groups as well as the mean RTs on correct (response on Go trial) and error (i.e., response on Stop trial) responses. Separate one-way ANOVAs revealed a significant main effect of context across

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

groups on both correct RTs, $F(2,72)=10.07$, $p<.001$, and error RT (responses on stop-signal trials), $F(2,72)=21.72$, $p<.00001$. Consistent with our hypothesis, we found a significant interaction between context condition and Probe SSD, $F(2.23,80.15)=3.60$, $p=.027$ (Fig. 3A). Shifting the mean of the prior distribution on the stop signal later into the trial led to delayed responding on Go trials as well as greater stopping success across Probe trial SSD's, exhibited by the rightward shift in the stop-curve and RT distributions in the Uniform and Late groups relative to the Early group. Thus, as predicted, participants can reliably learn to modulate their inhibitory control efficiency based on the probabilistic structure of prior stop-signal timing (see also Shenoy & Yu (2011)).

We next examined whether failed stop trials elicited any systematic changes in RT on subsequent trials. Figure 3B shows the immediate slowing and subsequent decay in RTs following a stop error, calculated with respect to the average RT on the ten trials that preceded the error. A one-way ANOVA revealed a significant effect of Context on the degree to which subjects slowed responses immediately following stop errors, $F(2,72)=4.27$, $p=.018$. Unlike the observed effects on RT and accuracy, which scaled with differences in the mean SSD in each Context, group differences in post-error slowing appeared to be driven by the variance of SSDs, with stop errors eliciting greater slowing in the Uniform context than in the Early and Late contexts (Fig. 3C). Collectively, these findings suggest that adaptive control is sensitive to multiple task dimensions and that these dimensions manifest in dissociable behavioral profiles.

Contextual modulation of decision process

The effect of recent experience on inhibitory control suggests that participants learn a prior on the probability of the SSD through trial-and-error (see also (Shenoy & Yu, 2011)), incorporating this prior into future decisions. In order to understand what stages of the decision process are targets of this learning, we fit the average stop-accuracy on Probe trials and correct and error RT distributions when responses are made, with the nested process model (as described in

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

(Dunovan et al., 2015)). Briefly, the nested process model assumes that an execution decision is made when an accumulating execution process crosses an upper decision boundary. On Stop trials, a nested braking process is initiated at the current state of the execution process at the time of the SSD and accumulates back towards the lower boundary with negative drift (see Fig. 1C). The model successfully cancels an action when the braking process reaches the lower bound before execution process terminates at the upper bound.

We previously found that contextual manipulations of the penalty structure and stop cue probability led to behavioral changes that were best described by modulations in the execution drift-rate when compared to other models with a single free parameter free to vary between conditions (Dunovan et al., 2015). To reduce the combinatorial space of possible model configurations, we first sought to replicate this finding by comparing alternative models in which only one parameter was free to vary across context condition (see Table 2) - either execution drift (υ_E), braking drift (υ_B), urgency (γ), or boundary height (α). Each model was evaluated using two complexity-penalized goodness-of-fit statistics: Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). A difference of 7-10 in the information criteria (IC) values for two models provides strong support for the model with the lower value. In line with our previous findings (Dunovan et al., 2015), leaving the execution drift-rate free provided a better account of Context-dependent changes in behavior compared to alternative single-parameter models (Best-Fit AICv_e=-369.26; Fig. 4A). The next best fit was afforded by allowing the urgency to vary across conditions (|AICv_E –AICγ| = 15.83).

To further test the relationship between execution drift-rate and context, we performed another round of fits to test for possible interactions between the execution drift-rate and a second free parameter, either boundary height (a), braking drift-rate (υ_B), or urgency gain (γ). The AIC and BIC scores from these fits showed that a combination of execution drift and boundary height (υ_E & a) provided the best overall fit to the data (Best Fit AICa,v_e = -381.52), reasonably

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

exceeding that of the drift-only model ($|AIC\upsilon_E - AIC\gamma| = 12.26$) to justify the added complexity of the dual parameter model. Figure 4B shows a qualitative assessment of the a & υ_E model's goodness of fit, revealing a high degree of overlap between the simulated and observed stop-accuracy and RT data in both Early and Late conditions. These results suggest that there may be two targets of learning in the decision process: a strong modulation of the execution drift rate and a more subtle modulation of the boundary height.

Dual learning mechanisms

It is not clear from the preceding analysis whether error-driven changes in the drift rate and boundary height are able to capture trial-to-trial adjustments of response speed and stop accuracy as statistics of the environment are learned experientially. Here we explore how drift-rate and boundary height mechanisms adapt on a trial-wise basis to different sources of feedback to drive context-dependent control and decision-making.

We implemented two forms of corrective learning - one targeting the execution drift-rate (υ_E) and another targeting the height of the decision boundary (a). On correct Go trials, the drift-rate parameter was updated according to signed RT errors calculated with respect to the Target RT (T^G) of 520ms, reducing the drift-rate to slow actions following a "fast" response (Fig. 5A left; $RT_t < 520$ ms) and increasing the drift-rate to speed future actions following a "slow" response (Fig. 5A middle, $RT_t > 520$ ms). This form of RT-dependent modulation in the drift-rate is motivated by recent findings demonstrating adaptation of action velocity by dopaminergic prediction error signaling in the striatum (Yttri & Dudman, 2016). In the context of the "believer-skeptic" framework (Dunovan & Verstynen, 2016), fast RT errors could reinforce the "skeptic" (i.e., indirect pathway) and suppress the "believer" (i.e., direct pathway) by decreasing dopaminergic tone in the striatum.

On failed stop trials (t_{err}) the boundary height was increased on the following trial (t_{err}+1) according to a delta function (Fig. 5B) and decayed

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

exponentially on each successive trial until reaching its baseline value or another failed stop was committed (see Methods for details). This form of adaptation in boundary height is motivated by physiological evidence that the STN plays a critical role in setting threshold for action execution and that this relationship is modulated by error commissions (Cavanagh et al., 2014). In Figure 5C, the model-predicted change in Go trial RT and Stop trial accuracy across the experiment is overlaid on the corresponding empirical measures, demonstrating a robust model fit to the observed data. To confirm that the trial-averaged behavior of the model was preserved after fitting the learning rates, the stop-accuracy curve and RT statistics were calculated from simulations of the adaptive model and overlaid on the average data from the Uniform condition (Figure 5D). The model's predictions are indeed closely aligned with all empirical statistics (adaptive χ^2=.0082, static χ^2 = .01). While this is not necessarily surprising, it is promising to confirm that introducing feedback-dependent adaptation in the drift-rate and boundary height parameters does not compromise the model's fit to trial averaged statistics.

The fits to the Uniform condition show two possible mechanisms for acquiring the prior on the SSD: adaptive modulation of response speed by the drift-rate and error-driven changes to boundary height. In order to confirm that these mechanisms work together to adaptively learn based only on the statistics of previous input signals, we took the average parameter scheme from the Uniform condition fits and simulated each subject in the Early and Late groups. If the context-dependent changes in the RT distributions and stop-accuracy are indeed a reflection of the proposed learning mechanisms, then the model simulations should reveal similar RT and accuracy time courses as in the observed behavior.

Figure 6A shows the simulated stop-curve and RT distributions generated by the adaptive model based on feedback in the Early and Late conditions. As in the observed data (Fig. 3A), adaptation to Early SSDs led to impaired stopping accuracy, but faster RT's relative to simulated predictions in the Late condition. In Figure 6B-C, the middle panels show the same trial-binned RT and stop-

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

accuracy means as in Figure 5C (Uniform condition), flanked by corresponding time courses from simulations to Early (left) and Late (right) conditions. The adaptive model predictions show a high degree of flexibility, conforming to idiosyncratic changes in the trial-wise behavioral dynamics within each Context SSD condition. For instance, the RTs in the Early condition exhibit a relatively minor and gradual decay over the course of the experiment (Figure 6B, left), contrasting markedly from the early increase and general volatility of RTs in the Late condition (Figure 6B, right). The adaptive model largely captures both patterns, underscoring feedback-driven adaptation in the drift-rate as a powerful and flexible tool for commanding inhibitory control across a variety of settings. In addition to predicting group differences in the time course of RTs, the simulations in Figure 6C show a striking degree of precision in the model-estimated changes in stop-accuracy, both over time and between groups.

Because the static model fits revealed marginal evidence for the drift-only model (Fig. 4A), we next asked whether this simpler model was able to account for the learning-related behavioral changes (see Fig. 5B) with the same precision as the dual learning (i.e., drift and boundary) model. To test this hypothesis, we ran simulations in which the boundary learning rate was set to zero, thereby leaving only the drift-rate free to vary in response to feedback. Figure 7A shows the error between observed and model-predicted estimates for each of the behavioral measures in Figure 2 (e.g., RT, stop accuracy, and post-error slowing) based on 25 simulations of the drift-only and dual learning models. Compared to the drift-only model, the dual learning model showed no significant benefits in terms of fit to the trial-wise RT, $t(24)=1.09$, $p=.28$, or accuracy $t(24)=.23$, $p=.82$, but showed a marked improvement in the fit to post-error slowing, $t(24)=-6.91$, $p<.00001$ (Fig. 7A). Importantly, the interaction of drift-rate and boundary adaptation in the dual learning model not only reduced the error in the model fit, but recovered the same qualitative pattern of post-error slowing across Contexts observed in the data (Fig. 7B). In contrast, the drift-only model predicted the largest post-error slowing effect in the Early condition (Fig. 7B, left). This is

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

particularly revealing since no information about the observed post-error slowing was included in the adaptive cost function when fitting the learning-rate parameters. Collectively, these results suggest that goal-directed tuning of movement timing (RT) and control (stop accuracy) is best described by feedback-driven changes in the drift-rate and boundary-height parameters of accumulation-to-bound decisions.

Discussion

The current results demonstrate the existence of two separable, yet interacting, learning mechanisms for adapting inhibitory control. Adaptation to errors in the timing of action execution was mediated by adjustments in the drift rate that progressively improved the precision of RTs with respect to the target response time. Inhibition errors, i.e., executed responses on trials requiring a stop, had a post-error slowing effect, mediated by an increase in the execution threshold that decayed over subsequent trials. These two mechanisms allowed for principled, context-specific adjustments in behavioral control to conflicting sources of task error, i.e., go timing and stop accuracy. Relative to the Uniform condition, subjects in the Early condition exhibited faster RTs at the expense of accuracy on probe Stop trials (see Fig. 3B). Subjects in the Early condition benefited from predictably short SSDs, making it easier to reactively cancel actions on Stop trials without sacrificing the the precision of action timing on Go trials. Over time, these subjects experience relatively few stop errors and thus focused on minimizing errors between their Go trial RT and the target response times (see Fig. 6B, left). In contrast, subjects in the Late condition slowed their RT on Go trials to accommodate the higher probability of a stop-cue late in the trial, introducing greater long term costs due to more frequent stop errors, and thus, delay responding on Go trials to improve inhibition accuracy. This principled adaptation of control is only possible due to the presence of dual learning mechanisms on the action control algorithm.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Dual Learning Mechanisms for Adaptive Control

In accumulation-to-bound models of decision-making, the drift-rate and boundary height parameters are functionally dissociated as representing the strength of evidence and response caution, respectively (Brown & Heathcote, 2008; Ratcliff, Smith, Brown, & McKoon, 2016) . However, this dissociation is only useful insofar as evidence is clearly defined and can be manipulated independently of the additional factors bearing on behavior (e.g. caution, expectation, attention). With the exception of perceptual decision-making tasks (Roitman & Shadlen, 2002), where evidence can be interpreted with respect to the strength of sensory information provided by the stimulus, it is often unclear what sources of information are actually encoded as evidence in the decision process and what are used to set the boundary height. In the following section we discuss evidence that, within the context of go/no-go control tasks, drift-rate and boundary height mechanisms are driven by corticostriatal and corticosubthalamic systems, respectively; and furthermore, that both participate in feedback-dependent learning but in the service of distinct behavioral goals.

Drift-Rate: Dopaminergic Modulation of Corticostriatal Pathways

At the computational level, the drift-rate parameter reflects the log-likelihood ratio of evidence for alternative hypotheses. In the context of the current task, the execution drift-rate can be interpreted as representing the relative evidence for go and no-go decisions encoded by the circuit-level competition between the direct and indirect pathways (Bahuguna, Aertsen, & Kumar, 2015). Indeed, studies combining behavioral modeling with single-unit recordings (Ding & Gold, 2010), optogenetics in animals (Yttri & Dudman, 2016), and neuroimaging in humans (van Maanen, Fontanesi, Hawkins, & Forstmann, 2016) have found reliable links between behaviorally derived estimates of drift-rate and activity in the striatum (Brody & Hanks, 2016). Crucially, the dynamics of competition between direct and indirect pathways is sensitive to dopaminergic signals that provide important feedback about the environmental consequences of recent actions to drive

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

behavior in the direction of the agent's current goal (Cox et al., 2015; Kravitz, Tye, & Kreitzer, 2012; Shan, Ge, Christie, & Balleine, 2014; Vicente, Galvão-Ferreira, Tecuapetla, & Costa, 2016). Feedback-dependent reweighting of corticostriatal connections has primarily been studied in the context of action-value learning; however, new evidence suggests a more nuanced role in tuning task-relevant movement parameters (Dudman & Krakauer, 2016; Rueda-Orozco & Robbe, 2015; Yttri & Dudman, 2016). Yttri and Duman (2016) demonstrated this by stimulating direct or indirect pathway neurons in the mouse striatum based on the velocity of a recently executed lever press and measuring the effects on future movements. Similar to the opponent effects of dopaminergic error signals that mediate action-value associations (Collins & Frank, 2014; Kravitz et al., 2012), they found that stimulation of the direct pathway following high-velocity presses further increased the velocity of future movements whereas stimulation of indirect pathway neurons decreased velocity. While the current study was not concerned with action velocity, per se, the adaptation of the drift-rate parameter to errors in action timing resembles a similar behavioral dynamic to that observed by Yttri and Dudman (2016). Indeed, a recent study by Soares et al., (2016) found that dopaminergic neurons in the mouse midbrain were not only necessary for accurate temporal perception, but that the perception of time could be systematically sped up or slowed down through optogenetic suppression and stimulation of these neurons. Future studies will be needed to confirm the proposed dependency of the drift-rate on striatum in which model-fits to behavior are performed in the presence of dopaminergic weighting at direct and indirect synapses.

Based on previous evidence that proactive control is mediated by the striatum, we have argued that the feedback-dependent modulation of execution drift-rate from to changes in the stop-cue prior is, at least in part, a reflection of the competition between the direct and indirect pathways, and that the observed adaptation in this parameter is a reflection of dopaminergic feedback about errors in action timing. In addition to the dopamine hypothesis, an alternative possibility

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

is that adaptation the drift-rate stems from top-down changes in the background excitability of the striatum, driven by diffuse inputs from premotor regions such as supplementary motor area (SMA) and pre-SMA (Forstmann et al., 2008; Murakami, Shteingart, Loewenstein, & Mainen, 2017; Murakami, Vicente, Costa, & Mainen, 2014; van Maanen et al., 2016). It remains unclear what functional differences may exist between premotor and dopaminergic representations of time or how they might differentially influence the encoding of action timing within the striatum. Integrating the behavioral and modeling techniques defined here with electrophysiological and optogenetic manipulations can better distinguish the nature of the training signal that modulates striatal activity during action control.

Boundary Height: Post-Error Slowing by the STN

The striatum has long been the focus of investigations into the neural basis of feedback-dependent learning; however, BG pathways have multiple targets of plasticity beyond the striatum. Recently, similar links have been identified between activity fluctuations in the STN and adaptive changes in behavior (Cavanagh et al., 2014; Herz et al., 2016; Justin Rossi et al., 2017; Wessel et al., 2016), raising new and interesting questions about the extent to which striatal and subthalamic learning signals independently influence behavior and how they might interact (Tewari, Jog, & Jog, 2016). Numerous studies have implicated the STN in setting the height of the decision threshold (Cavanagh et al., 2011; Frank et al., 2015; Herz et al., 2017, 2016; Zavala et al., 2016), controlled by diffuse excitatory inputs to the output nucleus of the BG and further suppressing motor thalamus to delay action execution (see Fig. 1A). Due to the monosynaptic connections between cortex and the STN that make up the hyperdirect pathway (Nambu, Tokuno, & Takada, 2002), unexpected sensory events (e.g., stop signals) can be quickly relayed through the STN to raise the decision threshold for ongoing action plans to prevent execution (Wessel & Aron, 2017). In addition to this rapid cortically-mediated form of adaptation, evidence suggests that strategic adjustments in decision threshold are achieved via activity-dependent plasticity in

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

the connections between STN and GPe (Baufreton & Bevan, 2008; Wei & Wang, 2016). In the current study, adaptive changes in the boundary height accounted for the observed post-error slowing in responses following failed stop trials, motivated by neuroimaging and electrophysiological evidence STN-mediated slowing of responses (Cavanagh et al., 2014; Frank et al., 2015; Herz et al., 2016). For simplicity, boundary adaption was restricted to being unidirectional - increasing after a stop-error and decaying back to, but never below, its original value. However, some evidence suggests that STN exerts bidirectional control over decision threshold, capable of promoting the adoption of both speed and accuracy policies (Herz et al., 2017). Thus, relating the adaptive threshold in the nested process model to recordings in the STN will likely require a more nuanced approach in order to generalize beyond the current task. Future studies will be needed to examine how each of these neural mechanisms are recruited to modify behavior, the relevant contexts and task dimensions they are sensitive to and the timescales they operate on.

Conclusion

Computational modeling of RT and accuracy in an adaptive stop-signal task revealed two independent learning mechanisms underlying feedback-dependent changes in control - one responsible for gradually tuning the execution drift-rate to timing errors on Go trials and another for increasing caution after a failed Stop trial by raising the execution boundary. The tuning of the drift-rate and boundary height parameters supports recent evidence of striatal- and STN-driven changes in goal-directed behavior following errors in action timing and inhibitory control respectively. While cognitive modeling approaches are unable to capture the complexity of neural information processing that underlies adaptive action control, they provide a rich description of the component operations and can thus be exploited for the purpose of more keenly parsing the functional role of different neural substrates. The current study shows how a straightforward hybridization of two cognitive modeling frameworks - accumulation-to-bound

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

dynamics with basic principles of reinforcement learning - provides confirmatory evidence for a dual-mechanism account of feedback dependent learning in inhibitory control.

Materials and Methods

Participants

Neurologically healthy adult participants (N=75, Mean age 22 years) were recruited from the Psychology Research Experiment System at Carnegie Mellon University. All procedures were approved by the local Institutional Review Board. All subjects were compensated for their participation through course credit toward fulfillment of their semester course requirements.

Adaptive Stop-Signal Task

All subjects completed a stop-signal task (N_{trials}=880) in which a vertically moving bar approached a white horizontal target line at the top of the screen (Fig. 2A). On 'Go' trials (N_{GO}=600) the subject was instructed to make a key press as soon as the bar crossed the target. The bar always intersected the target line at 520ms after trial onset. On each trial, the bar continued filling upward until a keypress was registered or until reaching the top of the screen, allowing a 680ms window for the subject to make a response. If no response was registered the subject received a penalty of (-100pts). On Go trials where a response was recorded before the 680 ms trial deadline, the subject received a score reflecting the precision of their response time relative to the target intersection, resulting in maximal points when RT=520ms. On Stop trials, the bar would stop and turn red prior intersecting the target line, prompting the subject to withhold their response. Successful and unsuccessful Stop trials yielded a reward of +200 points and penalty of -100 points, respectively. On the majority of Stop trials, the stop-signal delay (SSD) - the delay between trial onset and when the bar stopped – was sampled from a specific probability distribution (see Fig. 2B). We refer to these

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

trials as Context Stop trials ($N_{Context}$=200). Context SSD's in the Early and Late groups were sampled from Gaussian distributions with equal variance (σ=35ms), centered at μ_E=250ms and μ_L=350ms, respectively. Context SSDs in the Uniform group were sampled from a uniform distribution spanning a 10-520ms window. In Figure 2B, the sampled SSD times are plotted for a single subject in each Context – shown as dashes on a timeline ranging from 0-520ms. Finally, additional probe Stop trials (N_{Probe}=80) were included in which the bar stopped at 200, 250, 300, 350, or 400 ms after trial onset (16 trials per probe SSD), shown at the bottom of the Figure 2B timeline as red dashes.

Dependent Process Model

The nested process model (Fig. 1C; (Dunovan et al., 2015)) assumes that action-facilitating (i.e., direct) and action-suppressing (i.e., indirect) signals are integrated over time as a single execution process (θ_e), with a drift-rate that increases with the ratio of direct-to-indirect pathway activation. The linear drift and diffusion (φ_e) of the execution process is described by the stochastic differential equation in equation 1, accumulating with a mean rate of v_e (i.e., drift rate) and a standard deviation described by the dynamics of a white noise process (e.g., dW) with diffusion constant σ. The execution process is fully described by equation 2 in which the linear accumulation described by equation 1 is scaled by an urgency signal, modeled as a hyperbolic cosine function of time with gain γ.

$$d\varphi_e = \upsilon_e dt + \sigma dW \qquad (1)$$

$$\theta_e(t) = \varphi_e(t) \cdot cosh(\gamma \cdot t) \qquad (2)$$

A response is recorded if θ_e reaches the execution boundary (a) before the end of the trial window (680ms) and before the braking process reaches the lower (0) boundary (see below). In the event of a stop cue, the braking process (θ_b) is initiated at the current state of θ_e with a negative drift rate ($-v_b$). If θ_b reaches the

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

0 boundary before θ_e reaches the execution boundary then no response or RT is recorded from the model. The in θ_b over time is given by equation 3, expressing the same temporal dynamics of φ_e but with a negative drift rate ($-|v_b|$and no urgency signal in the absence of the dynamic bias signal. The dependency between θ_b and θ_e in the model is described by the conditional statement in equation 4, declaring that the initial state of θ_b (occurring at t = SSD) is equal to the state of θ_e(SSD).

$$d\theta_b = \upsilon_b dt + \sigma dW \qquad (3)$$

$$\theta_b(SSD) = \theta_e(SSD) \qquad (4)$$

In order to determine which of the model parameter(s) best accounted for the observed behavioral effects across Contexts, we first fit the model to the average data in the Uniform group, leaving all parameters free (see Table 1). Using the optimized Uniform parameter estimates to initialize the model, we then fit different versions of the model to data in the Early and Late groups allowing only one or two select parameters to vary between conditions. This form of model comparison provides a straightforward means of testing alternative hypotheses about the mechanism underlying Context-specific adaptation. The fitting procedure utilized a combination of global and local optimization techniques (Bogacz & Cohen, 2004; Dunovan et al., 2015). All fits were initialized from multiple starting values in steps to avoid biasing model selection to unfair advantages in the initial settings. Given a set of initial parameter values, all model parameters – Execution Drift-Rate (v_e), Braking Drift-Rate (v_b), Execution onset delay (tr), boundary height (a) and dynamic gain (γ) were optimized by minimizing a weighted cost function (see eq. 5 below) equal to the summed and squared error between an observed and simulated (denoted by ^) vector containing the following statistics: probability (P) of responding on Go trials (g), probability of stopping at each Probe SSD (d={200, 250, 300, 350, 400ms}), and

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

RT quantiles (q={.1, .2, …, .9}) on correct (RT^C) and error (RT^E) trials.

$$\chi^2_{static} = w_g \cdot (P_g - \hat{P}_g)^2 + \sum_{d}^{5} w_d \cdot (P_d - \hat{P}_d)^2 + \sum_{q}^{9} w_q^C \cdot (RT_q^C - \hat{RT}_q^C)^2 + \sum_{q}^{9} w_q^E \cdot (RT_q^E - \hat{RT}_q^E)^2 \quad (5)$$

The cost-function weights (*w*) were derived by first taking the variance of each summary measure included in the observed vector (across subjects), then dividing the mean variance by the full vector of variance scores. This approach represents the variability of each value in the vector as a ratio (Ratcliff & Tuerlinckx, 2002), where values closer to the mean are assigned a weight close to 1, and values associated with higher variability a weight <1, lower variability a weight >1 (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Dunovan et al., 2015). Weights applied to the RT quantiles were calculated by estimating the variance for each of the RT quantiles (Maritz & Jarrett, 1978) and then dividing the mean variance by that of each quantile. Stop accuracy weights were calculated by taking the variance in stop accuracy at each Probe SSD (across subjects) and then dividing the mean variance by that of each condition.

In order to obtain an estimate of fit reliability for each model we restarted the fitting procedure from 20 randomly sampled sets of initial parameter values. Each initial set was then optimized to average data in the Uniform condition using the Basinhopping algorithm (Wales & Doye, 1997) to find the region of global minimum followed by a Nelder-Mead simplex optimization (Nelder & Mead, 1964) for fine tuning globally optimized parameter values. The simplex-optimized parameter estimates were then held constant except for one or two designated context-dependent parameter(s) that were submitted to a second Simplex run in order to find the best fitting values in the Early and Late conditions.

Adaptive DPM

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

On correct Go trials, the drift-rate (v_t) was updated according to equation 6 to reflect the signed difference between model's response time on the current trial and the Target time (T^G = 520ms), increasing the drift-rate following "slow" responses (i.e., $RT_t > T^G$, Fig. 5A left) and decreasing the drift-rate following "fast" responses (i.e., $RT_t < T^G$, Fig. 5A middle). On failed stop trials, v_t was updated according to the same equation but with the error term reflecting the difference between RT_t and the trial response deadline (T^S=680), thus, slowing the drift-rate to reduce the probability of failed stops in the future.

$$v_{t+1} = v_t \cdot e^{\alpha_t (RT_t - T^{G/S})} \tag{6}$$

The boundary height (a_0) adapted to failed stop trials (Fig. 5A right) by increasing according to a delta function with height β_t and decaying exponentially on each subsequent trial (a_{terr}) until reaching its baseline value a_0 or until another stop error occurred (eq. 7).

$$a_{t_{err}} = a_0 + \beta_t e^{-t_{err}} \tag{7}$$

On all correct Go trials and the first failed stop trial, the timing errors were scaled by the same learning rate (α_0). An additional parameter was included to modulate the sensitivity (π) to stop errors over time, calculated according the power rule shown in equation 8. The trial-wise stop-error sensitivity was used to scale α_t (eq. 9) and β_t (eq. 10) learning rates on failed stop trials before calculating the drift (eq. 6) and boundary height (eq. 7) updates.

$$\pi_t = \begin{cases} \left(\frac{t}{10}\right)^p, & \text{if stop trial} \\ 1, & \text{if go trial} \end{cases} \tag{8}$$

$$\alpha_t = \pi_t \alpha_0 \tag{9}$$

$$\beta_t = \pi_t \beta_0 \tag{10}$$

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

To obtain estimates for the learning rate parameters for adaptation in the drift-rate (α), boundary-height (β), and sensitivity to stop errors (*p*) this adaptive form of the model was re-fit to the RT and stop accuracy data in the Uniform condition, holding all previously estimated parameters constant. Because standard parameter optimization for accumulator models requires information about the variance of response-times across trials, these approaches are poorly suited for investigating how decision parameters respond to error on a trial-wise basis. To overcome this issue, the cost function was modified to identify the values for α, β, and *p* that minimized the difference between the average observed and model-predicted RT and stop accuracy over a moving window of 30 trials (eq. 11). By averaging the behavior in 30-trial bins (30 bins total), this ensured that multiple Stop trials were included in each bin while still allowing relatively high-frequency behavioral changes to be expressed in the cost function.

$$\chi^2_{adapt} = \sum_{i}^{30}(\mu_{acc,i} - \hat{\mu}_{acc,i})^2 + \sum_{i}^{30}(\mu_{rt,i} - \hat{\mu}_{rt,i})^2 \tag{11}$$

These fits were performed by iteratively simulating the same trial sequence as observed for each individual subject, and fitting the average simulated subject to the average observed subject. This ensures that direct comparisons can be made between the trajectory of learning in the model and actual behavior.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Tables

Table 1. Average Uniform Parameters and Static Model Fit Statistics

	a	v_E	v_B	tr	γ	χ^2	**AIC**	**BIC**
Mean	.629	1.146	-1.012	.0746	1.149	.013	-173.19	-178.61
95% CI	.022	.065	.046	.009	.27	8e-4	1.06	1.06

Table 2. Static Fit Statistics for Early and Late Contexts

Context Parameter	χ^2 Best (Mean, 95%CI)	**AIC** Best (Mean, 95%CI)	**BIC** Best (Mean, 95%CI)
Execution Drift (v_E)	.020 (.026, .002)	-369.26 (-357.49, 4.06)	-365.51 (-353.75, 4.06)
Boundary Height (a)	.036 (.039, .001)	-341.02 (-337.66, 1.04)	-337.28 (-333.91, 1.04)
Braking Drift (v_B)	.036 (.038, .001)	-341.74 (-338.55, 1.02)	-337.10 (-334.81, 1.02)
Urgency Gain (γ)	.028 (.030, .001)	-353.43 (-350.13, 1.13)	-349.69 (-346.38, 1.13)
*a & v_E	.014 (.025, .005)	-381.52 (-357.84, 7.52)	-374.04 (-350.35, 7.52)
v_B & v_E	.021 (.026, .002)	-364.11 (-353.66, 4.11)	-356.63 (-346.18, 4.11)
γ & v_E	.018 (.025, .003)	-371.02 (-355.77, 4.57)	-363.54 (-348.29, 4.57)

* Denotes the best-fitting model

Table 3. Adaptive Model Parameters

α_0	β_0	p	**AIC**	**BIC**
.24	.031	.002	-173.19	-178.61

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Figures

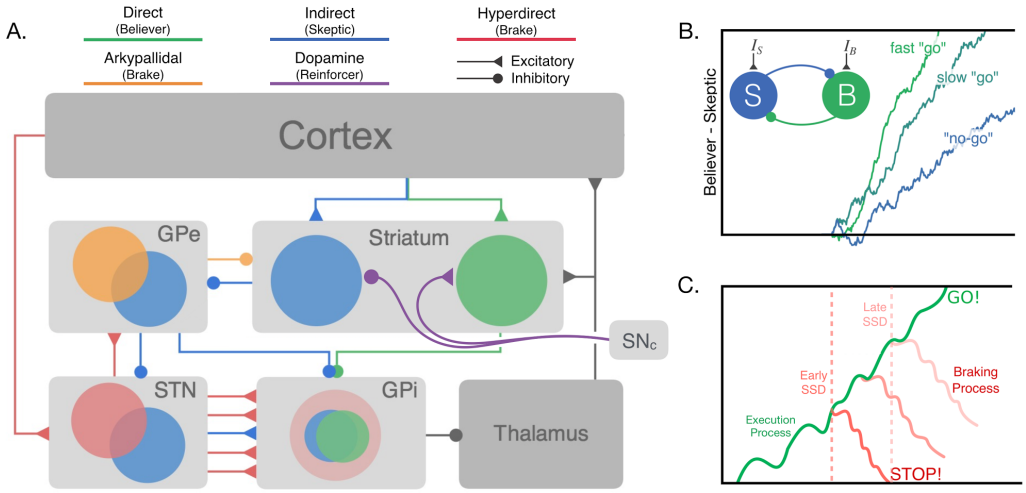

Figure 1. Basal Ganglia Circuitry and Dependent Process Model. (A) Major cortico-basal ganglia (BG) pathways and hypothesized control functions. Projections from the substantia nigra (SNc) deliver dopaminergic feedback signals to the striatum that have opposing effects on the excitability of direct and indirect pathway neurons. **(B)** Competition between direct and indirect pathways represented by mutually inhibiting "believer" (green; direct pathway) and "skeptic" (blue; indirect pathway) populations. Circuit-level dynamics of this competition modulate the rate of evidence accumulation leading up to action execution, leading to faster actions when competition is dominated by "believer". **(C)** The Dependent Process Model (DPM) assumes that the state of an accumulating execution process at the time a stop cue is registered determines initial state of the braking process, making it more difficult to cancel actions closer to the execution boundary. Panels **A** and **B** have been adapted with permission from Dunovan and Verstynen (2016)**.** Panel **C** has been adapted with permission from Dunovan et al., (2015).

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

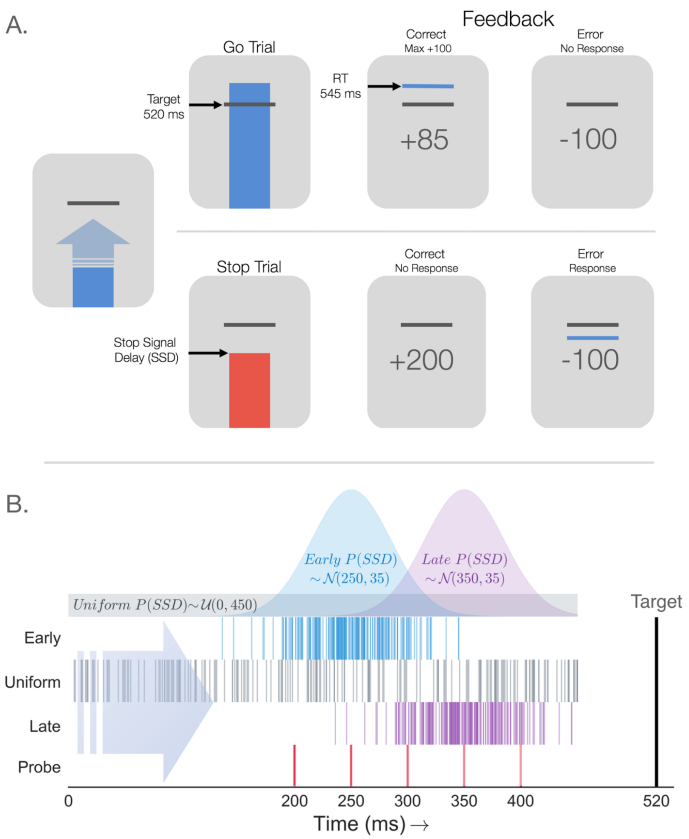

Figure 2. Adaptive Stop-Signal Task and Contextual SSD Statistics. (A) Anticipatory stop-signal task. On Go trials (upper) subjects were instructed to press a key when the ascending bar crossed a target line, always occurring on 520 ms after trial onset. Feedback was given informing the subject if their response was earlier or later than the Go Target (max +100 points). On Stop trials (lower), the bar stopped and turned red prior to reaching the Target line. If no response was made (correct), the subject received a bonus of +200 points. Failure to inhibit the keypress resulted in a -100 point penalty. **(B)** Stop-Signal statistics across Contexts. Distributions show the sampling distributions for SSDs on Context trials in the Early (blue), Uniform (gray), and Late (purple) groups. Early and Late SSDs were Normally distributed (parameters stated as in-figure text N(μ, σ)). Below the distributions, each row of tick-marks shows the Context SSDs for a single example subject each group. Bottom row of red tick-marks shows the five Probe SSDs included for all subjects regardless of context.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

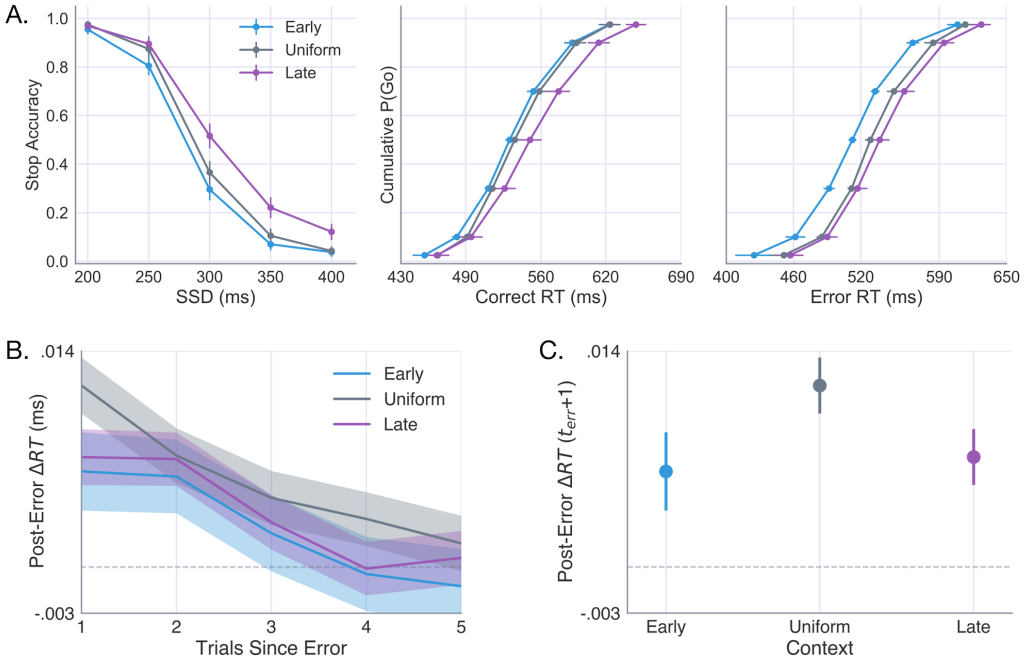

Figure 3. Effects of Context on Stop Accuracy and Response Times. (A) Subject-averaged stop-accuracy (left) and cumulative RT distributions for correct (Go trials; middle) and error (Stop trials; right) responses in the Early (blue), Uniform (gray), and Late (purple) Contexts. **(B)** Post-error slowing following failed Stop trials in each context and subsequent decay over five trials **(C)** The post-error slowing observed immediately after a failed stop (t_{err}+1) in each Context (e.g., first data point in panel **B**). Error bars and shaded area reflect the 95% confidence interval (CI) calculated across subjects.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

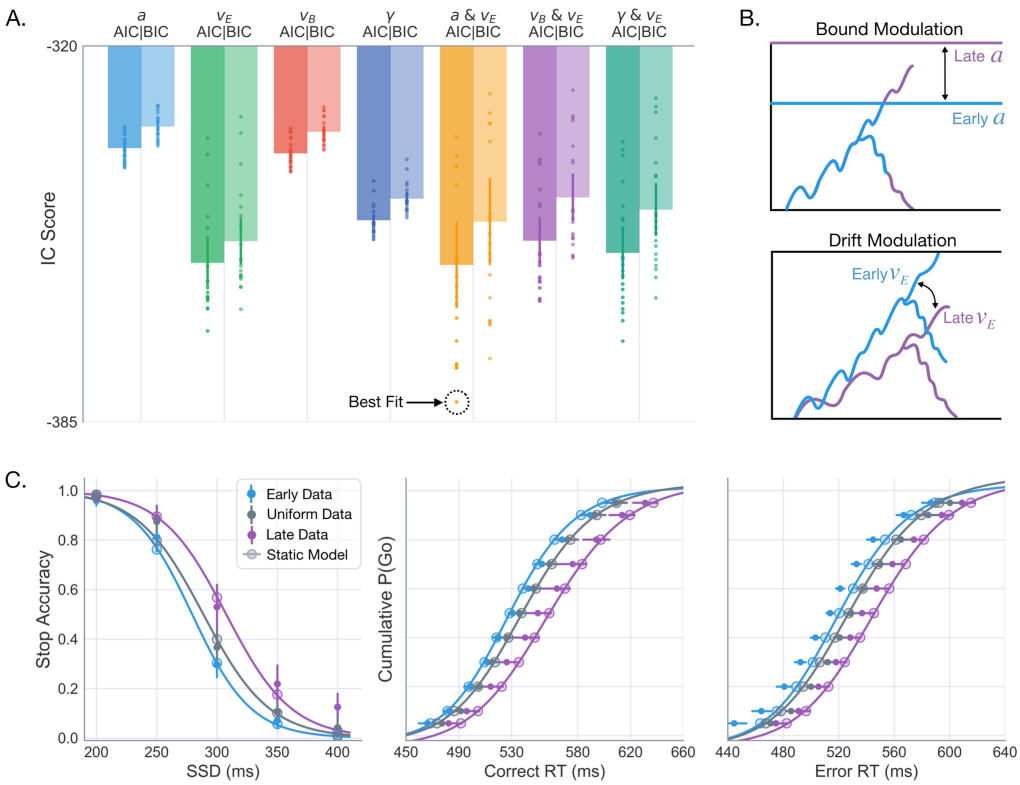

Figure 4. Model Comparison and Best-Fit Predictions Across Context. (A) AIC (dark) and BIC (light) scores for all single-parameter models, allowing either execution drift-rate (v_e; green), boundary height (a; cyan), braking drift-rate (v_b; red), or urgency (γ; blue) to vary across Context conditions. Three dual-parameter models were also included to test for possible benefits of allowing v_e (best-fitting single parameter model) to vary along with either a (yellow), v_b (purple), or γ (teal). Each dot shows the information criterion (IC) value for one of twenty fits performed with each model. Error bars show the 95% CI. **(B)** Qualitative effects of context on boundary height (top) and drift-rate parameter estimates (bottom) in the Early and Late Contexts. **(C)** Model predicted data (lines and larger transparent circles) simulated with best-fit parameters from the a & v_e (corresponding to dotted circle in **A**) overlaid on the average empirical data for Early (blue), Uniform (gray), and Late (purple) context conditions.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

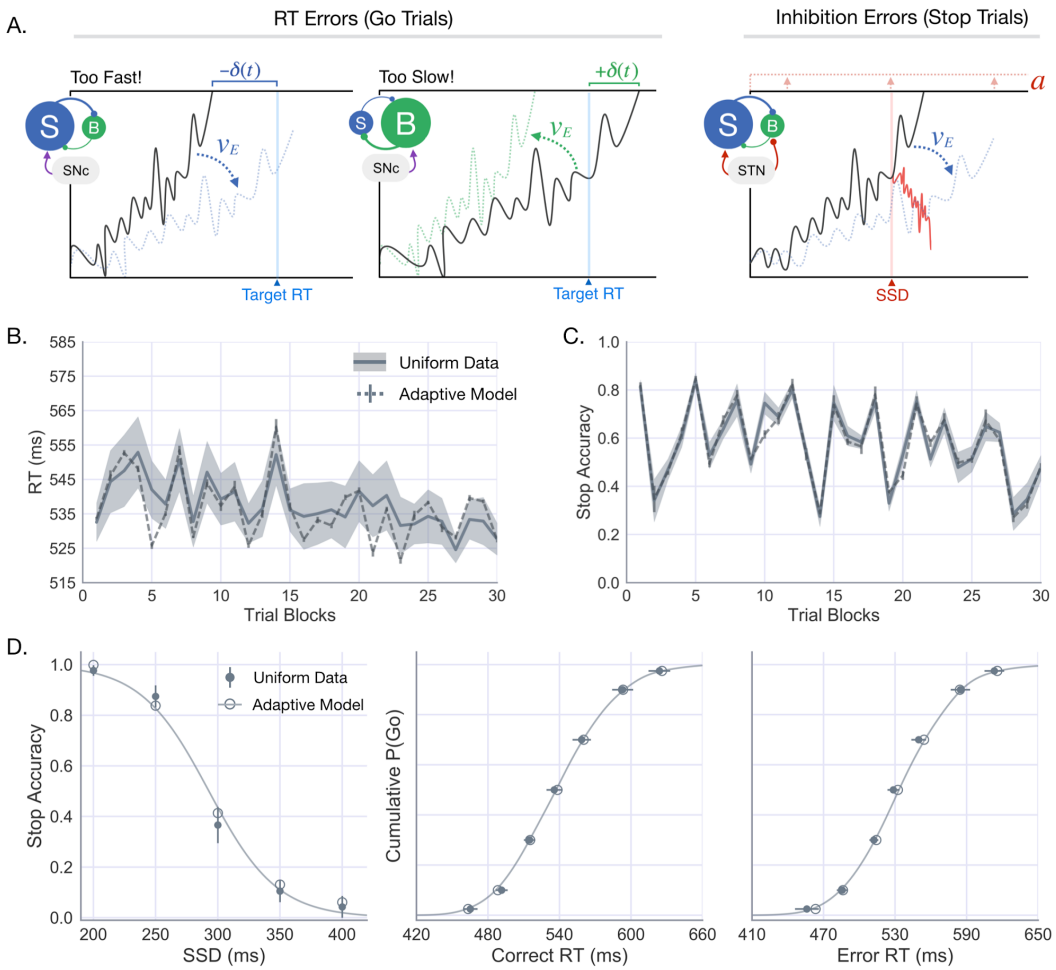

Figure 5. Adaptive DPM and Predicted Learning Trajectory in Uniform Condition. (A) Schematic showing how the execution drift-rate is modulated following timing errors on Go trials (left) and how the boundary height is modulated following failed inhibitions on Stop trials. **(B)** Subject-averaged timeseries (dark line) and 95% CI (transparent gray) showing the RT on Go trials (left) accuracy percentage on Stop trials (right). Both timeseries are 30 points in total, each calculated as by taking the average RT/stop-accuracy over successive~30-trial windows). The corresponding model predictions are overlaid (light gray line), averaged over simulations to each individual subject's data. **(C)** Average empirical stop-accuracy and RT statistics in the Uniform condition, (same as shown in Figure 3) with the predictions generated from simulations with the adaptive DPM.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

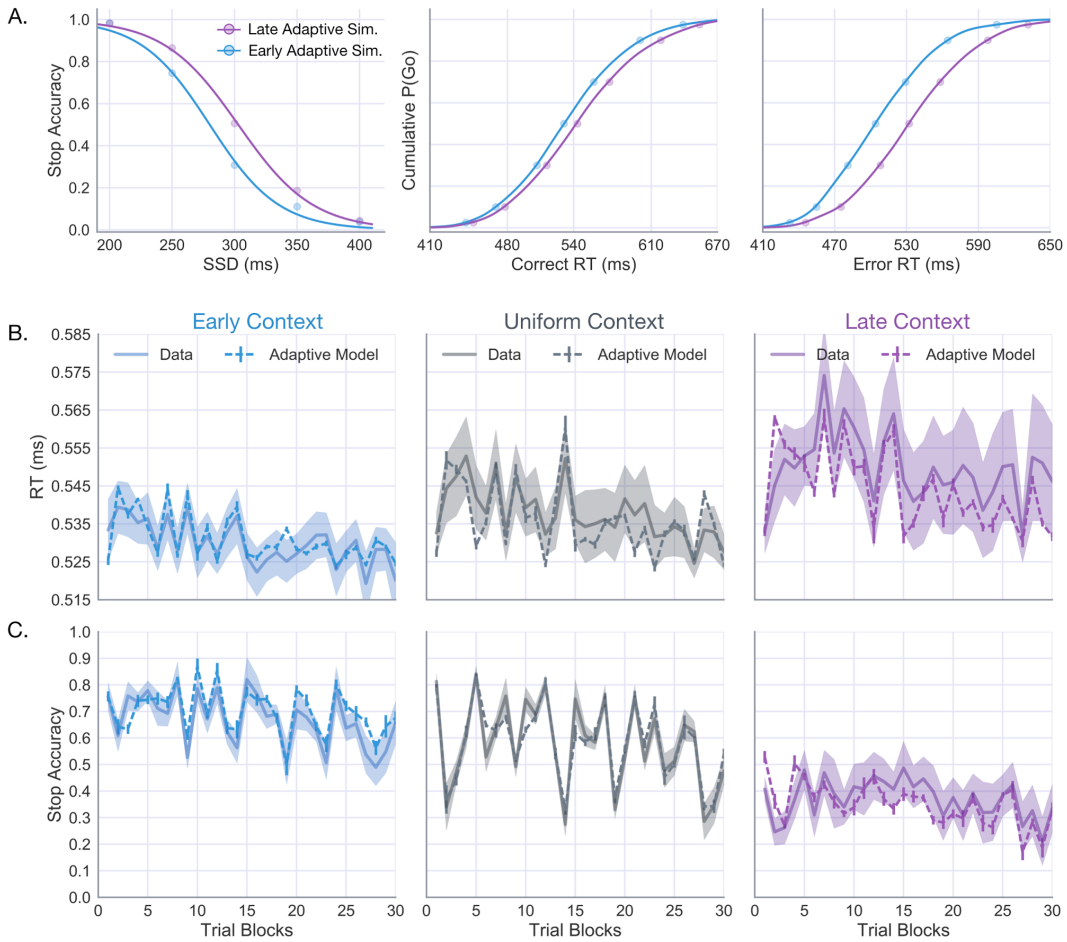

Figure 6. Adaptive DPM Modulates Behavior to Context-Specific Control Demands. (A) Average stop-accuracy curves (left) and correct (middle) and error (right) RT distributions predicted by adaptive model simulations in the Early (blue) and Late (purple) Contexts (initialized with the optimal parameters of the Uniform Context). **(B)** Empirical timeseries of Go RT's with model predictions overlaid (calculated using the same method described for Figure 5B-C for Early (left), Uniform (middle, same as in Figure 5C), and Late (right) Contexts. **(C)** Empirical and model predicted timeseries of stop-accuracy for the same conditions as in **B**.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

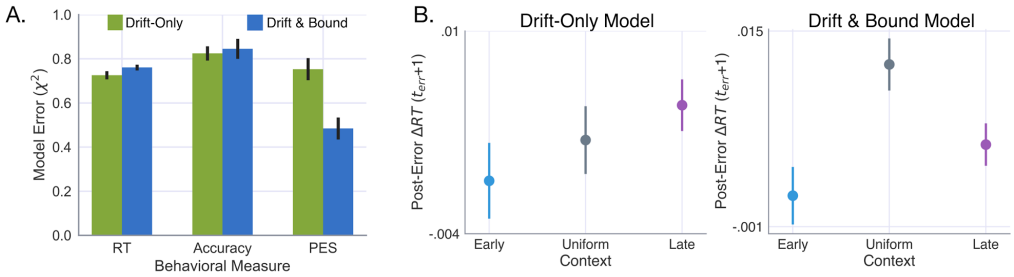

Figure 7. Predictive utility of including boundary adaptation compared to drift-only model (**A**) Relative error of simulated compared to observed RT, accuracy, and post-error slowing measures based on twenty simulated data sets for the drift-only and drift & bound adaptive models. Post-error slowing in each context condition as predicted by the **(B)** drift-only and **(C)** drift & bound models. Error bars reflect the 95% CI around the mean.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

References

Bahuguna, J., Aertsen, A., & Kumar, A. (2015). Existence and control of Go/No-Go decision transition threshold in the striatum. *PLoS Computational Biology*, *11*(4), e1004233.

Baufreton, J., & Bevan, M. D. (2008). D2-like dopamine receptor-mediated modulation of activity-dependent plasticity at GABAergic synapses in the subthalamic nucleus. *The Journal of Physiology*, *586*(8), 2121–2142.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765.

Bogacz, R., & Cohen, J. D. (2004). Parameterization of connectionist models. *Behavior Research Methods, Instruments, & Computers: A Journal of the Psychonomic Society, Inc*, *36*(4), 732–741.

Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation*, *19*(2), 442–477.

Bogacz, R., & Larsen, T. (2011). Integration of reinforcement learning and optimal decision-making theories of the basal ganglia. *Neural Computation*, *23*(4), 817–851.

Brittain, J.-S., Watkins, K. E., Joundi, R. A., Ray, N. J., Holland, P., Green, A. L., … Jenkinson, N. (2012). A role for the subthalamic nucleus in response inhibition during conflict. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *32*(39), 13396–13401.

Brody, C. D., & Hanks, T. D. (2016). Neural underpinnings of the evidence accumulator. *Current Opinion in Neurobiology*, *37*, 149–157.

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

time: linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178.

Cavanagh, J. F., Sanguinetti, J. L., Allen, J. J. B., Sherman, S. J., & Frank, M. J. (2014). The subthalamic nucleus contributes to post-error slowing. *Journal of Cognitive Neuroscience*, *26*(11), 2637–2644.

Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., & Frank, M. J. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience*, *14*(11), 1462–1467.

Collins, A. G. E., & Frank, M. J. (2014). Opponent actor learning (OpAL): modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive. *Psychological Review*, *121*(3), 337–366.

Cox, S. M. L., Frank, M. J., Larcher, K., Fellows, L. K., Clark, C. A., Leyton, M., & Dagher, A. (2015). Striatal D1 and D2 signaling differentially predict learning from positive and negative outcomes. *NeuroImage*, *109*, 95–101.

Ding, L., & Gold, J. I. (2010). Caudate encodes multiple computations for perceptual decisions. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(47), 15747–15759.

Ding, L., & Gold, J. I. (2012). Separate, Causal Roles of the Caudate in Saccadic Choice and Execution in a Perceptual Decision Task. *Neuron*, *75*(5), 865–874.

Dudman, J. T., & Krakauer, J. W. (2016). The basal ganglia: from motor commands to the control of vigor. *Current Opinion in Neurobiology*, *37*, 158–166.

Dunovan, K., Lynch, B., Molesworth, T., & Verstynen, T. (2015). Competing basal-ganglia pathways determine the difference between stopping and deciding not to go. *eLife*, *4*, e08723.

Dunovan, K., & Verstynen, T. (2016). Believer-Skeptic Meets Actor-Critic: Rethinking the Role of Basal Ganglia Pathways during Decision-Making and Reinforcement

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Learning. *Frontiers in Neuroscience*, *10*, 106.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(45), 17538–17542.

Frank, M. J., & Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: Computational analysis. *Cerebral Cortex* , *22*(3), 509–526.

Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. (2015). fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(2), 485–494.

Gurney, K. N., Humphries, M. D., & Redgrave, P. (2015). A New Framework for Cortico-Striatal Plasticity: Behavioural Theory Meets In Vitro Data at the Reinforcement-Action Interface. *PLoS Biology*, *13*(1), e1002034.

Heitz, R. P., & Schall, J. D. (2013). Neural chronometry and coherency across speed–accuracy demands reveal lack of homomorphism between computational and neural mechanisms of evidence …. *Philosophical Transactions of*. Retrieved from http://rstb.royalsocietypublishing.org/content/368/1628/20130071.short

Herz, D. M., Tan, H., Brittain, J.-S., Fischer, P., Cheeran, B., Green, A. L., … Brown, P. (2017). Distinct mechanisms mediate speed-accuracy adjustments in cortico-subthalamic networks. *eLife*, *6*. https://doi.org/10.7554/eLife.21481

Herz, D. M., Zavala, B. A., Bogacz, R., & Brown, P. (2016). Neural Correlates of Decision Thresholds in the Human Subthalamic Nucleus. *Current Biology: CB*, *26*(7), 916–920.

Justin Rossi, P., Peden, C., Castellanos, O., Foote, K. D., Gunduz, A., & Okun, M. S.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

(2017). The human subthalamic nucleus and globus pallidus internus differentially encode reward during action control. *Human Brain Mapping*, *38*(4), 1952–1964.

Keuken, M. C., Van Maanen, L., Bogacz, R., Schäfer, A., Neumann, J., Turner, R., & Forstmann, B. U. (2015). The subthalamic nucleus during decision-making with multiple alternatives. *Human Brain Mapping*, *36*(September), 4041–4052.

Kravitz, A. V., Tye, L. D., & Kreitzer, A. C. (2012). Distinct Roles for direct and indirect pathway striatal neurons in reinforcement. *Nature Neuroscience*, *15*(6), 816–818.

Majid, D. S. A., Cai, W., Corey-Bloom, J., & Aron, A. R. (2013). Proactive selective response suppression is implemented via the basal ganglia. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *33*(33), 13259–13269.

Mallet, N., Schmidt, R., Leventhal, D., Chen, F., Amer, N., Boraud, T., … Boraud, T. (2016). Arkypallidal Cells Send a Stop Signal to Striatum Report Arkypallidal Cells Send a Stop Signal to Striatum. *Neuron*, *89*(2), 308–316.

Maritz, J. S., & Jarrett, R. G. (1978). the Variance of the A Note on Estimating Sample Median, *73*(361), 194–196.

Murakami, M., Shteingart, H., Loewenstein, Y., & Mainen, Z. F. (2017). Distinct Sources of Deterministic and Stochastic Components of Action Timing Decisions in Rodent Frontal Cortex. *Neuron*, *94*(4), 908–919.e7.

Murakami, M., Vicente, M. I., Costa, G. M., & Mainen, Z. F. (2014). Neural antecedents of self-initiated actions in secondary motor cortex. *Nature Neuroscience*, *17*(11), 1574–1582.

Nambu, A., Tokuno, H., & Takada, M. (2002). Functional significance of the cortico--subthalamo--pallidal “hyperdirect”pathway. *Neuroscience Research*, *43*(2), 111–117.

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Nelder, B. J. a., & Mead, R. (1964). A simplex method for function minimization.

Pedersen, M. L., Frank, M. J., & Biele, G. (2016). The drift diffusion model as the choice rule in reinforcement learning. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-016-1199-y

Purcell, B. A., & Palmeri, T. J. (2016). Relating accumulator model parameters and neural dynamics. *Journal of Mathematical Psychology*. https://doi.org/10.1016/j.jmp.2016.07.001

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion Decision Model: Current Issues and History. *Trends in Cognitive Sciences*, *20*(4), 260–281.

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, *9*(3), 438–481.

Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *22*(21), 9475–9489.

Rueda-Orozco, P. E., & Robbe, D. (2015). The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nature Neuroscience*, *18*(3), 453–460.

Schall, J. D., Palmeri, T. J., & Logan, G. D. (2017). Models of inhibitory control. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *372*(1718). https://doi.org/10.1098/rstb.2016.0193

Shan, Q., Ge, M., Christie, M. J., & Balleine, B. W. (2014). The acquisition of goal-directed actions generates opposing plasticity in direct and indirect pathways in dorsomedial striatum. *The Journal of Neuroscience: The Official Journal of the*

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Society for Neuroscience, *34*(28), 9196–9201.

Shenoy, P., & Yu, A. J. (2011). Rational decision-making in inhibitory control. *Frontiers in Human Neuroscience*, *5*(May), 48.

Smith, Y., Bevan, M. D., Shink, E., & Bolam, J. P. (1998). Microcircuitry of the direct and indirect pathways of the basal ganglia. *Neuroscience*, *86*(2), 353–387.

Soares, S., Atallah, B. V., & Paton, J. J. (2016). Midbrain dopamine neurons control judgment of time. *Science*, *354*(6317), 1273–1277.

Sutton, R. S., & Barto, A. G. (1998). *Introduction to Reinforcement Learning*. Camgridge: MIT Press.

Tewari, A., Jog, R., & Jog, M. S. (2016). The Striatum and Subthalamic Nucleus as Independent and Collaborative Structures in Motor Control. *Frontiers in Systems Neuroscience*, *10*. https://doi.org/10.3389/fnsys.2016.00017

Turner, B. M., van Maanen, L., & Forstmann, B. U. (2015). Informing cognitive abstractions through neuroimaging: the neural drift diffusion model. *Psychological Review*, *122*(2), 312–336.

van Maanen, L., Fontanesi, L., Hawkins, G. E., & Forstmann, B. U. (2016). Striatal activation reflects urgency in perceptual decision making. *NeuroImage*. https://doi.org/10.1016/j.neuroimage.2016.06.045

Verbruggen, F., & Logan, G. D. (2009). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience and Biobehavioral Reviews*, *33*(5), 647–661.

Vicente, A. M., Galvão-Ferreira, P., Tecuapetla, F., & Costa, R. M. (2016). Direct and indirect dorsolateral striatum pathways reinforce different action strategies. *Current Biology: CB*, *26*(7), R267–R269.

Wales, D., & Doye, J. P. K. (1997). Global Optimization by Basin-Hopping and the

bioRxiv preprint doi: https://doi.org/10.1101/153676; this version posted June 22, 2017. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

Lowest Energy Structures of Lennard-Jones Clusters Containing up to 110 Atoms. *The Journal of Physical Chemistry. A*, *101*(97), 5111–5116.

Wei, W., & Wang, X.-J. (2016). Inhibitory Control in the Cortico-Basal Ganglia-Thalamocortical Loop: Complex Regulation and Interplay with Memory and Decision Processes. *Neuron*, *92*(5), 1093–1105.

Wessel, J. R., & Aron, A. R. (2017). On the Globality of Motor Suppression: Unexpected Events and Their Influence on Behavior and Cognition. *Neuron*, *93*(2), 259–280.

Wessel, J. R., Jenkinson, N., Brittain, J.-S., Voets, S. H. E. M., Aziz, T. Z., & Aron, A. R. (2016). Surprise disrupts cognition via a fronto-basal ganglia suppressive mechanism. *Nature Communications*, *7*, 11195.

Yttri, E. A., & Dudman, J. T. (2016). Opponent and bidirectional control of movement velocity in the basal ganglia. *Nature*, *533*(7603), 402–406.

Zavala, B., Tan, H., Little, S., Ashkan, K., Green, A. L., Aziz, T., … Brown, P. (2016). Decisions Made with Less Evidence Involve Higher Levels of Corticosubthalamic Nucleus Theta Band Synchrony. *Journal of Cognitive Neuroscience*, *28*(6), 811–825.