

The Cost of Adjusting Cognitive Control: A Dynamical Systems Approach

Ivan Grahek^{*,1} (ivan_grahek@brown.edu)
Xiamin Leng^{*,1} (xiamin_leng@brown.edu)
Sebastian Musslick^{1,2} (sebastian_musslick@brown.edu)
Amitai Shenhav¹ (amitai_shenhav@brown.edu)

¹Department of Cognitive, Linguistic, and Psychological Sciences
Carney Institute for Brain Science
Brown University
Providence, RI, USA

²Institute of Cognitive Science
Osnabrück University
Osnabrück, Germany

*These authors contributed equally to this work



Abstract:

Flexible behavior requires adjustments of cognitive control. While the costs of switching *between* different tasks are well studied, much less is known about the costs of adjusting the level of control allocation *within* a given task. Here, we develop a model of cognitive control dynamics that assumes adjustment costs arise due to the time needed to change control signals, and that these costs constrain one's selection of optimal control signals. We empirically test two predictions of the model. We show that control adjustment costs depend on the time allotted to alter control levels, and that optimal control signals are modulated by the expected costs of adjusting levels of cognitive control.

Keywords: cognitive control; control costs; drift diffusion model; task-set inertia; switch costs

Introduction

Goal-directed behavior requires flexible adjustments of stimulus-response mappings (task sets) and the cognitive control strategies that guide information processing. One of the hallmarks of cognitive control are the costs associated with the adjustment of information processing (Alport et al., 1994; Monsell, 2003). So far, these costs have been studied in situations where people switch *between* task sets. It remains unclear whether the costs of adjusting cognitive control generalize to cases where people adjust information processing *within* a given task. Here, we seek to develop and empirically test a computational model in which adjustment costs arise from dynamical changes in levels of control allocation. We confirm predictions of our model in experiments requiring switches between performance goals within the same task. We show that adjustment costs scale with the time available to adjust control (Experiment 1), and that people weigh such costs prospectively (based on expected frequency of goal switching) when setting levels of control (Experiment 2).

Dynamical Model of Adjustment Costs

In this study, we consider adjustments in two types of control signals: processing efficiency (drift rate: v) and response caution (threshold: a). These dimensions form a 2D space (Fig. 1A-left) in which optimal control signals (v^*, a^*) can be identified (Bogacz et al., 2006). Optimal control signals (Eq. 1 & 2) maximize reward rate (RR), by balancing the benefits and costs of control. Benefits are determined by the probability of correct and incorrect responses (ER), weights on these outcomes (w_1, w_2), and the time needed to make a response (decision time: DT , and non-decision time: t). Intensity cost depends on the drift rate (cf. Leng et al., 2021).

$$RR = \frac{w_1 \times (1 - ER(v, a)) - w_2 \times ER(v, a)}{DT(v, a) + t} - v^2 - e^{\sqrt{c_v(v-v_0)^2 + c_a(a-a_0)^2}} \quad (1)$$

$$v^*, a^* = \operatorname{argmax}_{v, a} [RR | v_0, a_0, w_1, w_2, t] \quad (2)$$

We recently showed that frequent changes in levels of control allocation (e.g., switching between the goal of being fast vs. being accurate) carry a performance cost (Grahek et al., 2022). Here we provide a mechanistic account of this cost (Eq. 3 & 4, Fig. 1A-middle) by casting control dynamics as gradual (τ) and noisy (c^2) adjustments from the current state of the control system (v_0, a_0), toward the optimal state (v^*, a^*):

$$dv = -\tau \times (v - v^*)dt + N(0, c^2 dt) \quad (3)$$

$$da = -\tau \times (a - a^*)dt + N(0, c^2 dt) \quad (4)$$

There are two central assumptions in the model. First, changing control levels from the current to desired state takes time. Second, the time required to adjust control incurs a cost, which serves to regularize changes in optimal control levels (exponential term in Eq. 1). Thus, the cognitive system optimizes control allocation by maximizing reward rate, while minimizing adjustment costs. At task onset, the current control state is implemented (Fig. 1A-right; Navarro & Fuss, 2009), irrespective of whether it amounts to the optimal state.

The model makes at least two sets of testable predictions. First, if suboptimal control states are caused by insufficient time to transition between control states, control levels will be closer to optimality as the time allowed for adjustment increases (Experiment 1; Fig. 1B-left). Second, adjustment costs should increase with more frequent (within-task) goal switches. Thus, optimal control signals should be prospectively adjusted based on the expected (cued) frequency of goal switches (Experiment 2; Fig. 1B-right).

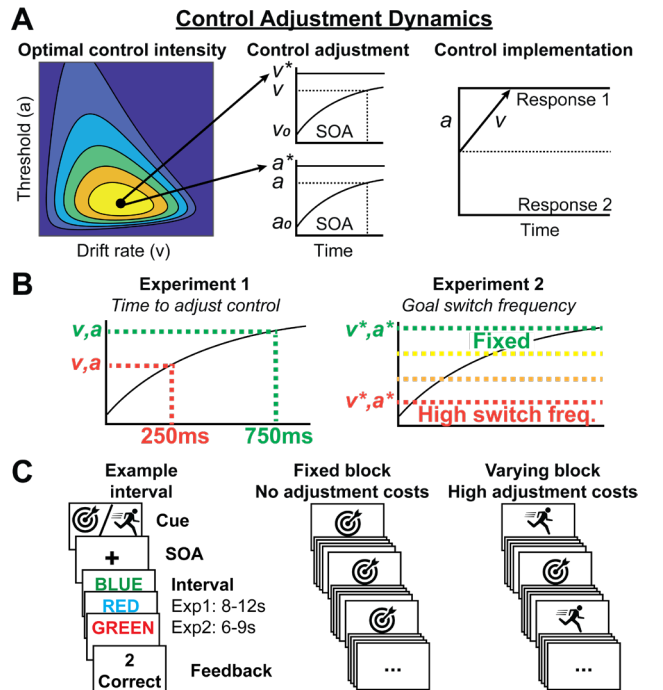


Figure 1. Computational model (A), predictions tested in two experiments (B), and task design (C).

Results

To manipulate adjustment costs, we had subjects perform the Stroop task as their performance goals varied, thus requiring them to adjust levels of control. The Stroop was performed over fixed time intervals during which subjects could complete as many trials as they wished (Fig. 1C-left). Before each interval, a cue instructed them to be as *accurate* (Accuracy goal; requires higher v and a), or as *fast* as possible (Speed goal; requires lower v and a). Across blocks, goals were either fixed (requiring no control adjustment), or varying (requiring frequent adjustments; Fig. 1C). To measure dynamic adjustments in control, we compared control signals (drift rate and threshold) between the two performance goals in fixed vs. varying blocks. Drift rates and thresholds were estimated with a hierarchical Bayesian drift diffusion model (Wiecki et al., 2013).

Time to Adjust Control

Our model assumes that control signals are adjusted incrementally from the current to the desired state. We predict that shorter adjustment windows result in suboptimal control signals because the desired state is not reached before task onset (Fig. 2A-left). To simulate this prediction, we first computed the optimal control states for the two performance goals, obtained by performing inverse reward-rate optimization on the group-level data to identify parameters of Eq. 1 (cf. Leng et al., 2021). Adjustments of control states involve switching between these optimal states. We could then simulate levels of control obtained when control signals transition from one state to another was terminated 250 or 750 timesteps after cue onset. We held constant adjustment cost weights for drift and threshold ($C_v = 0.1$, $C_a = 1$), as well as the dynamics parameters ($\tau = 0.6$, $dt = 0.005$; $c^2 = 0$).

To test this prediction, participants (N=50) performed the interval task described above (Fig. 1C-left) with varying time between cue and interval onset (SOA=250 vs. 750ms). Overall, drift rates ($b=0.15$; 95% CrI [0.06, 0.24]; $p_{b<0}<0.01$) and thresholds ($b=0.14$; 95% CrI [0.08, 0.19]; $p_{b<0}<0.01$) were higher in the Accuracy relative to Speed condition. The difference in thresholds between these goals was larger in fixed relative to varying blocks ($b=-0.07$; 95% CrI [-0.1, -0.03]; $p_{b>0}<0.01$), evidencing adjustment costs. Critically, this difference was lower when participants had more time to adjust control ($b=-0.78$; 95% CrI [-0.15, -0.01]; $p_{b>0}<0.01$; Figure B-left). For drift rates, we found no interaction between performance goal and block type, nor a 3-way interaction with SOA ($ps>0.3$). These results confirm our model's prediction that suboptimal adjustments can arise from insufficient time to transition between current and optimal levels of control (Fig 2A-left).

Expectation of Goal Switch Frequency

Our model postulates that control adjustment is regularized by prospective costs when optimizing control signals (v^* , a^*). To simulate this prediction, we implemented the same inverse optimization procedure and dynamics parameters as in Experiment 1. We then parametrically modulated adjustment cost weights on drift and threshold to simulate expected adjustment costs (highest when expecting to switch performance goal every 2 intervals; lowest when expecting no switches). These simulations assumed a relatively long SOA (1250ms). Our model predicts that, despite having sufficient *time* to adjust control, people should evaluate more frequent switches as incurring higher adjustment costs, and make smaller changes in levels of control signals on those blocks (Fig 2A-right).

In Experiment 2 (N=55), participants were explicitly instructed how often they would switch performance goals in a block (fixed vs. goal switches every 8, 4, or 2 intervals). As in Experiment 1, we found higher drift rates ($b=0.25$; 95% CrI [0.13, 0.37]; $p_{b<0}<0.01$) and thresholds ($b=0.13$; 95% CrI [0.09, 0.17]; $p_{b<0}<0.01$) in Accuracy vs. Speed intervals. Critically, the Accuracy-Speed difference decreased as the number of switches within a block increased (Figure 2B-right), for both drift rates and thresholds (interaction with the linear effect of block type: drift rates: $b=-0.20$; 95% CrI [-0.43, 0.03]; $p_{b>0}<0.05$; thresholds: $b=-0.12$; 95% CrI [-0.20, -0.04]; $p_{b<0}<0.01$). This result suggests that control optimization weighs expected control adjustment costs, as predicted by our model (Fig 2A-right).

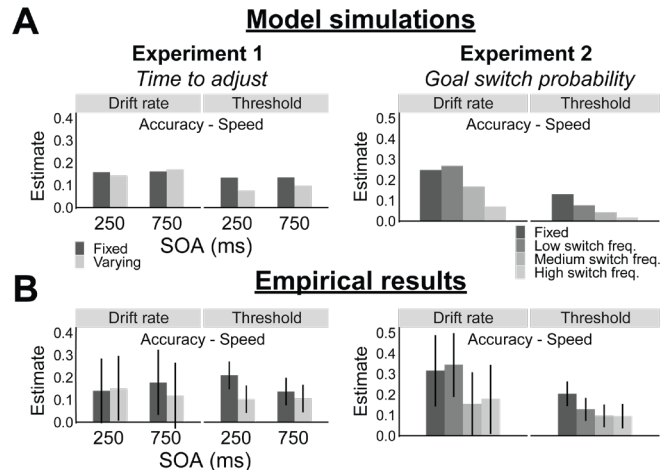


Figure 2. Model simulations (A) and parameter estimates with 95% credible intervals (B).

Conclusions

Both computational analyses and behavioral experiments suggest that within-task adjustments of cognitive control are subject to costs. Adjustment costs are proposed to arise due to the time needed to alter control signals, and control adjustments are regularized prospectively when allocating control.

Acknowledgments

This work was supported by the Training Program for Interactionist Cognitive Neuroscience T32-MH115895 (X.L.), NIMH R01MH124849 (A.S.), and NIMH R21MH122863 (A.S.).

References

- Alport, A., Styles, E. A., & Shulan, H. (1994). Shifting Intentional Set: Exploring the Dynamic Control of Tasks. In *Attention and performance XV: Conscious and nonconscious information processing*. The MIT Press.
- 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.
- Grahek, I., Leng, X., Fahey, M. P., Yee, D., & Shenhav, A. (2022). Empirical and Computational Evidence for Reconfiguration Costs During Within-Task Adjustments in Cognitive Control. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *44*(44).
- Leng, X., Yee, D., Ritz, H., & Shenhav, A. (2021). Dissociable influences of reward and punishment on adaptive cognitive control. *PLOS Computational Biology*, *17*(12), <https://doi.org/10.1371/journal.pcbi.1009737>
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140.
- Musslick, S., Bizyaeva, A., Agaron, S., Leonard, N., & Cohen, J. D. (2019). Stability-Flexibility Dilemma in Cognitive Control: A Dynamical System Perspective. *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*.
- Navarro, D. J., & Fuss, I. G. (2009). Fast and accurate calculations for first-passage times in Wiener diffusion models. *Journal of Mathematical Psychology*, *53*(4), 222–230.
- Wiecki, T., Sofer, I., & Frank, M. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics*, *7*.