

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Empirical and Computational Evidence for Reconfiguration Costs During Within-Task Adjustments in Cognitive Control

Permalink

<https://escholarship.org/uc/item/4xq0c263>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Grahek, Ivan
Leng, Xiamin
Fahey, Mahalia Prater
et al.

Publication Date

2022

Peer reviewed

Empirical and Computational Evidence for Reconfiguration Costs During Within-Task Adjustments in Cognitive Control

Ivan Grahek*, Xiamin Leng*, Mahalia Prater Fahey, Debbie Yee, & Amitai Shenhav

Department of Cognitive, Linguistic, and Psychological Sciences; Carney Institute for Brain Science, Brown University

Corresponding author: ivan_grahek@brown.edu

*These authors have contributed equally to the study

Abstract

To achieve goals, people leverage cognitive control to adjust how they process information. Here we show that frequent adjustments in information processing strategies (e.g., response threshold) within a single task give rise to reconfiguration costs. In two experiments we induced different performance goals in a Stroop task via explicit instruction or incentives, and these goals either varied or were fixed across different blocks. Across both experiments, we find smaller adjustments in control intensity when people frequently adjust the amount of control they exert, relative to blocks in which they don't. We show that these results can be accounted for with a model that maximizes reward rate while minimizing reconfiguration costs (proportional to the Euclidean distances between the previous and current control signals). These findings suggest that cognitive control adjustments are regularized to constrain larger adjustments in control, which has important implications for computational modeling and measurement of motivated cognitive control.

Keywords: cognitive control; reconfiguration costs; motivation; drift diffusion model

Introduction

In order to achieve their goals, people engage cognitive control processes to adjust which information they pay attention to, and how they process this information. Many environments require rapid switches between different tasks, and frequent changes in the type of information which is being processed. Such situations incur switch-costs – detriments in behavioral performance when people need to focus on one aspect of their environment, after having been focused on a different one (Alport et al., 1994; Monsell, 2003). A rich literature on task switching has focused on people's ability to rapidly change between task sets which require them to process different aspects of their environment (Kiesel et al., 2010). However, many situations don't require switching between tasks, but rather adjustments of information processing strategies within a single task. For example, when incentives in an environment change, people have to adjust how much attention they pay to the task at hand, or how cautious they are when making decisions on how to act. Here we investigate the costs associated with such adjustments in *control intensity*.

Computational models of cognitive control propose that control intensity is determined through the maximization of the value of cognitive control (Shenhav et al., 2013; Verguts et al., 2015; Manohar et al., 2015; Alexander et al., 2018).

These models are supported by a wide range of studies showing that people adjust their levels of attention and caution based on their current goals, as reflected in changes in drift rate and threshold within a Drift Diffusion model (e.g., focus on speed vs. accuracy; Forstmann et al., 2008; Ratcliff & Rouder, 1998). Crucially, people allocate more attention when they expect higher performance-contingent rewards (Padmala & Pessoa, 2011; Krebs et al., 2012), and they adjust both drift rates and thresholds according to the level of expected rewards and penalties (Leng et al., 2021). This work suggests that people exert optimal levels of control given the expected incentives in their environment. Recent models of cognitive control propose that such value-based adjustments in control come with a cost arising from the need to adjust control signals (Lieder et al., 2018; Musslick et al., 2015, 2019; Musslick & Cohen, 2021). However, direct empirical tests of this proposal are lacking.

Here we consider different environments in which people have to readjust their control settings (*control intensity* measured as the magnitude of drift rates and thresholds; Ratcliff & McKoon, 2008), while doing a single cognitive control task. We show that in such environments control intensity is not adjusted to the optimal level suggested by a reward-rate optimal model (Bogacz et al., 2006; Leng et al., 2021). Rather, there is an inertia in the control system which regularizes larger adjustments in control intensity. In other words, we show that cognitive control adjustments depend not only on the optimal level of control given the current environment, but also on the magnitude of control adjustment needed to reach the value-optimal level from the previous control level. This finding suggests that cognitive control adjustments are regularized in a way that prevents large adjustments in how people process information (Ritz et al., 2022). Such control regularization has important implications for computational models of cognitive control, as well as for measurement of motivated cognitive control.

Reward rate model with reconfiguration costs

To simulate reward-optimal adjustments in cognitive control we used a reward-rate model (Figure 1) which finds the optimal levels of drift rate and threshold given the current performance goals (Bogacz et al., 2006; Leng et al., 2021). This model (Eq. 1) calculates the expected value of control (*EVC*) by considering the weights on the current value of giving a correct response (w_1) and making an error (w_2).

These weights are specified by the current performance goals (e.g., valuing speed over accuracy) which generate different experimental conditions (i). The weights scale the rate of correct responses ($1 - ER$) and the error rate (ER) obtained from the Drift Diffusion model (DDM; Bogacz et al., 2006) for specific values of drift rate (v) and threshold (a). The expected value of control includes two types of costs. The opportunity cost (Kurzban et al., 2013; Otto & Daw, 2019) which discounts the overall outcome and includes both the decision time (DT) and non-decision time (NDT). The second cost is the intensity cost (Musslick et al., 2015; Shenhav et al., 2013) represents the quadratic cost on cognitive control intensity, operationalized as the drift rate. The expected value of control is scaled by the probability of the occurrence of different conditions.

$$EVC(V, A, W, NDT) = \sum_{i=1}^N p_i \left[\frac{w_{1i} \times (1-ER(v_i, a_i)) - w_{2i} \times ER(v_i, a_i)}{DT(v_i, a_i) + NDT} - v_i^2 \right] \quad (\text{Eq. 1})$$

The model then finds the optimal set of drift rates (V) and thresholds (A) which maximize the expected value of control for a given set of weights (W), as specified by the different condition, and a non-decision time.

$$V, A = \text{argmax}_{V, A} [EVC|W, NDT] \quad (\text{Eq. 2})$$

However, this model predicts no differences in drift rates and thresholds between the environment in which people have to switch between different conditions and the environment which doesn't require switches. Crucially, here we extended the previous model by including a reconfiguration cost associated with readjusting drift rate and threshold levels across different experimental conditions (e.g., speed and accuracy; cf. Lieder et al., 2018; Musslick et al., 2018). We implemented this cost (Eq. 3) of moving between different conditions as the exponentiated Euclidian distance between the drift rate and threshold configurations in the two different conditions i and j . We opted for the exponentiated Euclidian distance following previous computational implementations of the EVC model (Musslick et al., 2015), but we confirmed that the qualitative pattern of results produced by the model

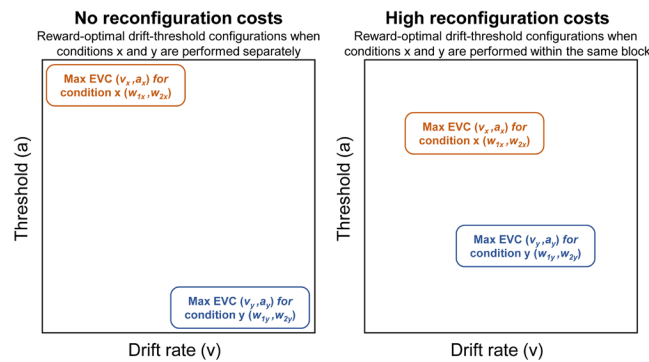


Figure 1. Configurations of drift rate and threshold which maximize the expected value of control in blocks with and without reconfiguration costs.

remains the same when using different distance metrics (e.g., Euclidian and Manhattan distance). This model reduces to the model in Equation 1 if the weight on the reconfiguration cost (RC) is set to 0. Otherwise, the model includes a certain level of reconfiguration cost. We included the weights on drift rate and threshold adjustments so that the model could capture the situation in which people weigh the reconfiguration cost on each dimension differently. We included C_v and C_a to capture potential bias between drift rate and threshold change in the cost of reconfiguration.

$$EVC(V, A, W, NDT, RC) = \sum_{i=1}^N p_i \left[\frac{w_{1i} \times (1-ER(v_i, a_i)) - w_{2i} \times ER(v_i, a_i)}{(DT(v_i, a_i) + NDT)} - v_i^2 \right] - RC \sum_{i=1}^N \sum_{j=1, j \neq i}^N \left[e^{\sqrt{C_v(v_i - v_j)^2 + C_a(a_i - a_j)^2}} \right] \quad (\text{Eq. 3})$$

The optimal values of drift and threshold are then calculated by maximizing the expected value of control (Eq. 2).

Experiment 1

Methods

Participants. We recruited 48 participants on Prolific, and excluded 4 of them due to failed attention checks, yielding the final sample of 44 participants (31 female; median age=30). The research protocol was approved by Brown University's Institutional Review Board.

Design. Participants performed a Stroop task in which they identified the ink color of the color word by pressing one of the 4 corresponding keys, while ignoring the word content. Participants performed the interval version of the Stroop task (Figure 2A) in which they completed as many trials as they wished during a fixed time period (8-12s). This task has been shown to produce reliable adjustments in both drift rates and thresholds in response to incentives, and participant's performance is well captured by a reward-rate model (Leng et al., 2021). Prior to each interval, participants were presented with a cue (1.5 s; Figure 2C) instructing them to perform the task either as quickly or as accurately as possible (cf. Forstmann et al., 2008; Ratcliff & Rouder, 1998). Participants received feedback (1.5 s) on how many correct responses they gave. Crucially, there were two types of blocks (Figure 2B) which yielded different levels of reconfiguration costs. In Fixed blocks, participants performed only one condition on every interval (e.g., focus on speed). In Varying blocks, participants performed multiple conditions within the block (e.g., focus on speed, followed by focus on accuracy). We predicted that Fixed blocks will not incur reconfiguration costs, while Varying blocks will. Participants performed 4 blocks, each of which included 20 intervals. For half of the blocks, the instructed condition was fixed over the entire block, meaning that in these blocks participants always received a cue telling them to focus on the same dimension of performance (e.g., speed).

The other half of the blocks were Varying, meaning that the performance goals changes within a block. Before each block, participants were informed whether this will be a Fixed or a Varying block, these blocks were intermixed, and **their order was counterbalanced across participants.**

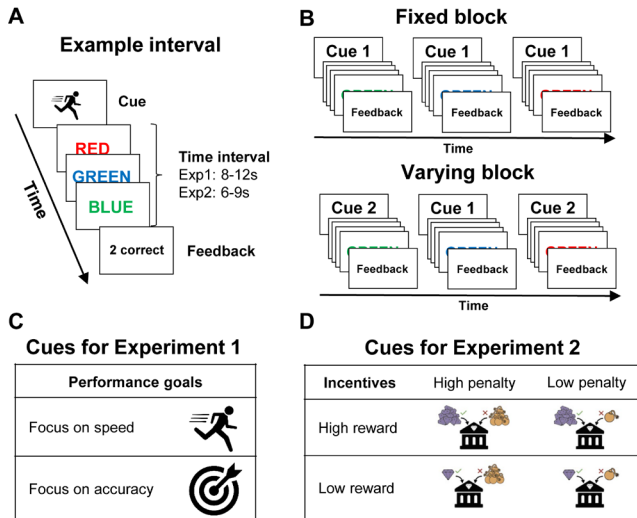


Figure 2. Experimental design. **A. One interval.** Each interval in the experiment consisted of a fixed time interval during which participants could complete as many Stroop trials as they wished. **B. Block types.** Fixed blocks refer to blocks in which a given condition remained constant throughout the block, whereas Varying blocks refer to blocks in which that condition (instruction in Experiment 1, incentive level in Experiment 2) varied across the block. **C. Experimental conditions in Experiment 1.** Participants were presented with cues instructing them to perform the task as fast or as accurately as they can. **D. Experimental conditions in Experiment 2.** Participants were presented with cues which indicated that they could earn high or low rewards for correct responses, and that they could receive high or low penalties for errors.

Statistical analyses and Drift Diffusion modeling. To predict reaction times and accuracies on each trial we fitted hierarchical linear mixed models (lme4 package in R; Bates et al., 2015) and included congruency, interval length, interval and block type and their interaction as predictors. All of the main effects and the intercept were also included as random effects. Reaction times were log-transformed for these analyses. We then fitted a hierarchical Bayesian drift diffusion model (HDDM; Wiecki et al., 2013) to the reaction time and accuracy data. The fitted model includes the effect of congruency on drift rate, as well as the effects of instruction type (speed vs. accuracy), block type (Varying vs. Fixed), and the interaction between instruction and block type on drift rate. We also included the same effects on the threshold, but without the effect of congruency. Finally, we included the effect of instruction type on non-decision time. We ran 5 MCMC simulations (chains; 10000 iterations; 8000 warmup) to estimate the model. We confirmed chain convergence by examining trace plots and the ratio of

variances between and within chains (Gelman-Rubin statistic), and posterior predictive checks. We summarized the obtained posterior distributions by reporting their means, 95% credible intervals, and the probability that the parameter of interest is higher than 0 (e.g., $p_{b < 0} = 0.01$).

Computational model. In order to simulate the environment in Experiment 1 we used the model presented in Equation 3. In this experiment, participants were instructed to focus either on speed or on accuracy. We relied on the empirical DDM fits in order to simulate the performance in the similar range as the empirically observed values. We employed an inverse reward-rate optimization procedure (Leng et al., 2021) to set the baseline values for the value of giving a correct response (w_1) and the value of committing an error (w_2). The inverse reward-rate optimization was based on the empirically obtained group-level estimates of the non-decision time, and the drift rate and threshold levels in the speed and accuracy conditions. This inverse optimization procedure yielded the weights for each of the 2 incentive conditions. With these weights we simulated the blocks in which the condition of each interval is fixed as speed or accuracy ($RC = 0.05$) and the blocks in which the condition varies between speed and accuracy randomly ($RC = 0.4$). We fixed the relative weight of drift rate adjustment in reconfiguration cost (C_r) as 1 and varied the weight for threshold adjustment to generate predictions matched with empirical results. The model matches with the empirical finding with $C_a = 10$. We then used a gradient descent algorithm to find the configuration of drift rates and thresholds which maximize the expected value of control. The drift rate and threshold configurations obtained in this way were then compared to the empirical findings.

Results

Reaction times and accuracy. As shown in Figure 3A, participants were faster ($b=0.13$; 95% CI [0.09, 0.18]; $p < 0.001$) and less accurate ($b=6.88$; 95% CI [4.13, 11.45]; $p < 0.001$) when instructed to focus on speed relative to when instructed to focus on accuracy. Crucially, when instructed to focus on speed, they were more accurate in the Varying compared to Fixed blocks ($b=1.33$; 95% CI [1.12, 1.59]; $p < 0.001$). When instructed to focus on accuracy there was a non-significant trend toward being more accurate in the Fixed relative to the Varying blocks ($b=0.85$; 95% CI [0.70, 1.03]; $p=0.104$). Further, in the speed condition participants were faster to respond in Fixed than in Varying blocks ($b=0.02$; 95% CI [0.70, 1.03]; $p=0.104$), and there were no significant differences in the accuracy condition ($b=0.00$; 95% CI [-0.01, 0.01]; $p=0.888$). These results indicate that the performance in the speed condition was more similar to the performance in the accuracy condition when participants had to switch between these strategies within one block relative to when they had only one instruction per block.

Drift diffusion model. When they were instructed to focus on accuracy, relative to when they were instructed to focus on speed (Figure 3B – left), participants had higher drift rates ($b=0.68$; 95% CrI [0.43, 0.94]; $p_{b<0}<0.01$), thresholds ($b=0.39$; 95% CrI [0.23, 0.55]; $p_{b<0}<0.01$), and non-decision times ($b=0.04$; 95% CrI [0.02, 0.07]; $p_{b<0}<0.001$). Crucially, this difference between accuracy and speed in thresholds was smaller in blocks that included both conditions compared to blocks with only one condition ($b=0.07$; 95% CrI [0.01, 0.13]; $p_{b<0}<0.01$). We found no such block differences in drift rates ($b=0.05$; 95% CrI [-0.11, 0.20]; $p_{b<0}=0.27$).

Model simulation results. Results of the model simulations (Figure 3C) show that the model was able to capture the main pattern of the empirically observed effects (Figure 3B). The model replicated the larger adjustment of thresholds in the Fixed relative to the Varying blocks, and only a small adjustment of drift rates.

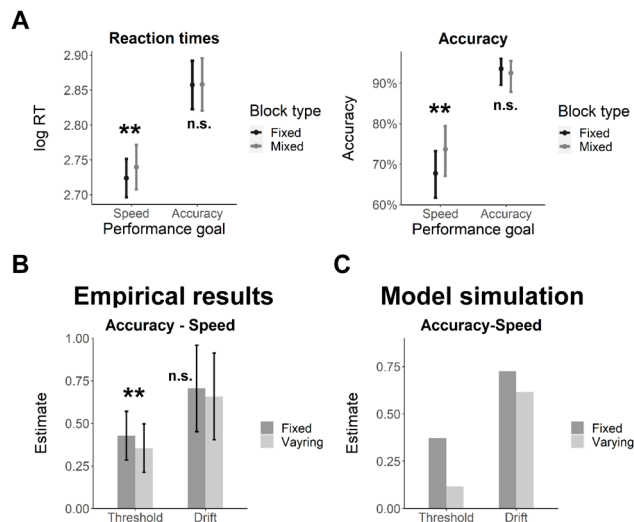


Figure 3: Experiment 1 results and model simulations. **A. Reaction times and accuracies.** Regression estimates of the reaction times (left) and accuracies (right) for the speed and accuracy conditions in blocks in which participants either have one performance goal (Fixed) or are switching between two performance goals (Varying). Error bars represent 95% confidence intervals and **: $p<0.001$. **B. Drift diffusion modeling results.** Drift rate and threshold differences between the accuracy and the speed condition in Fixed and Varying blocks. Error bars represent 95% credible intervals and **: the ratio of posterior samples on two sides of 0 is <0.001 . **C. Model predictions.** Model simulations for the difference in drift rates and thresholds between the accuracy and speed conditions in Fixed and Varying blocks.

Discussion

The results of Experiment 1 show that inducing the performance goals of focusing on either speed or accuracy in a classical cognitive control task produces changes in both response thresholds and drift rates. While the focus on

accuracy relative to speed increased both the thresholds and drift rates, this difference was reduced for thresholds when the participants had to switch between the two performance goals within a block compared to the blocks in which they only had one performance goal. This pattern of results can be successfully simulated with a computational model that assumes that there is a reconfiguration cost operationalized as the distance in the drift-threshold space between the two conditions. These empirical findings and model simulations provide evidence for the existence of reconfiguration costs in situations in which participants are switching between explicitly induced performance goals within a single task. In Experiment 2 we sought to provide further evidence for the existence of reconfiguration costs in the absence of explicitly induced performance goals. Rather, we implicitly induced performance goals by differentially incentivizing performance with high or low monetary incentives.

Experiment 2

Methods

Participants. We recruited 80 participants who were paid to perform the task on Prolific. The final sample included 69 participants (34 female; median age=21; 11 excluded due to failed attention checks). The research protocol was approved by Brown University's Institutional Review Board.

Design. Participants performed the same Stroop task as in Experiment 1 (Figure 2A), completing as many trials as they wished within a fixed time interval (6-9s in this experiment). Crucially, rather than explicitly inducing performance goals (via cued instructions as in Experiment 1), we differentially incentivized performance goals by offering rewards (earning 100 points vs. 1 point) for correct responses, and penalties (losing 100 points vs. 1 point) for errors (cf. Leng et al., 2021). Prior to each interval, participants saw a cue (1 s) informing them about the reward and penalty levels within the interval (Figure 2D). After each interval they received feedback (1 s) about how many rewards and penalties they earned. Within a block, one incentive type always varied (e.g., high vs. low reward), while the other was fixed (e.g., low penalty). This meant that in some blocks the reward manipulation induced reconfiguration costs, while the penalty manipulation did not, and vice versa (Figure 2B). Participants performed 4 blocks, out of which 2 were with varying rewards, and 2 with varying penalties. Each block consisted of 15 intervals, and at block onset participants were informed which incentive will be fixed and at which level (e.g., low penalty), and which incentive will be varying (e.g., switching between high and low rewards). The block order was randomized across participants.

Statistical analyses and Drift Diffusion modeling. We used the same procedure as in Experiment 1 to analyze the reaction time and accuracy data, and fit the Drift Diffusion model. The fitted model included the effect of congruency on drift rate, as well as the effects of the reward and penalty levels, the

effect of reward vs. penalty being fixed, and the interaction between the last two predictors. We allowed threshold to vary according to the all of the same predictors except congruency.

Computational model. In order to simulate the environment in Experiment 2 we used the model presented in Equation 3. In this experiment participants received high or low rewards for correct responses ($w_{1HighReward}$ and $w_{1LowReward}$), and high or low penalties for errors ($w_{2HighPenalty}$ and $w_{2LowPenalty}$). We relied on the empirical group-level DDM fits in order to simulate performance in the similar range as the empirically observed values. We employed an inverse reward-rate optimization procedure (Leng et al., 2021) to set the baseline values for the weights on the correct responses (w_1) and errors (w_2) based on the empirically obtained estimate of non-decision time, and the drift rate and threshold levels for each pair of reward and penalty values. This inverse optimization procedure yielded the weights for each of the 4 incentive conditions. With these weights we simulated the blocks in which reward levels varied (Reward-Varying; $RC = 0.4$) and the blocks in which reward levels were fixed, but penalty levels varied (Reward-Fixed; $RC = 0.05$). We did the same for the Penalty-Varying vs. Penalty-Fixed. We then found the values of drift rate and threshold which maximized the expected value of control and compared them to the empirically observed results.

Results

Reaction times and accuracy. Participants performed faster ($b=-0.03$; 95% CI [-0.04, -0.03]; $p<0.001$; Figure 4A), but less accurately on intervals with high relative to low rewards ($b=0.79$; 95% CI [0.68, 0.91]; $p<0.001$). Conversely, they were slower ($b=0.02$; 95% CI [0.02, 0.03]; $p<0.001$), but more accurate ($b=1.75$; 95% CI [1.51, 2.04]; $p<0.001$) on high relative to low penalty intervals. Importantly, we found substantial performance differences between the Fixed and the Varying blocks. The difference between high and low reward intervals was larger in Reward-Varying relative to Reward-Fixed blocks for both reaction times ($b=-0.01$; 95% CI [-0.02, -0.01]; $p<0.001$) and accuracies ($b=0.70$; 95% CI [0.55, 0.89]; $p<0.05$). The same was true for penalty effect when comparing Penalty-Fixed and Penalty-Varying intervals on both reaction times ($b=-0.02$; 95% CI [-0.03, -0.01]; $p<0.001$) and accuracies ($b=0.60$; 95% CI [0.47, 0.77]; $p<0.001$). These results show that the difference in performance between high and low incentive conditions was larger in blocks in which the relevant incentive was fixed, compared to blocks in which it was varying.

Drift diffusion model. Consistent with the previous findings with this task (Leng et al., 2021), participants exhibited lower thresholds for high relative to low reward ($b=-0.12$; 95% CrI [-0.18, -0.05]; $p_{b<0}<0.001$; Figure 4B), while maintaining a similar drift rate ($b=0.04$; 95% CrI [-0.07, 0.16]; $p_{b=0}=0.23$). For high relative to low penalties, they exhibited higher

thresholds ($b=0.15$; 95% CrI [0.07, 0.24]; $p_{b<0}<0.001$) and higher drift rates ($b=0.14$; 95% CrI [0.02, 0.26]; $p_{b<0}<0.001$). Most importantly, the difference between the threshold for high relative to low reward was smaller in Reward-Varying relative to Reward-Fixed blocks ($b=-0.06$; 95% CrI [-0.11, -0.01]; $p_{b<0}<0.001$). Similarly, the difference in threshold between high and low penalty levels was greater for Penalty-Fixed than for Penalty-Varying blocks ($b=0.10$; 95% CrI [0.05, 0.15]; $p_{b<0}<0.001$). We did not find significant effects of block type on drift rate adjustments.

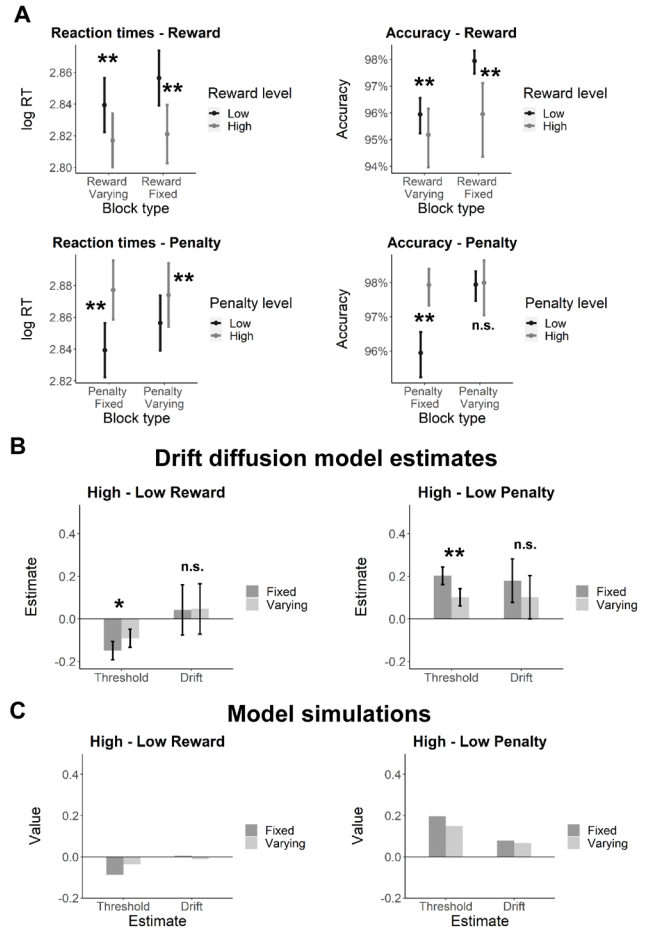


Figure 4: Experiment 2 results and model simulations.

A. Reaction times and accuracies. Regression estimates of the reaction times and accuracies for the reward (up) and the penalty (down) conditions for the blocks in which the relevant incentive is either fixed (Fixed) or changing (Varying). Error bars represent 95% confidence intervals and **: $p<0.001$. **B. Drift diffusion modeling results.** Drift rate and threshold estimates of the differences between the high and low value of the relevant incentive in blocks in which that incentive is either fixed or varying. Error bars represent 95% credible intervals and **/*: the ratio of posterior samples on two sides of 0 is $<0.001/0.05$. **C. Model predictions.** Model simulations for the difference in the drift rates and thresholds between the high and low value of the relevant incentive when that incentive is either fixed or varying.

Model simulation results. Predictions from the model simulations (Figure 4C) fully capture the observed empirical results (Figure 4B). The model predicts lower thresholds for high relative to low rewards, and predicts that this effect will be more pronounced in the fixed relative to the varying blocks (Figure 4C left). Further, the model predicts higher thresholds and drift rates for high relative to low penalty conditions, and, crucially, that this difference will be more pronounced in fixed relative to varying blocks (Figure 4C middle). This prediction is in line with the observed results.

Discussion

In Experiment 2 we sought to provide further evidence for the existence of reconfiguration costs within a single task by implicitly inducing different performance goals. We incentivized participants' performance by offering high or low performance-based monetary reward and penalties. This manipulation led to lower thresholds when high relative to low rewards were on offer, and higher thresholds and drift rates when high relative to low penalties were expected. Crucially, we found that the changes in thresholds for both the reward and the penalty effect were reduced in blocks in which these incentives were varying (e.g., switching between high and low rewards in one block) compared to blocks in which the incentive was fixed (e.g., always high reward, but varying levels of penalties). This pattern of findings was well captured by a computational model that included different levels of rewards and penalties, but also, crucially, reconfiguration costs which regularized the movement in the drift-threshold space. These findings provide further support for the existence of reconfiguration costs arising from control intensity adjustments within a single task.

General discussion

People engage cognitive control to adjust which information they process, and how they process it, in order to accomplish their goals. While research on task switching has demonstrated the costs associated with changing which information is processed between different tasks, much less is known about the costs associated with adjusting control intensity within a single task. Here we demonstrate that adjustments in cognitive control intensity within a single task are associated with reconfiguration costs. These costs regularize the magnitude of cognitive control adjustments and bias against large fluctuations in control.

Across two tasks, we induced explicit (Experiment 1) and implicit (Experiment 2) performance goals to create conditions in which different configurations of control signals (drift rates and thresholds) are optimal. Participants performed these conditions in blocks in which they either had to frequently switch between varying performance goals (high reconfiguration costs; e.g., focusing on speed after having just focused on accuracy), or had only one fixed performance goal (low reconfiguration costs; e.g., focus on speed on every interval in one block). In both experiments we show that adjustments in response thresholds were smaller in

blocks with low relative to high reconfiguration costs. We used a reward-rate optimal model (Bogacz et al., 2006; Leng et al., 2021) to simulate performance of artificial agents in these two experiments. This model is able to capture most of the empirically observed patterns, but only if it includes reconfiguration costs (operationalized as the Euclidian distance on the difference between the drift rate and threshold levels across conditions in a block). This finding provides support for the computational models of control which include reconfiguration costs (Lieder et al., 2018; Manohar et al., 2015; Musslick et al., 2015, 2019), and extends them to include costs on threshold adjustments.

While our empirical findings and modeling results support the existence of reconfiguration costs within a single task, several important questions are left open. Our current experimental designs cannot fully rule out the possibility that participants are sometimes making the mistake of having the wrong performance goal for a given interval in the varying blocks (e.g., focusing on speed despite the cue telling them to focus on accuracy). One argument against this idea is that a computational model which would perform our tasks in this way would have difficulty capturing the asymmetric patterns of reconfiguration costs we observed for threshold relative to drift rate adjustments. However, further experimental work is needed in order to fully rule out this possibility.

The observed asymmetries in reconfiguration costs are themselves notable. While our model was able to capture them by placing a higher reconfiguration cost on threshold than drift, an important question for future work is whether this asymmetry reflects a general property of these control signals or is specific to the drivers of adjustments in our current tasks. Given that prior modeling of the task in Experiment 2 (Leng et al., 2021) found an important role for intensity costs on drift rates (but not necessarily for thresholds), it would also be valuable to examine whether there is a more general trade-off between these two forms of costs. Future experiments could address this question by investigating reconfigurations costs in experiments that encourage larger changes in drift rates (e.g., by varying reward magnitude substantially while maintaining low penalties). Finally, here we have focused on the costs associated with adjusting the levels of attention and caution, and fixed the other DDM parameters, but future work should examine the possibility that adjustments in other parameters (e.g., diffusion noise) could also incur reconfiguration costs.

Here we offer a way of measuring reconfiguration costs, which can be used to compare different models and implementations of reconfiguration costs. The existence of reconfiguration costs is critical for improved measurement of incentivized cognitive control. Classical paradigms which measure the effects of motivation on cognitive control (Botvinick & Braver, 2015; Parro et al., 2018; Yee & Braver, 2018) commonly include varying incentives. Reconfiguration costs will arise in these situations, decreasing the experimenter's sensitivity to these effects. We formalize the cause of this decrease, and offer a way to maximize measurement precision on the effects of interest.

Acknowledgements

This work was supported the Training Program for Interactionist Cognitive Neuroscience T32-MH115895 (X.L.), an NSF Graduate Research Fellowship (M.P.F.), Training Program for Computational Psychiatry T32-MH126388 (D.Y.), NIMH R01MH124849 (A.S.), and NIMH R21MH122863 (A.S.).

References

- Alexander, W. H., Brown, J. W., Collins, A. G. E., Hayden, B. Y., & Vassena, E. (2018). Prefrontal Cortex in Control: Broadening the Scope to Identify Mechanisms. *Journal of Cognitive Neuroscience*, *30*(8), 1061–1065. https://doi.org/10.1162/jocn_a_01154
- 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. <https://doi.org/10.7551/mitpress/1478.001.0001>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*, 1–48. <https://doi.org/10.18637/jss.v067.i01>
- 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. <https://doi.org/10.1037/0033-295X.113.4.700>
- Botvinick, M., & Braver, T. (2015). Motivation and Cognitive Control: From Behavior to Neural Mechanism. *Annual Review of Psychology*, *66*(1), 83–113. <https://doi.org/10.1146/annurev-psych-010814-015044>
- Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Cramon, D. Y. von, Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, *105*(45), 17538–17542. <https://doi.org/10.1073/pnas.0805903105>
- Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching—A review. *Psychological Bulletin*, *136*(5), 849–874. <https://doi.org/10.1037/a0019842>
- Krebs, R. M., Boehler, C. N., Roberts, K. C., Song, A. W., & Woldorff, M. G. (2012). The Involvement of the Dopaminergic Midbrain and Cortico-Striatal-Thalamic Circuits in the Integration of Reward Prospect and Attentional Task Demands. *Cerebral Cortex*, *22*(3), 607–615. <https://doi.org/10.1093/cercor/bhr134>
- Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, *36*(6), 661–679. <https://doi.org/10.1017/S0140525X12003196>
- Leng, X., Yee, D., Ritz, H., & Shenhav, A. (2021). Dissociable influences of reward and punishment on adaptive cognitive control. *PLOS Computational Biology*, *17*(12), e1009737. <https://doi.org/10.1371/journal.pcbi.1009737>
- Lieder, F., Shenhav, A., Musslick, S., & Griffiths, T. L. (2018). Rational metareasoning and the plasticity of cognitive control. *PLOS Computational Biology*, *14*(4), e1006043. <https://doi.org/10.1371/journal.pcbi.1006043>
- Manohar, S. G., Chong, T. T.-J., Apps, M. A. J., Batla, A., Stamelou, M., Jarman, P. R., Bhatia, K. P., & Husain, M. (2015). Reward Pays the Cost of Noise Reduction in Motor and Cognitive Control. *Current Biology*, *25*(13), 1707–1716. <https://doi.org/10.1016/j.cub.2015.05.038>
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140. [https://doi.org/10.1016/S1364-6613\(03\)00028-7](https://doi.org/10.1016/S1364-6613(03)00028-7)
- 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*. <https://par.nsf.gov/biblio/10125021-stability-flexibility-dilemma-cognitive-control-dynamical-system-perspective>
- Musslick, S., Botvinick, M. M., Shenhav, A., & Cohen, J. D. (2015). *A computational model of control allocation based on the Expected Value of Control*. 6.
- Musslick, S., & Cohen, J. D. (2021). Rationalizing constraints on the capacity for cognitive control. *Trends in Cognitive Sciences*, *25*(9), 757–775. <https://doi.org/10.1016/j.tics.2021.06.001>
- Otto, A. R., & Daw, N. D. (2019). The opportunity cost of time modulates cognitive effort. *Neuropsychologia*, *123*, 92–105. <https://doi.org/10.1016/j.neuropsychologia.2018.05.006>
- Padmala, S., & Pessoa, L. (2011). Reward Reduces Conflict by Enhancing Attentional Control and Biasing Visual Cortical Processing. *Journal of Cognitive Neuroscience*, *23*(11), 3419–3432. https://doi.org/10.1162/jocn_a_00011
- Parro, C., Dixon, M. L., & Christoff, K. (2018). The neural basis of motivational influences on cognitive control. *Human Brain Mapping*, *39*(12), 5097–5111. <https://doi.org/10.1002/hbm.24348>
- Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. <https://doi.org/10.1162/neco.2008.12-06-420>

- Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, 9(5), 347–356.
<https://doi.org/10.1111/1467-9280.00067>
- Ritz, H., Leng, X., & Shenhav, A. (2022). Cognitive Control as a Multivariate Optimization Problem. *Journal of Cognitive Neuroscience*, 23.
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function. *Neuron*, 79(2), 217–240.
<https://doi.org/10.1016/j.neuron.2013.07.007>
- Verguts, T., Vassena, E., & Silvetti, M. (2015). Adaptive effort investment in cognitive and physical tasks: A neurocomputational model. *Frontiers in Behavioral Neuroscience*, 9, 57.
<https://doi.org/10.3389/fnbeh.2015.00057>
- Wiecki, T., Sofer, I., & Frank, M. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics*, 7.
<https://www.frontiersin.org/article/10.3389/fninf.2013.00014>
- Yee, D. M., & Braver, T. S. (2018). Interactions of motivation and cognitive control. *Current Opinion in Behavioral Sciences*, 19, 83–90.
<https://doi.org/10.1016/j.cobeha.2017.11.009>