

***Title:* Quantifying the cost of cognitive stability and flexibility**

Authors:

Danae Papadopetraki^{1,2} Monja I. Froböse¹, Andrew Westbrook³, Bram B. Zandbelt^{1,2}, Roshan Cools^{1,2}

Affiliations:

¹Radboud University, Donders Institute for Brain, Cognition and Behaviour, Centre for Cognitive Neuroimaging, Nijmegen

²Radboud university Medical Center, Department of Psychiatry, Nijmegen

³Brown University, Department of Cognitive, Linguistics, and Psychological Sciences, Providence

Correspondence:

Danae Papadopetraki, Kapittelweg 29, 6525 EN, Nijmegen.

E-mail: d.papadopetraki@donders.ru.nl.

Data and analysis code: <https://github.com/danae1968/stabflex2019>

Acknowledgements: This research was supported by a VICI grand from NWO (Grant No. 453-14-015). We thank Rebecca Calcott^{1,2} and Lieke Hofmans^{1,2} for comments that greatly improved the manuscript.

Abstract

Exerting cognitive control is well known to be accompanied by a subjective effort cost and people are generally biased to avoid it. However, the nature of cognitive control costs is currently unclear. Recent theorizing suggests that the cost of cognitive effort serves as a motivational signal to bias the system away from excessive focusing (i.e. cognitive stability) and towards more cognitive flexibility. We asked whether the effort cost of cognitive stability is higher than that of cognitive flexibility. Specifically, we tested this prediction in the domain of working memory by using (i) a delayed response paradigm that allows us to manipulate demands for stability (distractor resistance) and flexibility (flexible updating) of working memory representations, as well as (ii) a subsequent cognitive effort discounting paradigm that allows us to quantify the subjective effort costs assigned to performing the delayed response paradigm. We show strong evidence, in two different samples (28 and 62 participants respectively) that subjective cost increases as a function of demand. Moreover, we demonstrate that the subjective cost of performing a task requiring cognitive stability (distractor resistance) is higher than that requiring flexible updating, supporting the hypothesis that the subjective effort cost of cognitive stability is higher than that of flexibility.

Introduction

Cognitive control often refers to the set of mechanisms required to focus on and pursue a goal, especially in the face of distraction, temptation or conflict. Succeeding to exert cognitive control and focusing on the task at hand is highly valued in our industrialized society, as it allows us to complete our tasks and achieve our long-term goals. Despite its importance, failures of cognitive control are very common. Procrastinating, failing to meet deadlines and performance decrements after fatigue are such examples familiar to most of us.

Why do people fail to exert cognitive control? Focusing on a task carries an effort cost, making people biased to avoid it¹, even if such avoidance implies forgoing rewards^{2,3}. The mechanisms underlying these cognitive effort costs remain elusive. While poor performance on cognitive control tasks has often been explained as limitations in cognitive capacity, more recent accounts shift the focus from capacity to motivation⁴. These accounts are supported by experiments that show that performance decrements (caused by effort) can be overcome by increases in incentive motivation, for example as a function of monetary rewards⁵. According to some such resource allocation accounts, the subjective cost of cognitive effort represents a motivational signal to remain open to alternative opportunities, thus promoting flexibility even at the expense of reduced engagement in a current ongoing task^{1,6-8}. As our attentional resources are limited^{9,10}, focusing on a given task means that we have to give up on other tasks that require the same set of mechanisms, thus incurring an opportunity cost⁶. Hence, failures of cognitive control can be viewed as stemming not just from failures in *implementation*, but also as a *choice* to pursue alternative tasks that may be more rewarding.

Such a motivational mechanism would be adaptive, given that our constantly changing environment requires a dynamic balance between the cognitive states of focus and flexibility^{8,11}. Focusing is crucial for completing our goals, but flexibility is essential when goals change. Flexibility also allows us to explore alternative ways to solve a problem and come up with new ideas, i.e. to be innovative and creative. Mind-wandering, for example, has been associated with

more “a-ha” moments when problem- solving^{12–15} and practicing voluntary mind-wandering has been proposed as a training method to boost creativity¹⁶.

According to current theorizing¹⁷, the stability/flexibility tradeoff in working memory is moderated by the strength of current task representations. Strong representations facilitate focusing on a current task-set at the cost of reduced flexibility, for example when task- switching. Weak representations in contrast allow flexible adaptation but reduce focused intensity.

How do we decide when to be focused and when to relax constraints in order to be flexible? We have previously argued that we arbitrate between a focused (closed) state versus a flexible (open) one, based on a cost-benefit analysis in which the benefit of cognitive effort corresponds to increased focus and is weighted against its (e.g. opportunity) cost, corresponding to reduced flexibility¹¹. We thus reasoned that the cost of cognitive stability is higher than that of cognitive flexibility. Here, we investigate this hypothesis by using a novel version of the cognitive effort discounting paradigm (COGED)³ that allowed us to measure the cost that people assign to performing tasks requiring cognitive stability or flexibility. Specifically, rather than asking participants to discount monetary offers to perform the N-back task, which requires both focusing and flexibility at the same time, we asked participants to discount offers to perform a working memory task requiring either stable maintenance and distractor resistance or flexible updating. This design allowed us to separately quantify the subjective costs of a task with demands for greater stability or flexibility, respectively. We obtained two independent datasets to replicate, and robustly establish the predicted differences between the costs of focus and flexibility.

As in the case of the original COGED paradigm, our paradigm consisted of two stages. In the first stage, subjects performed variants of a well-established colour wheel working memory task¹⁸. Participants experienced different demands (set sizes 1 to 4) of the two conditions of the task. One condition required flexible updating; the other condition required focused distractor resistance (stability). In the second stage, participants made a series of choices between repeating one of the working memory conditions in return for monetary rewards. Some trials required choices between either one of the (stability or flexibility) task conditions versus taking a break.

Other trials required direct comparisons between the two (flexibility versus stability) task conditions.

Results

Working memory task performance

We investigated the effect of demands for working memory stability versus flexibility using a modified color wheel task (Figure 1A, see Methods section for more details). Participants were exposed to conditions requiring either distractor resistance (i.e. ignore condition) or flexible updating (i.e. update condition). Every trial of the paradigm consisted of three phases that were separated by two delay periods. In the first phase (*encoding*), participants saw coloured squares, which they always had to memorize. Then after a delay of two seconds, participants saw new colours in the same square locations (*interference phase*). In the ignore condition, participants were instructed to maintain in their memory the colors from the encoding phase and not be distracted by the new interfering colors. In the update condition, participants had to let go of their initial representations and update the new stimuli into their working memory. We manipulated the working memory demand by varying the number of stimuli that needed to be remembered. During the response phase, participants had to match the color of one of the relevant squares by clicking with the mouse on a color wheel.

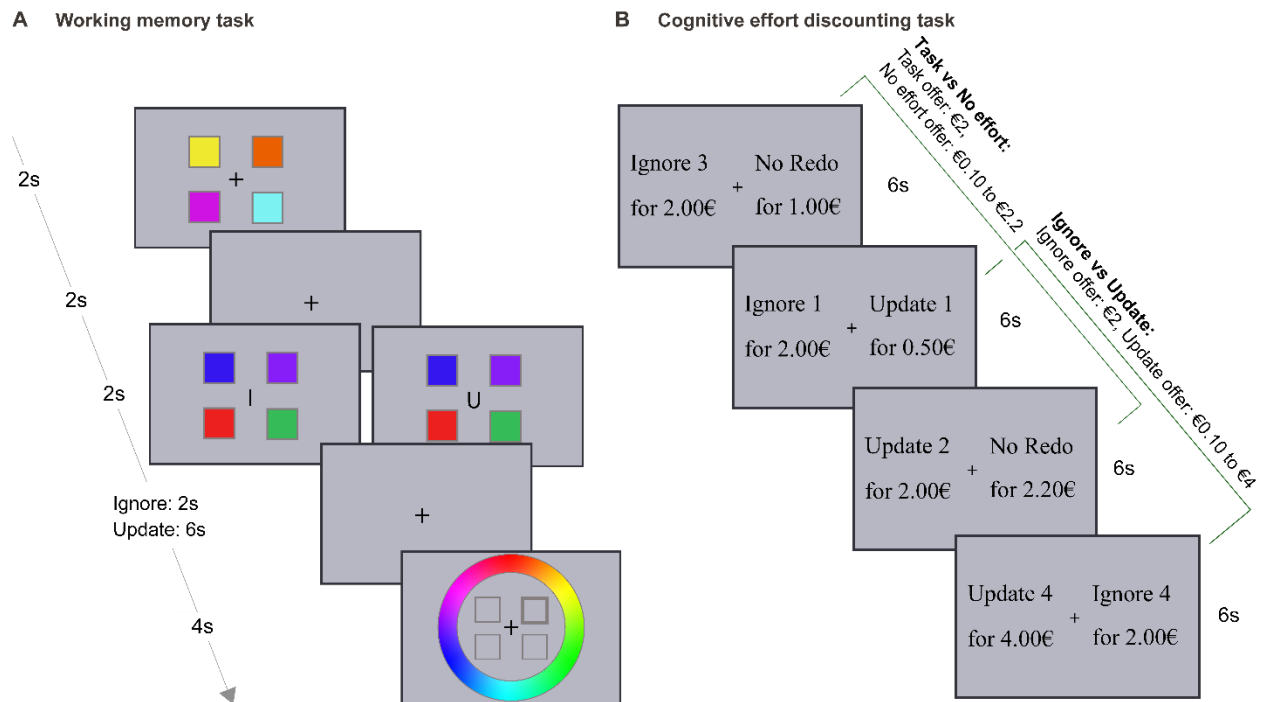


Figure 1. A An illustration of the color wheel working memory task. Every trial of the task consists of three phases. In the encoding phase (2 s), participants need to memorize colored squares. After a delay of 2 s, during the interference phase (2 s), a letter indicates if it is an ignore (I for ignore) or an updating (U for update) trial. In ignore trials, participants need to maintain colors from the encoding phase and ignore intervening stimuli. In update trials, participants have to let go of their previous representations and update into their memory the stimuli from the interference phase. Another delay separates interference from the response phase. This delay is 2 s for ignore and 6 s for update trials to match the time that the target stimuli are maintained between conditions. During response phase, participants see a color wheel and black frames of the same squares; they have 4 s to click on the target color for the highlighted square. Demand is manipulated by varying the number of squares from one to four. The example displayed here is of the highest demand. **B** Example trials of the COGED task. Participants perform two versions of choices. In the “task vs no effort” version, participants have to choose between repeating a level of ignore or update and not repeating the color wheel task at all (“No Redo”). The task option offer remains fixed at €2 and the no effort “No Redo” option varies from €0.1 to €2.2. In the “Ignore vs Update” trials, participants must choose between repeating either the ignore or update condition of the same demand. Ignore offers are always fixed at €2 and update offers vary from €0.1 to €4. Trial duration is 6 s. The trials are intermixed.

Accuracy

Performance on the working memory task was sensitive to the demand (i.e. set size) manipulation and, in line with earlier studies contrasting ignore and update trials, participants performed more

poorly in the ignore compared with the update condition^{19,20} (Figure 2A&B; Supplemental Table 1 for descriptive statistics; Supplemental Figure 1 for precision indices). This observation was supported by Bayesian model comparison (Table 1), showing strongest support for the model including set size and condition in both studies ($BF_{10} = 24876$ and $BF_{10} = 5.5e+12$, respectively). The runner-up model was the one including both main effects and their interaction, which was ~ 3.2 and ~ 2 times less likely than the model with the main effects only for experiment 1 and 2 respectively. The effects analysis confirmed the conclusion based on model comparison, showing that accuracy decreased with increasing set size (Experiment 1: $F_{1.52,39.6} = 6.510$, $p = 0.007$, $BF_{INC} = 83$, Experiment 2: $F_{1.63,97.7} = 16.998$, $p = 2.8e-6$, $BF_{INC} = 1.6e+10$) and that participants performed better on update compared with ignore trials (Experiment 1: $F_{1,26} = 11.068$, $p = 0.003$, $BF_{INC} = 448$, Experiment 2: $F_{1,60} = 24.095$, $p = 7.4e-6$, $BF_{INC} = 939$) (Table 2). Evidence for an interaction effect was not conclusive (Experiment 1: $F_{2,14,78} = 2.205$, $p = 0.116$, $BF_{INC} = 1.3$, Experiment 2: $F_{2,32,139.3} = 3.238$, $p = 0.035$, $BF_{INC} = 1.9$). We conclude that accuracy decreased as a function of set size and, across set sizes, was worse on ignore relative to update trials.

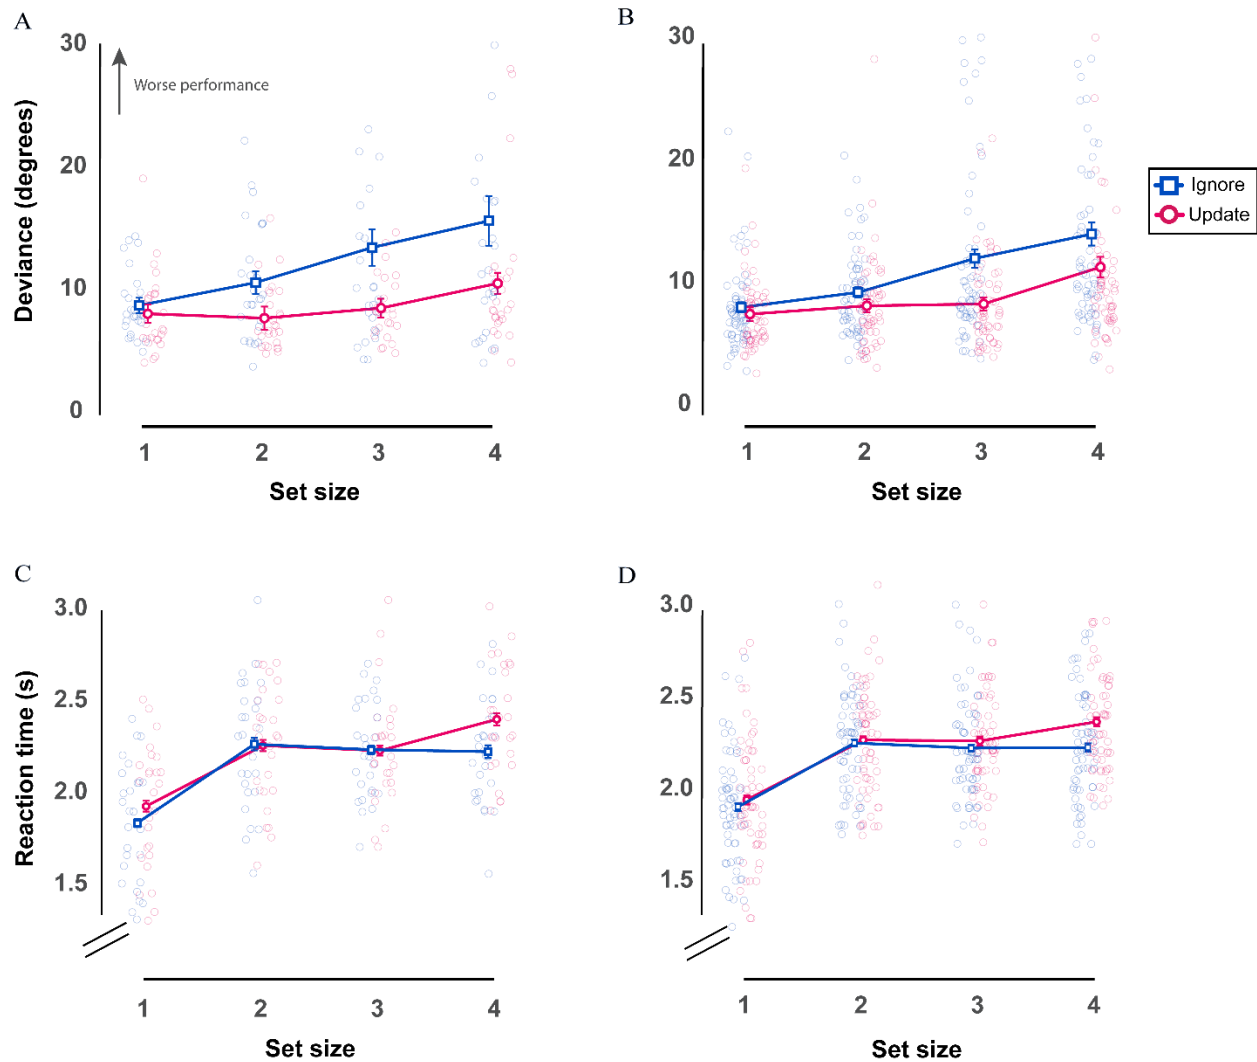


Figure 2. Performance on the color wheel working memory task. **A** Median deviance for experiment 1 (27 participants). **B** Median deviance for replication experiment 2 (61 participants). Deviance in degrees from the correct color is displayed here as a function of set size for ignore and update trials. **C&D** Median reaction times as a function of set size for ignore and update conditions. **C** Experiment 1 (27 participants). **D** Experiment 2 (61 participants). Error bars indicate within-participant SEM^{21,22}.

Table 1. Model comparison relative to the null model for deviance in the working memory task

Models	BF ₁₀	
	Experiment 1	Experiment 2
Condition+ Set size	24875.5	5.5e+12
Condition+ Set size+ Condition* Set size	8023.0	2.7e+12
Set size	48.3	5.8e +9
Condition	263.2	346.2

Table 2. Effects analysis for deviance in the working memory task

Effects	BF _{inclusion}		P _{value}	
	Experiment 1	Experiment 2	Experiment 1	Experiment 2
Condition	448.2	939.4	0.003	7.4e-6
Set size	83.2	1.6e+10	0.007	2.8e-6
Condition* Set size	1.3	1.9	0.116	0.035

Reaction times

Statistical analyses suggest that RTs varied as a function of demand and task condition; participants were responding faster on trials that presented fewer squares (i.e. lower set size) and in the ignore (versus update) condition (Figure 2C&D; Supplemental Table 2 for descriptive statistics). In the first experiment, Bayesian model comparison (Table 3) showed that the best model was the one including condition, set size and the interaction between the two (BF₁₀ = 4.8e+31, ~1.4 times better than the one also including the interaction). Effects analyses (Table 4) confirmed that participants were faster on ignore compared with update trials ($F_{1,26} = 16.436$, $p = 4.1e-4$, BF_{INC} = 44.6), a very strong set size effect ($F_{3,78} = 64.739$, $p = 4.1e-21$, BF_{INC} = ∞) and an interaction effect ($F_{2,41,62.7} = 5.643$, $p = 0.003$, BF_{INC} = 26). In the second experiment, the main effects were in the same direction (set size: $F_{2,4,141} = 90.386$, $p = 1.6e-28$, BF_{INC} = ∞, condition: $F_{1,60} = 16.179$, $p = 1.6e-$

4, $BF_{INC} = 88$), but the evidence for an interaction was weaker ($F_{2.7,165} = 4.405$, $p = 0.007$, $BF_{INC} = 3.8$). The model that only involves condition and set size was marginally better than the one including the interaction. Thus, the dependence of the set size effect on task demand is not clear.

Table 3. Model comparison relative to the null model for RTs in the working memory task

Models	BF ₁₀	
	Experiment 1	Experiment 2
Condition+ Set size	6.7e+30	1.4e+53
Condition+ Set size+ Condition* Set size	4.8e+31	1.3e+53
Set size	8.2e+29	2.0e+51
Condition	0.84	3.5

Table 4. Effects analysis for RTs in the working memory task

Effects	BF inclusion		P _{value}	
	Experiment 1	Experiment 2	Experiment 1	Experiment 2
Condition	44.6	87.8	4.1e-4	1.6e-4
Set size	∞	∞	4.1e-21	1.6e-28
Condition* Set size	25.5	3.8	0.003	0.007

Cognitive effort discounting: To repeat or to avoid?

Next, we quantified the subjective cost participants assigned to performing the update and ignore trials. The design of this task was inspired by the temporal and cognitive effort discounting literature^{3,23}. To assess subjective value, participants made choices about repeating a level of the color wheel task for a monetary reward (effort option) or not repeating it for a usually smaller reward (no effort option) (Figure 1B). We decided to contrast the task offer against a break (instead of a lower load offer) as this reflects real- life choices more closely and incurs higher opportunity costs. The task offer was fixed at €2 and the “no effort” offer varied from €0.1 to €2.2. Every choice was sampled three times to account for response variability. Participants were instructed that after all choices were completed, one of them would be randomly selected and they would repeat a few blocks of that set size and mostly that condition (to reduce predictability). If the “no effort” option was selected, they were instructed that they should remain in the testing room for the same amount of time, but they could use their time as they pleased, e.g. make use of their phone or lab computer. They were also informed that receiving the monetary reward would not be contingent on their performance, as long as they put effort into doing the task.

We computed participants’ indifference points (IPs) to estimate subjective value. Indifference points reflect the monetary amount offered for the presumably less effortful option at which participants are equally likely to choose one or the other, thus the probability of accepting either option would be 0.5. We calculated the probabilities of accepting the presumably easier offer using binomial logistic regression analysis.

Figures 3A&B depict the logistic regression curves of an example participant for whom it was possible to estimate indifference points (the participant selected both the task and no effort options enough times to fit a logistic regression, see Methods section) for both update (A) and ignore (B) conditions. The indifference point (IP) represents the degree of discounting of the high-effort offer, where an IP of 2 corresponds to subjective equivalence (given that the offer of the discounted task was always €2, IP = €2 implies that the participant finds the task and the no effort option equally costly). IPs smaller than €2 represent greater discounting (the participant finds the task option to be more costly than the no-redo option) and thus reduced subjective value.

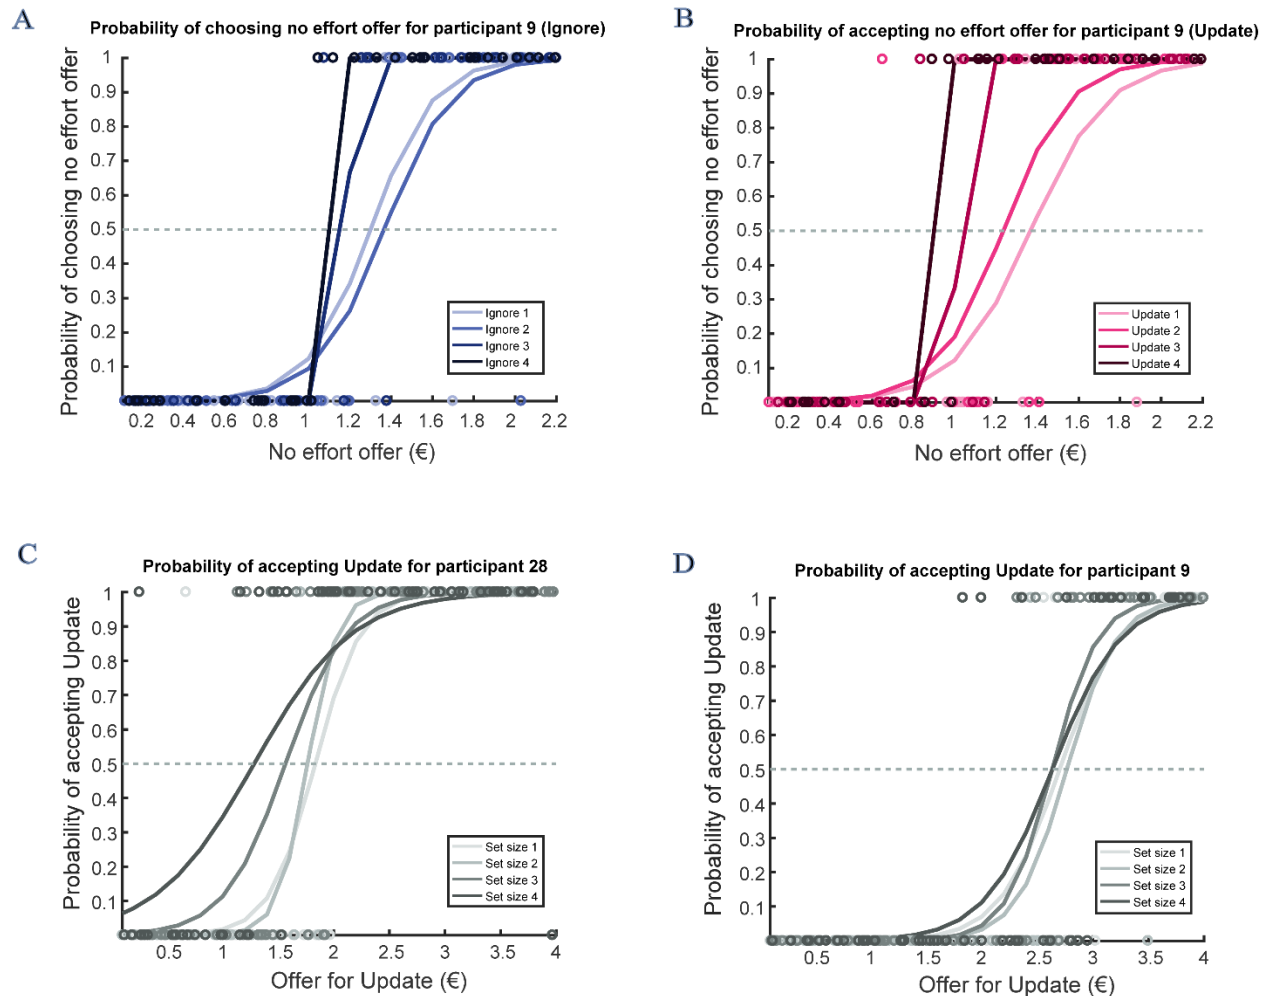


Figure 3. Example participant logistic regression curves. **A, B** Logistic regression curves for “task vs no effort” choices of one participant for update (A) and ignore (B) condition. The probability of accepting the “no effort” (i.e. no task) offer (y-axis) is plotted as a function of the amount of money offered for “no effort” (x-axis). Task offer was always €2 for both conditions and all set sizes and the “no effort” offer varied from €0.10 to €2.20. The estimated indifference point is the offer for “no effort” where the probability of choosing to do the task or the “no effort” option is equal (i.e. 0.5). Indifference points decreased with increasing set size. **C, D** Example logistic regression curves for “ignore vs update” indifference points. The probability of choosing the update offer is shown to vary as a function of the amount of money offered for update. Ignore offer was always €2 for all set sizes, while the update offer varied from €0.10 to €4. The indifference point is the update offer for which the acceptance probability is 0.5, i.e. subjective equivalence. **C:** Example participant who discounted rewards in order to avoid ignore trials (preference for update). **D:** Example

participant who discounted rewards in order to avoid the more demanding levels of update trials (preference for ignore).

Next, we analyzed IPs using Bayesian and classical 2x4 repeated measures ANOVAs to assess whether the subjective value of an offer decreased with demand. Indifference points are displayed as a function of set size and experiment in Figure 4A&B (Supplemental Table 3 for descriptive statistics). Overall, the results show that participants found higher demands of the task more costly as the subjective value decreased with increasing task difficulty (i.e. set size). Moreover, in line with our hypothesis, the subjective cost of performing the ignore condition is higher than that of the update condition (Figure 4A to D). On average, participants found the no effort option less costly than the task option, for both conditions. Analyses of data from Experiment 1 (Table 5) showed that the winning model, which included set size and condition ($BF_{10} = 5006$) was four times more likely than the runner-up model which included set size alone ($BF_{10} = 1229$). Our second experiment replicated this finding, with the same winning model ($BF_{10} = 9.7e+19$) being ~19 times more likely than the runner up (Table 5). Individual effects analyses (Table 6) strengthened these model comparison-based inferences: they provide very strong evidence for a set size effect (Experiment 1: $F_{1,3,31} = 5.666$, $p = 0.016$, $BF_{INC} = 1246$, Experiment 2: $F_{1,57,77} = 22.230$, $p = 2.8e-7$, $BF_{INC} = 6.0e+15$), indicating that participants find higher set sizes to be increasingly costly. In Experiment 1, there was anecdotal evidence that the subjective value of the ignore condition was lower than that of the update condition ($F_{1,23} = 10.924$, $p = 0.003$, $BF_{10} = 3.1$). The more powerful replication study showed extreme evidence for a lower subjective value of ignore versus update ($F_{1,49} = 18.216$, $p = 9.0e-5$, $BF_{10} = 1684$), indicating that participants found the ignore condition subjectively more costly than the update condition. Finally, there is limited evidence against an interaction effect (Experiment 1: $F_{3,69} = 1.798$, $p = 0.168$, $BF_{INC} = 0.2$, Experiment 2: $F_{2,66,130} = 2.167$, $p = 0.102$, $BF_{10} = 0.5$).

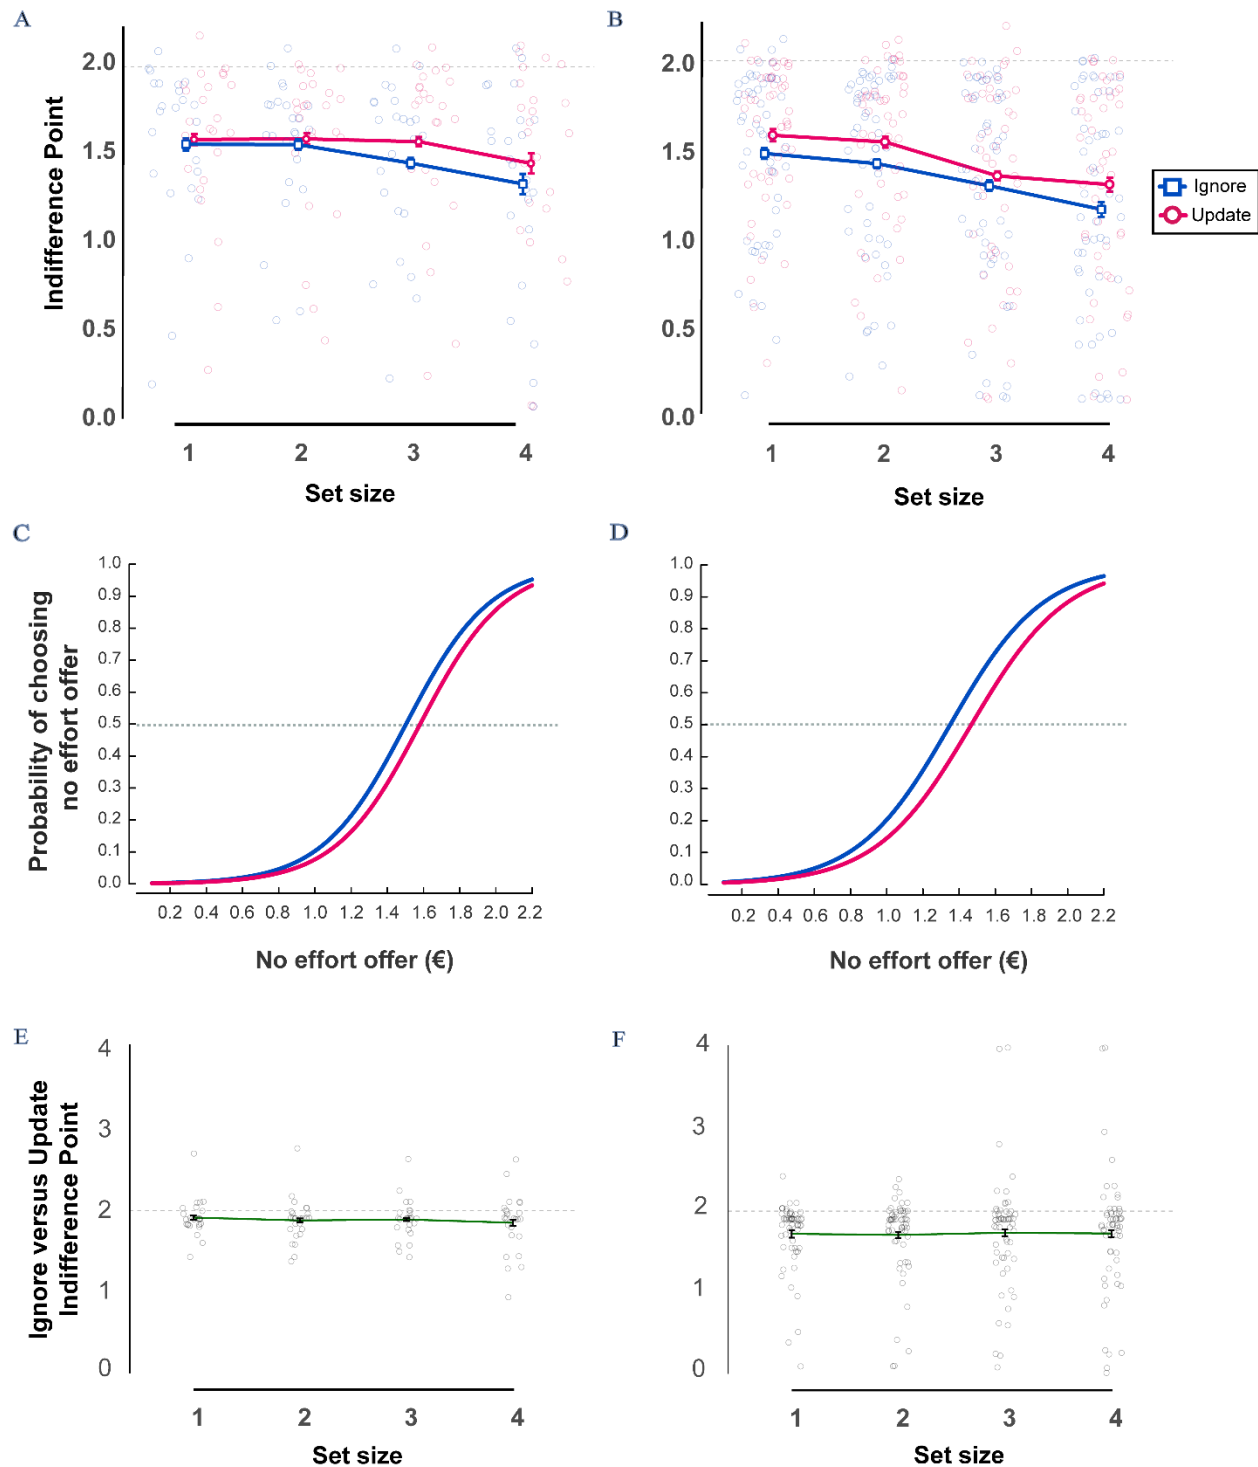


Figure 4. **A, B** “Task vs no effort” indifference points as a function of set size. **A** Experiment 1 (24 participants). **B** Experiment 2 (50 participants). The more the indifference points deviate from 2, the more participants discounted the task option (the task offer was fixed at €2). **C, D** Logistic regression curves for “task vs no effort” choices per condition across set size. The probability of accepting the “no effort” (i.e. no task) offer (y-axis) is plotted as a function of the amount of money

offered for “no effort” (x-axis). Logistic curves estimated using the lme4 package²⁴ in R²⁵. **C** Experiment 1 (24 participants). **D** Experiment 2 (50 participants). **E, F** Indifference points for “ignore versus update” choices as a function of set size. **E** Data from experiment 1 (26 participants). **F** Data from experiment 2 (58 participants). Indifference points smaller than 2 indicate a preference for update over ignore (offer for ignore was fixed at €2). Error bars indicate within-participant SEM^{21,22}.

Table 5. Model Comparison relative to the null model for “task vs no effort” indifference points

Models	BF ₁₀	
	Experiment 1	Experiment 2
Condition+ Set size	5002.7	9.7e+19
Condition+ Set size+ Condition* Set size	722.8	5.2e+18
Set size	1229.2	4.0e+16
Condition	2.7	243.7

Table 6. Effects analysis for “task vs no effort” indifference points

Effects	BF _{inclusion}		P _{value}	
	Experiment 1	Experiment 2	Experiment 1	Experiment 2
Condition	3.1	1683.9	0.03	9.0e-3
Set size	1241.0	6.0e+15	0.016	2.8e-7
Condition* Set size	0.5	0.2	0.168	0.102

Cognitive effort discounting: To Ignore or to Update?

Next, we assessed choices that involved direct comparison between performing the ignore and the update trials. The offer for ignore was fixed at €2 and the offer for update varied from €0.1 to €4. Accordingly, an IP < 2 indicates a preference for (increased subjective value of) update vs ignore, while an IP > 2 represents a preference for ignore vs update (see Methods section for more

details). Figures 3C&D depict logistic regression curves of two example participants, one preferring the update condition and exhibiting an effect of set size (left panel) and the other preferring the ignore condition and not exhibiting an effect of set size (right panel).

Descriptive statistics are presented in Supplemental Table 4 and one-sample t-test output in Table 7. In Figures 4E&F, we report the average indifference points per set size. In accordance with our second hypothesis, the overall average subjective value of ignore versus update choices was less than 2 (1.88), indicating a preference for update over ignore. The support in the data for this hypothesis is ~4.8 times higher than the null (T-test (IP<2) $t_{25} = -2.440$, $p = 0.011$, $BF_{-0} = 4.8$). In Experiment 2, the average subjective value was 1.73 and a preference for update over ignore was ~65 times more supported by the data than no preference (T-test (IP<2): $t_{57} = -3.535$, $p = 4.1e-6$, $BF_{-0} = 65$). The output of the one-way repeated-measures ANOVA shows very strong evidence for the data under the null hypothesis that subjective value is not influenced by set size (Experiment 1: $F_{1,8,45} = 0.961$ $p = 0.382$, $BF_{10} = 0.149$, Experiment 2: $F_{1,2,69} = 0.069$, $p = 0.840$, $BF_{10} = 0.023$). Our results provide confidence in our second hypothesis that participants discount rewards in order to repeat flexible updating trials over distractor resistance and this does not vary with set size.

Table 7. One-sample t-test results that "Ignore vs Update" IPs are smaller than 2

Set size	BF ₋₀		P _{value}	
	Experiment 1	Experiment 2	Experiment 1	Experiment 2
Across	4.8	64.9	0.011	4.1e-4
1	2.4	6698.0	0.026	2.6e-6
2	4.3	1105.0	0.013	1.8e-5
3	4.2	13.2	0.013	0.002
4	3.0	10.9	0.02	0.003

Having established that ignore is both more difficult and perceived as more costly for most participants, we next asked whether variability in preference for update varies with variability in task performance. Plotting deviance for ignore versus update against preference for ignore

versus update reveals little correlation ($r=-0.079$, $BF_{10}=0.181$, $p=0.503$, Figure 5 & Supplemental Figure 3). We also assessed a relationship between preference and performance using mixed effects logistic regression (see methods). We compared the models with and without the main effect of performance (deviance). For both experiments, adding deviance did not improve model fit significantly (Experiment 1: model without deviance: BIC: 4474.3, AIC: 4376.5; full model: BIC: 4482.5, AIC: 4377.9, $p(\text{pr}>\text{Chisq}) = 0.430$; Experiment 2: model without deviance: BIC: 10535, AIC: 10426; full model: BIC: 10544, AIC: 10427, $p(\text{pr}>\text{Chisq}) = 0.463$). Additionally, in the full model, which includes deviance, the effect of condition is still present in Experiment 2 ($p = 0.061$ for Experiment 1; $p = 0.0001$ for Experiment 2). The above suggest that variability in performance does not explain away differences in preference for update versus ignore.

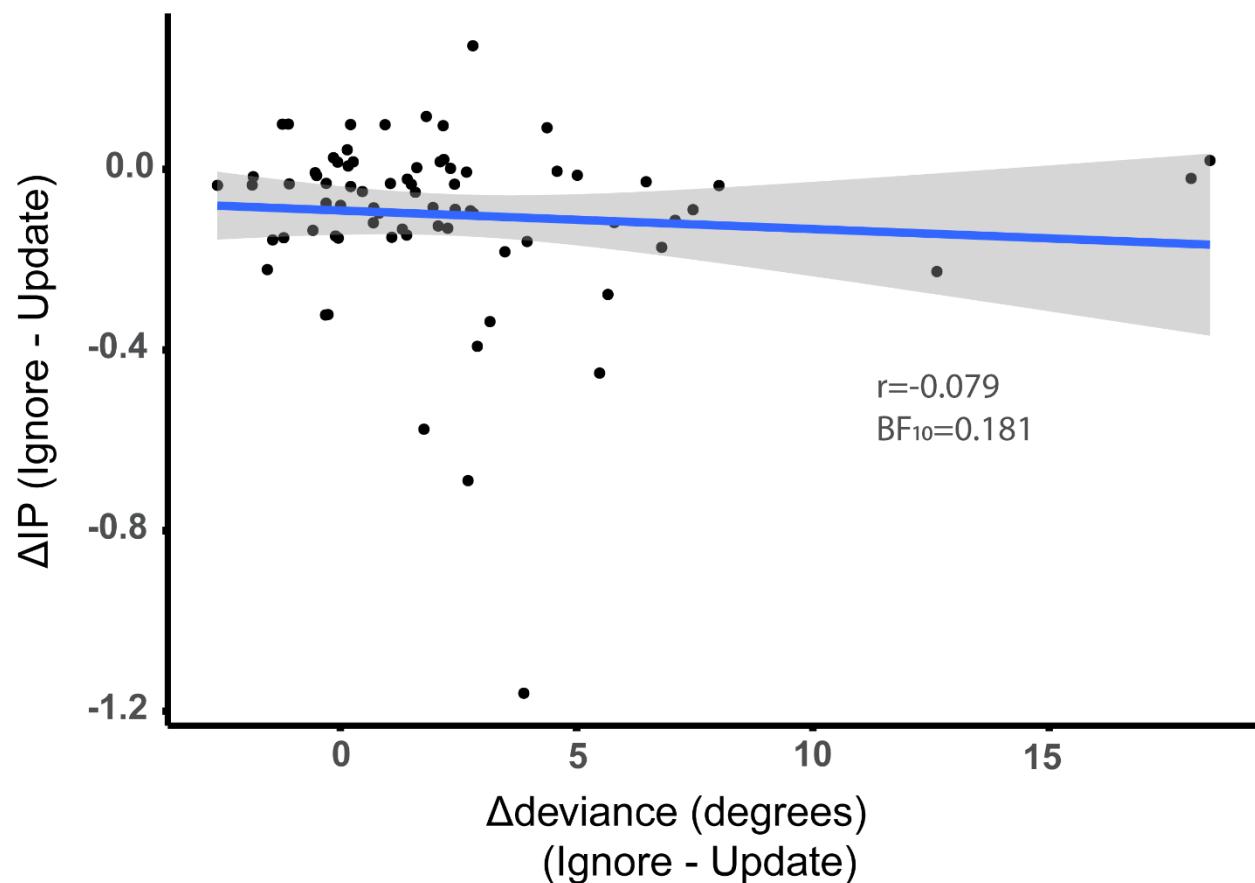


Figure 5. Correlation between condition differences in performance and preference (IP). Ignore minus update performance does not covary with ignore minus update IP. Depicted data pooled from both Experiments (74 participants).

Discussion

In this project, we set out to quantify the subjective value of cognitive stability and cognitive flexibility in the domain of working memory. We asked not only whether these working memory processes are associated with higher subjective costs when demand increases, but also whether tasks requiring cognitive flexibility carry a lower subjective cost than do tasks requiring cognitive stability. In keeping with prior work^{2,26–28}, we demonstrate highly robust and monotonic discounting of delayed response task value with parametrically increasing working memory load (i.e. set size). Most critically, the results provide strong evidence that the ignore version of the task with high stability demands is more costly than is the update version of the task with high flexibility demands: participants are willing to forgo higher monetary offers in order to avoid repeating performing ignore compared with update trials. This finding is evident both indirectly when participants had to choose between the task and a break, and also directly when they had to choose between ignore and update. This result was replicated in the second independent sample and concurs with our primary prediction that the cognitive effort cost of cognitive stability is higher than that of cognitive flexibility.

Depending on one's perspective, this effect of condition on effort cost might be very intuitive or surprising. We might be surprised, because the update trials were longer, and required encoding and gating into working memory twice the number of stimuli compared with the ignore trials. Moreover, many studies have shown that tasks with high demands for cognitive flexibility, like task switching and set-shifting, are accompanied by robust (residual) costs^{29,30}. However, the effect might be considered intuitive, if we recognize that reorienting to salient stimuli can be considered a bottom-up process. In this task, updating is a relatively automatic process, while ignoring requires the withholding of intervening stimuli and thus resolution of conflict, that is,

the core function of cognitive control^{10,31–33}. This then brings us back to the original question: What makes cognitive control costly?

One possibility is that this effect reflects a difference in opportunity costs. In our task, the more subjectively costly ignore trials were 4 seconds shorter than were the cheaper update trials, thus opportunity costs are unlikely to map directly to time costs³⁴. However, we speculate that the effect of task demands on subjective cost reflects an opportunity cost of focusing: the cognitive strategy required for accurate ignore versus update performance differs in the degree to which it allows novel input and thus, alternative opportunities, to impinge on current processing. More generally, it is possible that the brain is more strongly biased against tasks that demand stable focusing compared with flexible opening given that focusing will incur higher opportunity costs across environments.

The observation that the subjective cost of repeating ignore is higher than that of update is in line with the finding that participants perform more poorly on ignore compared with update trials. This finding concurs with previous results from studies using an analogous task with ignore and update conditions^{19,35,36}. In those prior studies, however, the task-relevant delay between the to-be-remembered items and the probe was shorter in the update than the ignore condition, rendering inference about the cognitive mechanism underlying the performance difference difficult. Here we show that the ignore condition is accompanied by worse performance than the update condition, even if task-relevant delay is matched between conditions.

A key question that is raised by the performance difference between task conditions is whether the condition effect on subjective effort cost reflects differences in the degree of (aversion to) anticipated performance error. We argue, however, that an increase in the anticipated - performance error is unlikely to account fully for the increase in subjective effort cost of the ignore versus update condition, for the following three reasons. First, while instructing participants, we highlighted that monetary rewards would not be contingent on performance during the 'redo session', so that performance error should not have influenced participants' choices in our design. Second, in a statistical mixed-effects model that took into account

accuracy, the effect of condition was still present, as a trend in Experiment 1 and highly significant in the more powerful Experiment 2. Third, there was no evidence for a clear association between performance error and measured preference (Figure 5). In future studies, we might consider matching performance between the two conditions or provide “fake” feedback to influence participants’ beliefs about their performance.

Notably, participants responded not only more accurately, but also more slowly on the update than the ignore trials. We are puzzled by this finding, and consider it possible that the slowing reflects a reduction in a nonspecific orienting response to the intervening stimulus, which might have acted as a warning signal. Warning signals are known to induce slowing of responses as a function of foreperiod (delay), which in our case is longer in update trials³⁷. This hypothesis is supported indirectly by the observation that reaction times in flexible update conditions in previous studies in which cue delays were matched between conditions were indeed faster than in ignore conditions^{19,35}. We also consider an alternative explanation, namely that the effect of condition on reaction times reflects a modulation of a decision threshold rather than of attentional orienting, trading time for higher accuracy³⁸ in the update condition in which the memory is more robustly maintained and such a strategy would be beneficial. Here we should note that in both experiments, the time of the mouse click was used as an index of reaction time. However, a clearer picture could be formed if we also had data on initial response times (mouse move) and decision times (move to click). This is a limitation that should be addressed in future studies.

In addition to disentangling the subjective value of distracter resistance and flexible updating task performance, the present results strengthen and extend previous studies on the value of cognitive engagement. First, we confirm that, on average, people are averse to cognitive demand, are ‘cognitive misers’, even willing to decline rewards in order to avoid demanding tasks. This strengthens earlier work showing that participants prefer to avoid more demanding N-back tasks³, detection tasks²⁷ or sustained attention tasks². Our results further generalize these conclusions to the most classic of working memory tasks: the delayed response task. A distinct strength of our design is the fact that our implementation of the discounting procedure

takes into account the observation that choices are probabilistic³⁹. Unlike prior studies on cognitive effort which used staircase procedures sampling every choice option only once^{2,3}, we sampled the full discounting curve and every choice option multiple times. Furthermore, unlike prior studies, in which on the first trial a lower monetary offer was made for the low effort option than for the high effort option, we avoided (potential) anchor effects by presenting offers randomly. Finally, unlike previous discounting studies we also gave participants the opportunity to choose the effortful option for less money. As expected, most participants declined this offer, but the subjective value of four participants (total in both samples) was higher than 2 for at least one of the two working memory processes, indicating a preference for repeating the working memory task, suggestive of effort seeking⁴⁰.

The neural mechanisms underlying the considerable individual variation in subjective cognitive effort costs should be addressed in future work. Past work has shown that administration of the catecholamine reuptake blocker methylphenidate improves distracter resistance at the expense of flexible updating on a task analogous to the one employed here²⁰. However we also know that effects of psychostimulants vary greatly with individual baseline measures of dopamine⁴¹. How does a preference for ignore versus update relate to baseline measures of dopamine and psychostimulant effects on cognition? A role for dopamine in effort-based decision-making would be consistent with studies in physical effort, where it has been shown that in Parkinson's patients dopamine medication increases selection of high effort/high reward trials⁴², while dopamine depletion decreases willingness to exert effort in humans and rodents^{43,44}. Another potentially relevant neurotransmitter is noradrenaline that seems to be involved in switching modes between task engagement (focusing) and disengagement (distractibility)^{45,46}. Indeed, there is recent evidence showing that amphetamine and methylphenidate, both altering catecholamine transmission, modify cognitive demand avoidance in rodents⁴⁷ and humans⁴⁸ respectively. These findings together suggest that catecholamines contribute to valuation of cognitive effort⁴⁹ and we believe our paradigm is suited to aid uncovering such effects in future research.

In conclusion, this study provides new insights to the novel and growing fields of cognitive effort discounting and value-based decision-making. Specifically, we showed that with increasing demand on working memory processes, the subjective valuation decreased, both for the process of distractor resistance and flexible updating. We also show strong evidence that distractor resistance is perceived as relatively costlier than flexible updating.

Methods

Participants

For Experiment 1, 32 participants (22 women), aged between 18-29 years old were tested in total. Participants had normal or corrected-to-normal vision. Colorblind participants were excluded. Four data sets were lost during data transfer, so we ended up with 28 data sets (20 women, 18-33 years old, mean: 24). For Experiment 2, we sought to replicate the finding that update is more costly than ignore. We performed a sequential sampling power calculation using the BFDA⁵⁰ package in R. We set the minimum sample size to 60⁵¹, the maximum to 100 and set the boundary of sampling at a Bayes Factor of 10 for either the null or the alternative hypothesis given the effect size estimated from Experiment 1. We collected 62 data sets (37 women, 20-44 years old, mean: 25.6, standard deviation: 4.3), at which sample size the boundary was already reached. The study was approved by the local ethics committee (CMO region Arnhem/Nijmegen, The Netherlands, CMO2001/095) and all participants provided written informed consent, according to the declaration of Helsinki.

Exclusion criteria

We excluded participants based on four rules: 1) Failing to pass the color sensitivity test twice. 2) Striking evidence that they did not understand or will to perform the tasks. 3) Mean deviation exceeding 3 standard deviations from average for at least one of our main conditions (across demand) of the working memory task. 4) Failing to estimate reliable indifference points for at least one condition (across demand levels) of the effort discounting tasks.

Based on our criteria, one outlier was excluded from performance analysis of Experiment 1 for deviating more than 3 standard deviations from the mean for ignore ($\sim 3SD$) and one from Experiment 2 for deviating more than 3 standard deviations from the mean of both conditions ($\sim 5.4SD$ from ignore and $\sim 6.6SD$ from update mean). Four people were excluded from the analysis of task vs no effort indifference points in Experiment 1 and twelve in Experiment 2. In Experiment 1, all four were excluded because we could not estimate indifference points for at least one of the two conditions (ignore/update). Among the four that were excluded, one always chose the no effort option, one of them always chose the task option and one of them always chose no effort for update trials and task for ignore trials. In Experiment 2, eleven participants were excluded due to inadequate response variability and one because he was not performing the task. Out of the eleven whose IPs we could not estimate, one almost always chose the task option and the rest always preferred the no effort option. The other participant always responded using one of the two response buttons. This is a clear indication that he was not trying to perform the task because easy and hard offer presentation was counterbalanced across response buttons. We excluded two participants from the analysis of “ignore vs update” indifference points in Experiment 1 analysis; one because we could not estimate any indifference points (always chose ignore) and another because they deviated more than 3 standard deviations from the mean. Four participants were excluded in the replication for the same analysis. One always chose ignore, two always chose update and one did not do the task (see above).

Task design

All paradigms were entirely programmed in MATLAB (Mathworks, Natick, MA, USA)(release 2013a) using the Psychophysics Toolbox extension⁵² (version 3.0.12) on a Windows 7 operating system. The screen resolution was 1920x1080 pixels. The background color for all paradigms was grey (R: 200 G:200 B:200).

The experiment lasted about 130 minutes and consisted of four tasks performed at a computer and questionnaires that participants filled in at the end. The first task (~ 7 min) was a color sensitivity test aiming to check whether participants were sensitive to the colorful stimuli used in

the memory task. Participants then proceeded with the color wheel working memory task to acquire experience with varying demand of the two working memory processes of interest (~10min practice and 30min task). The third task (~5min practice and 55min task) was a cognitive effort discounting paradigm that was used to estimate subjective value and address our research questions. The last computer-based task was a redo of the color wheel task (~10min). Finally, participants filled in some experiment-related questionnaires (~5min).

Color sensitivity task

For our working memory task, we used color stimuli and a color wheel, so it was crucial that our participants' color vision was not impaired. To test their sensitivity to our manipulation we developed a version of the color wheel task without a working memory component. In this task participants viewed a colored square in the middle of the screen and the same color wheel used in the memory task. Their goal was to click on the color of the wheel that matched the colored square.

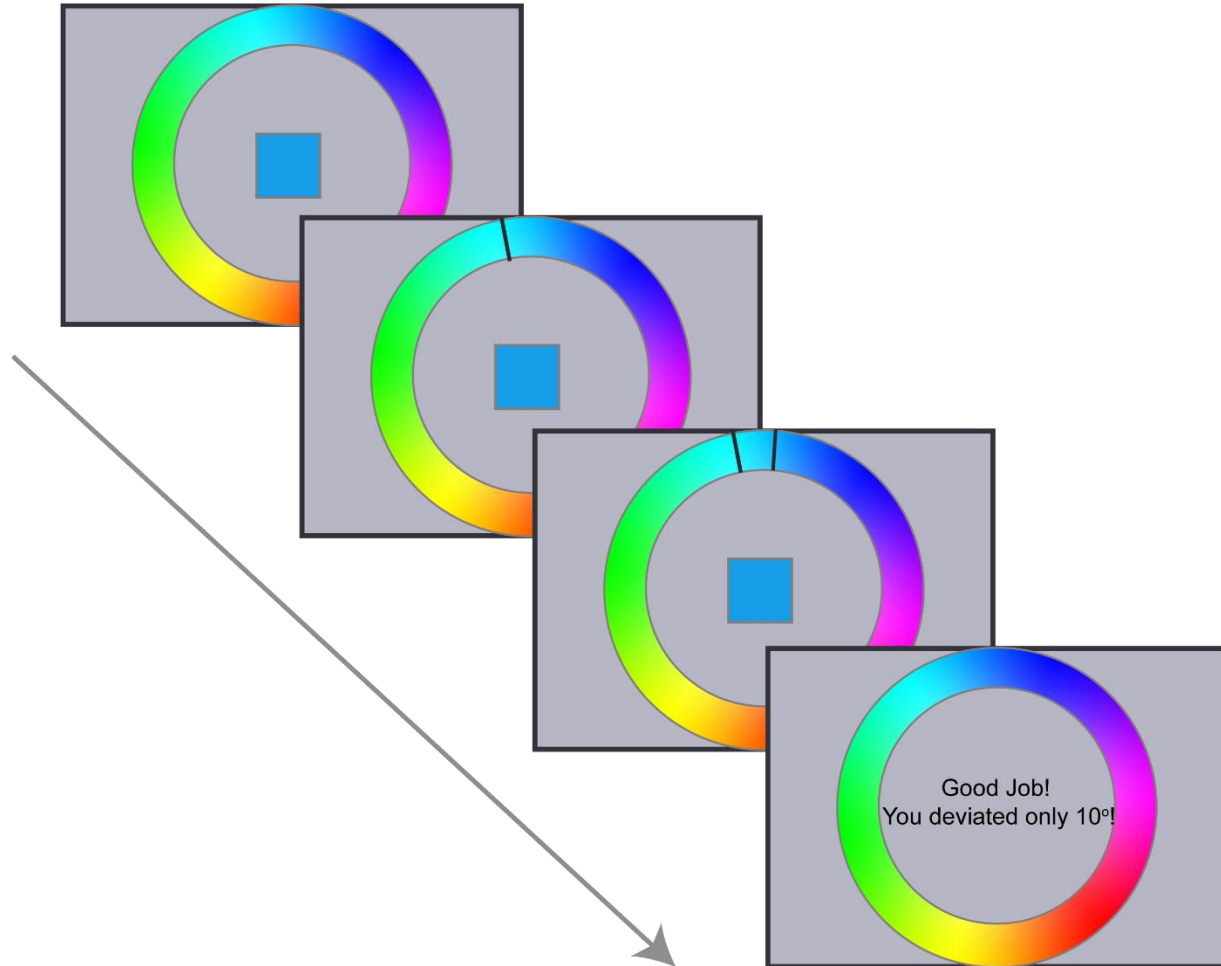


Figure 7. Example trial of the color sensitivity task. Participants view a colored square in the middle of the screen and they have to click with a mouse on the corresponding color on the color wheel. A black line indicates the selected color and a successive line the correct color. If the selected color deviated 10° or less from the correct color they receive feedback that they performed well. The task is self-paced (24 trials total). They successfully complete the task if their average deviance is less than 15° .

The stimuli used for the color sensitivity task were a color wheel, black lines and colored squares. The color wheel was created by 512 successive colored arcs of equal angle ($512/360^\circ = 1.42^\circ$), each arc carrying a different color. The radius of the wheel was 486 pixels. To form the wheel into a ring, a smaller circle was superimposed, whose radius was ~ 362 pixels. The centre of both the

wheel and the circle coincided with the centre of the screen. The 512 colors of the color wheel arcs were generated using the hsv MATLAB colormap. The black lines were 0.4° black arcs.

In every trial of this task, participants viewed the color wheel and a colored square in the middle of the screen (Figure 7). They were instructed to look at the color of the square and use the mouse to click on the corresponding shade on the color wheel. To indicate that their response was recorded a black line appeared on the color wheel and successively another black line appeared designating the location of the correct color. Feedback consisted of the actual deviance plus a positive message ('Good job! You deviated only ___ degrees.') and was provided only when responses deviated less than 10° .

To test a representative sample of the color wheel we split the wheel in 12 main arcs. Participants were tested in two different shades from each of the 12 color categories (arcs). So, they performed in total 24 trials of this task. The presentation of the trials as well as the orientation of the color wheel were randomized. The responses were self-paced and total task duration was approximately 7min.

The main dependent variable in this task was deviance in degrees from the correct color. If their average deviance was less than 15° by the end of the task, the experiment continued. Otherwise, they had one more chance to perform the color sensitivity task, but if failed again they would be excluded.

All participants from both experiments completed 24 trials of the color sensitivity task and they all met the criterion (average deviance from correct color below 15 degrees) to continue to the main paradigm. For Experiment 1, the average deviance from the target color was 6.63 degrees (SD = 1.23; median = 4.72, SD = 0.85) and for replication Experiment 2 mean deviance was 6.27 (SD=1.4; median=4.85, SD=1.08). We also reported the median for easy comparison with the color wheel working memory task results.

Color wheel working memory task

After successfully completing the colour sensitivity task, participants proceeded with the colour wheel working memory task. In this part, participants experienced varying demands of cognitive stability and cognitive flexibility. This task was based on a short-recall task¹⁸ and delayed-match-to-sample tasks¹⁹ that have previously been used to disentangle between the two working memory processes of interest.

The stimuli displayed during this paradigm were a color wheel, colored squares, black frames of squares, a fixation cross, black lines and central letter cues. The color wheel was generated as described in the color sensitivity section. The number of squares varied from one to four and they could be located in four different positions. The centres of the squares formed a rectangle with dimensions 248*384 pixels. Each of the four squares was 100x100 pixels in size. To choose the colors of the squares, we split the color wheel into 12 main arcs of 42 colors each and only used the 15 central colors of each arc. The arcs from which the colors would be sampled per trial were defined manually, but the exact shade (RGB values) was randomly selected. The letter cues were “I” and “U”, colored black and presented at the centre of the screen.

Every trial of the task consisted of three phases separated by two delay periods (Figure 1). During the **encoding** phase, participants viewed the fixation cross and one to four colored squares for two seconds. The number of squares displayed (set size 1-4) represented the demand of the trial. A delay of two seconds followed, during which only the fixation cross was displayed. Then the **interference** phase followed. In this phase, participants viewed the same number of squares as during encoding, at the same locations, but with different colors. Instead of a fixation cross, one of the two letter cues was presented during interference in the middle of the screen. The cue indicated the condition of the trial: “I” for ignore trials and “U” for update. The second delay duration depended on trial condition, and was two seconds for ignore and six seconds for update trials. Finally, during the **response** phase participants saw black frames of the same squares, one of which was highlighted, in addition to the color wheel and the fixation cross. If the participant responded within four seconds, a black line appeared on the color wheel, otherwise, they were instructed to respond faster (‘Please respond faster!’). The total duration of the response phase

was five seconds.

For the encoding phase, participants were instructed to always memorize the colors and locations of all presented squares. The instructions for the interference phase differed based on the condition as indicated by the letter cue. In ignore trials, participants needed to maintain in their memory the colors from the encoding phase and not be distracted by the new intervening stimuli. In flexible updating trials, participants had to let go of their previous representations and update into their memory the stimuli from the interference phase. Thus, the colors that needed to be remembered for the ignore condition were the ones from the encoding phase, while for the update condition they were the ones from the interference phase. To match the time that the relevant stimuli were maintained in memory for both conditions, the second delay was 4 seconds longer for update trials. Participants were to indicate the color for only the highlighted square. They had to identify the target color on the color wheel and click using a mouse, within four seconds. Only the first response counted. A black line indicated their response. Only during practice trials, a second line appeared at the correct color and positive feedback was displayed if they were performing well (as in color sensitivity section). During the task, no feedback was provided. We instructed participants to fixate in the middle of the screen throughout the task in order to dissuade them from adopting the strategy of closing their eyes during ignore trials in order to avoid being distracted.

Participants first underwent a practice session of 16 trials and then performed two blocks of the task. A block consisted of 64 trials, resulting from repeating each combination of difficulty (four levels: set size 1 – 4) and condition (two levels: ignore and update) eight times. Depending on the difficulty level of the trial, a group of two to eight colors was used to create the trial stimuli, each color coming from one of the 12 arcs. Colors of the same arc never appeared more than once in the same trial. To make sure that ignore and update trials were as similar and counterbalanced as possible, the color stimuli sets displayed and the target colors were the same for both conditions. Because the relevant colors appeared during encoding for ignore and during interference for update, we made sure that the same group of stimuli also appeared in reverse order between these two phases. So, the same groups of colored squares were presented four times per set size and in total 32 groups of colors were used. To decrease learning effects due to repetition, we split

the same stimuli groups between the two blocks. To control for differences between the two hemispheres in representation of color⁵³, target locations (left/right) were counterbalanced across conditions. Moreover, the same colors were highlighted for all four set sizes.

Cognitive effort discounting task

After participants gained experience with all four difficulty levels of update and ignore conditions of the color wheel working memory task, they proceeded with the third part of the experiment: the effort discounting task (Figure 1B). The aim of this paradigm was to quantify the subjective value that participants assigned to the color wheel task performance. There were two versions of choice trials to address our two research questions. In both versions, two options were accompanied by an amount of money and the options defined what participants would do in the last part of the experiment.

In every trial of the task, participants saw a rectangle containing two options and a fixation cross. The options could be “No Redo” or any set size of ignore or update, for example “Ignore 2”, corresponding to the ignore condition of the task and set size of 2. Below each option, a monetary reward was displayed, for example “for 2€”. Participants could choose the left or right option by pressing 1 or 2 on the keyboard and they had six seconds to respond. When participants made a choice, a black square surrounding the selected offer appeared to indicate that their response was recorded.

At this stage, participants were instructed that there were two more parts in the experiment. In the last part, they would have the opportunity to earn a bonus monetary reward by redoing one to three blocks of the color wheel task. However, the amount of the bonus and the type of trials they would repeat would be based on the choices they made on the choice task. To highlight the importance of every choice, we instructed them that of all the choices they made (of both versions) the computer would select only one randomly and the bonus and redo would be based on that single choice. To minimize effects of error avoidance on choices, we informed participants that accuracy during the redo part would not influence whether they receive the monetary reward, as long as their performance was comparable with the first time that they did the color wheel task (part 2 of experiment). Both the rewards and the redo were real and not hypothetical.

Task vs No effort: Choices between working memory task and no task

These trials addressed the first research question: whether the subjective values of ignore and update decrease as a function of task demand. Here, participants had to choose between repeating a level of ignore or update (effort offer) and not redoing the color wheel task at all (no effort offer). If they chose the no effort option (“No Redo”) they were instructed that they would be able to use their time as they pleased (e.g. by using their phones or lab computer) but they would still have to stay in the testing lab so that time spent on the experiment was the same for both options. Otherwise, if the option to repeat the task was selected, the redo trials would consist of mostly the selected choice condition and level. “Mostly” is important because if they always did the same condition during the redo, they would be able to predict whether they had to update or ignore. We emphasized that they should take their time to respond, consider both the money and their experience while doing the color wheel task as well as the importance of choosing their true preference and not try to please us.

Ignore vs Update: Choices between cognitive stability and cognitive flexibility

This COGED trial type aimed to investigate whether ignore is perceived as costlier than flexible updating by directly contrasting them. In these trials, participants had to choose between doing the same level of either ignore or update.

For the ‘task vs no effort’ version of the COGED, the amount offered for the no effort “No Redo” option varied from €0.10 to €2.20 in €0.20 steps (except the first step, which was €0.10), while the task option (effort offer) was always fixed at €2.00. The €2.20 option for “No Redo” was included to identify whether there were participants who strongly preferred performing the task, even if that meant forgoing rewards. As we hypothesized that ignore would be costlier, in the ‘ignore vs update’ version of the task, ignore (hard offer) was kept steady at €2 and update (easy offer) was varying from €0.10 to €4 in €0.20 steps (as above). There were 96 possible pairs for “task vs no effort” choices (12 amounts*2 conditions*4 set sizes) and 84 for “ignore vs update” choices (21 amounts*4 set sizes). Given the evidence that choice is probabilistic rather than deterministic³⁹, every pair of options was sampled three times. We decided on three repetitions

of the pairs based on a simulation analysis using pilot data (Supplemental Figure 4) in order to optimize the trade-off between indifference point estimation and task duration. Each participant performed three blocks that contained in total 288 trials of “task vs no effort” trials and 252 trials of “ignore vs update”. The trials of the two versions were interleaved (mixed) and randomized within each block. To avoid location effects, we counterbalanced the left-right presentation of the two options. Total task duration was about 55 minutes.

We decided to use fixed sets of offers and not a titrated staircase procedure to estimate subjective value because staircase procedures do not sample the entire logistic regression curve. This COGED version allowed us to sample the logistic regression curves adequately because all participants were faced with the entire range of offer options.

Redo

After participants finished three blocks of the discounting choice task, one of their choices was pseudo-randomly selected. Specifically, the computer only sampled from “ignore vs update” choices of level 3 or 4. Participants always did one block of 24 trials of the color wheel task. Two-thirds of these trials were their preferred condition (ignore/update). We decided to never select the no effort option to maintain experimenter credibility, so that participants discussing the task are convinced that the consequences are real. The redo data were not analysed and participants always received the bonus regardless of their performance.

Debriefing questions

After the end of the experiment we requested participants to complete questionnaires. We explicitly asked them to report their preference by asking “Which trials did you prefer?”.

Data analysis

We analysed our data using both frequentist and Bayesian statistics. All statistical analyses were performed using open source software JASP (version 0.7.5.6)^{54,55} on a Windows 7 operating system.

As skepticism against classical statistical tools increases⁵⁶, we turned to Bayesian statistics⁵⁷. This allowed us to quantify evidence for our hypotheses instead of forcing an all-or-none decision and an arbitrary cut-off of significance. Bayesian statistics can also provide evidence for the null hypothesis (H_0), thus distinguishing between undiagnostic data (“absence of evidence”) and data supporting H_0 (“evidence of absence”). Another important benefit is that we are able to monitor evidence as data accumulate and we can continue sampling without biasing the result. Due to all the above advantages, we decided that our main conclusions would be drawn based on the Bayesian analyses.

However, frequentist statistics are well-established and widely-acknowledged tools, so more scientists are familiar with their rationale and interpretation. To ensure that our results are interpretable for all and to allow comparison with earlier work, we additionally included classical statistics. Bayesian statistics allow model comparison, but also provide evidence for individual effects. When possible, we reported Bayesian model comparison (BF_{10} : Bayes factor of model against the null) as well as Bayesian and frequentist effects analyses ($BF_{INC(LUSION)}$: Bayes factor of Bayesian model averaging). We used the default JASP Cauchy priors for all Bayesian statistics⁵⁵. Regarding frequentist statistics, we considered a p-value of 0.05 or smaller as significant. In the cases where sphericity was violated, we reported the Greenhouse-Geisser corrected p-values.

Color sensitivity task analysis

The data from this task were used only to establish that participants are sensitive enough to our color wheel. We calculated the overall average deviance in degrees.

Color wheel task data analysis

We computed the median deviance and median reaction time for all levels of ignore and update. The rationale behind choosing the median was that it is less sensitive to extreme values. For example, 90° and 180° accuracy scores are both wrong responses, but the latter affects the mean much more strongly. We used the above indices for the statistical analysis using classical and Bayesian 2x4 repeated measures ANOVAs with condition (ignore/update) and set size (levels 1-4)

as within-subject factors. All participants in both experiments performed above chance level (mean deviance less than 90°).

Discounting choice task data analysis

As an estimate of subjective value, we computed participants' indifference points. The indifference points can be interpreted as the financial amount offered for the presumably less effortful option (no effort or update) at which participants are equally likely to choose one or the other, thus the probability of accepting either option would be 0.5. With the main dependent variable being choice, a dichotomous variable, we calculated the probabilities of accepting the presumably less effortful offer using binomial logistic regression analysis in MATLAB and extracted the indifference points for the different conditions.

Choices between working memory task and no effort

Having determined the indifference points for all levels of both working memory conditions per participant, we continued with the statistical analysis using classical and Bayesian 2x4 repeated measures ANOVAs to assess our first hypothesis that subjective value decreases with demand for ignore and update. Confirmation of this hypothesis would require that the model including set size is more likely than the null model, or the presence of a set size effect with p-value smaller than 0.05. We also performed Bayesian and classical one sample t-tests on the indifference points across levels for both conditions to assess whether the subjective value of the working memory functions was overall lower than the no task subjective value. The task offer was always €2, so a subjective value lower than 2 would imply that participants were discounting the task option.

Choices between Update and Ignore

We then computed participants' indifference points collapsing across levels of "ignore vs update" choice trials to evaluate our hypothesis that ignore has a lower subjective value than update using Bayesian and classical one sample t-tests. As ignore offer was set at 2€, subjective values lower than 2 indicate that participants were willing to forgo rewards to repeat update instead of ignore

trials. Additionally, we calculated indifference points for all levels separately and used a 1x4 ANOVA with set size as a factor to assess if the preference for update varies with demand.

Mixed effects analyses

We tested for a relationship between preference and performance with mixed effects logistic regression analyses, using the lme4 package²⁴ in R²⁵. In our model, we regressed preference on fixed effects of set size, condition, the money offered for the “no effort” option and deviance (accuracy index). We also included random intercepts and slopes for the effects of easy offer amount, condition and set size per participant. Continuous variables “easy offer amount” and “deviance” were log-transformed and standardized. We also assessed if the effect of condition remained significant after including deviance in the model. For that analysis, deviance was also added as a by- participant random slope. Model fits were compared using likelihood ratio chi-square tests.

To estimate the discounting curves across participants (Figure 3C&D) we used a mixed effects model per condition with offer amount as fixed factor and participant as a random factor.

Bibliography

1. Kool, W., McGuire, J. T., Rosen, Z. B. & Botvinick, M. M. Decision making and the avoidance of cognitive demand. *J. Exp. Psychol. Gen.* **139**, 665–682 (2010).
2. Massar, S. A. A., Lim, J., Sasmita, K. & Chee, M. W. L. Rewards boost sustained attention through higher effort: A value-based decision making approach. *Biol. Psychol.* **120**, 21–27 (2016).
3. Westbrook, A., Kester, D. & Braver, T. S. What is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. *PLoS One* **8**, e68210 (2013).
4. Botvinick, M. & Braver, T. Motivation and Cognitive Control : From Behavior to Neural

- Mechanism. *Annu. Rev. Psychol.* (2015). doi:10.1146/annurev-psych-010814-015044
5. Padmala, S. & Pessoa, L. Reward reduces conflict by enhancing attentional control and biasing visual cortical processing. *J. Cogn. Neurosci.* **23**, 3419–32 (2011).
 6. Kurzban, R., Duckworth, A., Kable, J. W. & Myers, J. An opportunity cost model of subjective effort and task performance. *Behav. Brain Sci.* **36**, 661–79 (2013).
 7. Inzlicht, M., Schmeichel, B. J. & Macrae, C. N. Why self-control seems (but may not be) limited. *Trends Cogn. Sci.* **18**, 127–133 (2014).
 8. Cools, R. The cost of dopamine for dynamic cognitive control. *Curr. Opin. Behav. Sci.* **4**, 152–159 (2015).
 9. Shenhav, A. *et al.* Toward a Rational and Mechanistic Account of Mental Effort. *Annu. Rev. Neurosci.* **40**, 99–124 (2017).
 10. Botvinick, M. M. & Cohen, J. D. The computational and neural basis of cognitive control: Charted territory and new frontiers. *Cogn. Sci.* **38**, 1249–1285 (2014).
 11. Cools, R. The costs and benefits of brain dopamine for cognitive control. *Wiley Interdiscip. Rev. Cogn. Sci.* **7**, 317–329 (2016).
 12. Mooneyham, B. W. & Schooler, J. W. The costs and benefits of mind-wandering: A review. *Can. J. Exp. Psychol. Can. Psychol. expérimentale* **67**, 11–18 (2013).
 13. Baird, B. *et al.* Inspired by Distraction. *Psychol. Sci.* **23**, 1117–1122 (2012).
 14. Ansburg, P. I. & Hill, K. Creative and analytic thinkers differ in their use of attentional resources. *Pers. Individ. Dif.* **34**, 1141–1152 (2003).
 15. Zedelius, C. M. & Schooler, J. W. Mind wandering “Ahas” versus mindful reasoning: alternative routes to creative solutions. *Front. Psychol.* **6**, 834 (2015).
 16. Zedelius, C. M. & Schooler, J. W. The Richness of Inner Experience: Relating Styles of Daydreaming to Creative Processes. *Front. Psychol.* **6**, 2063 (2015).
 17. Musslick, S., Jang, S. J., Shvartsman, M., Shenhav, A. & Cohen, J. D. Constraints associated

- with cognitive control and the stability-flexibility dilemma. (2017).
18. Zhang, W. & Luck, S. J. Discrete fixed-resolution representations in visual working memory. doi:10.1038/nature06860
 19. Fallon, S. J. & Cools, R. Reward acts on the pFC to enhance distractor resistance of working memory representations. *J. Cogn. Neurosci.* **26**, 2812–26 (2014).
 20. Fallon, S. J., van der Schaff, M. E., ten Huurne, N. & Cools, R. Methylphenidate improves cognitive stability at the expense of cognitive flexibility. *submitted* (2015).
 21. Cousineau, D. Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutor. Quant. Methods Psychol.* **1**, 42–45 (2005).
 22. Morey, R. D. Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutor. Quant. Methods Psychol.* **4**, 61–64 (2008).
 23. Kable, J. W. & Glimcher, P. W. The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* **10**, 1625–1633 (2007).
 24. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using **lme4**. *J. Stat. Softw.* **67**, 1–48 (2015).
 25. R Core Team. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>. (2014). Available at: <http://www.r-project.org/>. (Accessed: 30th October 2018)
 26. Westbrook, A., Kester, D. & Braver, T. S. What Is the Subjective Cost of Cognitive Effort? Load, Trait, and Aging Effects Revealed by Economic Preference. *PLoS One* **8**, 1–8 (2013).
 27. Chong, T. T.-J. *et al.* Neurocomputational mechanisms underlying subjective valuation of effort costs. *PLOS Biol.* **15**, e1002598 (2017).
 28. Apps, M. A. J., Grima, L. L., Manohar, S. & Husain, M. The role of cognitive effort in subjective reward devaluation and risky decision-making. *Sci. Rep.* **5**, 16880 (2015).
 29. Rogers, R. D. & Monsell, S. Costs of a predictable switch between simple cognitive tasks. *J.*

- Exp. Psychol. Gen.* **124**, 207–231 (1995).
30. Rubinstein, J. S., Meyer, D. E. & Evans, J. E. Executive control of cognitive processes in task switching. *J. Exp. Psychol. Hum. Percept. Perform.* **27**, 763–797 (2001).
 31. Feng, S. F., Schwemmer, M., Gershman, S. J. & Cohen, J. D. Multitasking vs. multiplexing: Toward a normative account of limitations in the simultaneous execution of control-demanding behaviors. *Cogn. Affect. Behav. Neurosci.* **14**, 129 (2014).
 32. Corbetta, M. & Shulman, G. L. Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* **3**, 201–15 (2002).
 33. Ernst, M., Daniele, T. & Frantz, K. New perspectives on adolescent motivated behavior: attention and conditioning. *Dev. Cogn. Neurosci.* **1**, 377–89 (2011).
 34. Niv, Y., Daw, N. D., Joel, D. & Dayan, P. Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology (Berl)*. **191**, 507–20 (2007).
 35. Fallon, S. J., van der Schaaf, M. E., ter Huurne, N. & Cools, R. The Neurocognitive Cost of Enhancing Cognition with Methylphenidate: Improved Distractor Resistance but Impaired Updating. *J. Cogn. Neurosci.* **29**, 652–663 (2017).
 36. Fallon, S. J., Mattiesing, R. M., Dolfen, N., Manohar, S. G. & Husain, M. Ignoring versus updating in working memory reveal differential roles of attention and feature binding. *Cortex*. **107**, 50–63 (2018).
 37. Loveless, N. E. & Sanford, A. J. The impact of warning signal intensity on reaction time and components of the contingent negative variation. *Biol. Psychol.* **2**, 217–226 (1975).
 38. Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U. & Nieuwenhuis, S. The neural basis of the speed–accuracy tradeoff. *Trends Neurosci.* **33**, 10–16 (2010).
 39. Rieskamp, J. The probabilistic nature of preferential choice. *J. Exp. Psychol. Learn. Mem. Cogn.* **34**, 1446–65 (2008).
 40. Inzlicht, M., Shenhav, A. & Olivola, C. Y. The Effort Paradox: Effort Is Both Costly and

- Valued. *Trends Cogn. Sci.* **22**, 337–349 (2018).
41. Cools, R. & D'Esposito, M. Inverted-U-shaped dopamine actions on human working memory and cognitive control. *Biol. Psychiatry* **69**, e113-25 (2011).
 42. Chong, T. T.-J. *et al.* Dopamine enhances willingness to exert effort for reward in Parkinson's disease. *Cortex* **69**, (2015).
 43. Salamone, J. D., Correa, M., Farrar, A. & Mingote, S. M. Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology (Berl)*. **191**, 461–82 (2007).
 44. Floresco, S. B., Tse, M. T. L. & Ghods-Sharifi, S. Dopaminergic and Glutamatergic Regulation of Effort- and Delay-Based Decision Making. *Neuropsychopharmacology* **33**, 1966–1979 (2008).
 45. Jepma, M., Te Beek, E. T., Wagenmakers, E.-J., van Gerven, J. M. A. & Nieuwenhuis, S. The role of the noradrenergic system in the exploration-exploitation trade-off: a psychopharmacological study. *Front. Hum. Neurosci.* **4**, 170 (2010).
 46. Aston-Jones, G. & Cohen, J. D. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* **28**, 403–50 (2005).
 47. Cocker, P. J., Hosking, J. G., Benoit, J. & Winstanley, C. A. Sensitivity to Cognitive Effort Mediates Psychostimulant Effects on a Novel Rodent Cost/Benefit Decision-Making Task. *Neuropsychopharmacology* **37**, 1825–1837 (2012).
 48. Froböse, M. I. *et al.* Catecholaminergic modulation of the avoidance of cognitive control. *J. Exp. Psychol. Gen.* **147**, 1763–1781 (2018).
 49. Froböse, M. I. & Cools, R. Chemical neuromodulation of cognitive control avoidance. *Curr. Opin. Behav. Sci.* **22**, 121–127 (2018).
 50. Schönbrodt, F. D. BFDA: Bayes factor design analysis package for R. (2016).
doi:<https://github.com/nicebread/BFDA>

51. Simonsohn, U. Small Telescopes. *Psychol. Sci.* **26**, 559–569 (2015).
52. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–6 (1997).
53. Gilbert, A. L., Regier, T., Kay, P. & Ivry, R. B. Whorf hypothesis is supported in the right visual field but not the left. *Proc. Natl. Acad. Sci. U. S. A.* **103**, 489–94 (2006).
54. JASP Team. JASP (version 0.7.5.6) (Computer Program). (2016).
55. Wagenmakers, E. *et al.* Bayesian Inference for Psychology . Part II : Example Applications with JASP. *prep* (2016).
56. Ioannidis, J. P. a. Why most published research findings are false. *PLoS Med.* **2**, 0696–0701 (2005).
57. Wagenmakers, E. *et al.* Bayesian Inference for Psychology . Part I : Theoretical Advantages and Practical Ramifications. *prep* 1–42 (2016).
58. Bays, P. M., Catalao, R. F. G. & Husain, M. The precision of visual working memory is set by allocation of a shared resource. *J. Vis.* **9**, 7.1-11 (2009).