AMERICAN PSYCHOLOGICAL ASSOCIATION

Journal of Experimental Psychology: General

ISSN: 0096-3445

2022, Vol. 151, No. 10, 2324–2341
https://doi.org/10.1037/xge0001208

Is the Juice Worth the Squeeze? Learning the Marginal Value of Mental Effort Over Time

A. Ross Otto[1], Senne Braem[2], Massimo Silvetti[3], and Eliana Vassena[4, 5]

[1] Department of Psychology, McGill University
[2] Department of Experimental Psychology, Ghent University
[3] Computational and Translational Neuroscience Laboratory, Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy
[4] Department of Experimental Psychopathology and Treatment, Behavioral Science Institute, Radboud University Nijmegen
[5] Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen

In keeping with the view that individuals invest cognitive effort in accordance with its relative costs and benefits, reward incentives typically improve performance in tasks that require cognitive effort. At the same time, increasing effort investment may confer larger or smaller performance benefits—that is, the marginal value of effort—depending on the situation or context. On this view, we hypothesized that the magnitude of reward-induced effort modulations should depend critically on the marginal value of effort for the given context, and furthermore, the marginal value of effort of a context should be learned over time as a function of direct experience in the context. Using two well-characterized cognitive control tasks and simple computational models, we demonstrated that individuals appear to learn the marginal value of effort for different contexts. In a task-switching paradigm (Experiment 1), we found that participants initially exhibited reward-induced switch cost reductions across contexts—here, task switch rates—but over time learned to only increase effort in contexts with a comparatively larger marginal utility of effort. Similarly, in a flanker task (Experiment 2), we observed a similar learning effect across contexts defined by the proportion of incongruent trials. Together, these results enrich theories of cost-benefit effort decision-making by highlighting the importance of the (learned) marginal utility of cognitive effort.

Keywords: effort, reward, marginal value, reward, cognitive control

Supplemental materials: https://doi.org/10.1037/xge0001208.supp

The resource-limited nature of cognitive processing prescribes that people should exert cognitive effort only when it is worthwhile—that is, when the costs incurred by expending additional effort are justified by the benefits it may confer (Hull, 1943; Kool & Botvinick, 2018; Shenhav et al., 2017). For example, reward incentives mobilize cognitive processing resources across diverse cognitive control tasks (Bijleveld et al., 2010; Frömer et al., 2021; Hübner & Schlösser, 2010; Otto & Vassena, 2021; Padmala & Pessoa, 2011), which is thought to arise from a cost-benefit tradeoff calculation (Shenhav et al., 2017): Increasing the benefits tied to effort exertion offsets the cognitive costs of exerting effort. People also tend to avoid exertion of cognitive effort when given the choice, in line with the idea that mentally demanding behavior is experienced as costly (Kool et al., 2010; Vogel et al., 2020; Westbrook et al., 2013) and further buttressing the notion that effort allocation arises from a cost-benefit calculus.

Less is known, however, about how individuals integrate cost and benefit information in making effort allocation decisions over time. Providing explicit information regarding task demands (i.e., effort costs) and available reward incentives (i.e., benefits) both modulate effort expenditure and alter individuals' "decisions" to engage in effortful processing (Braem, 2017; Krebs et al., 2012; Vassena et al., 2019). One important but open question concerns whether the marginal benefits of rewarded effort interact with the demand level of the environment. In other words, does an individual's inclination to adjust effort expenditure over time further depend on the context-specific benefits of this effort allocation? Intuitively, efficient allocation of effort should consider the net utility (benefits minus costs) of effort expenditure: In contexts where increasing control allocation is costly, reward incentives

This article was published Online First April 7, 2022.

A. Ross Otto https://orcid.org/0000-0002-9997-1901

This research was supported by an NSERC Discovery Grant and a CFI Infrastructure Grant to A. Ross Otto. Senne Braem is supported by an ERC Starting Grant (European Union Horizon 2020 program, Grant 852570). This work has not been presented previously at a conference or other research venue. All raw data pertaining to this study can be accessed via the Open Science Framework at https://osf.io/7zsrv/. The authors are grateful to Nick Yeung and Guido Gendolla for helpful discussions in the development of this work.

Correspondence concerning this article should be addressed to A. Ross Otto, Department of Psychology, McGill University, 2001 McGill College Avenue, Montréal, QC H3A 1G1, Canada. Email: ross.otto@mcgill.ca

should exert less influence on effort outlay than contexts where increasing control allocation is inexpensive.

Building on this framework, we propose that reward-guided effort allocation should depend on the marginal utility of increasing effort in a particular context—that is, the additional task performance benefits gained from increasing effort allocation—and that the computation of this marginal utility should drive learning of efficient effort allocation, over time. While marginal utility is a foundational idea for understanding how, for example, a decision maker takes into account their current wealth level in making reward-related choices (Li & Hsee, 2021), marginal utility may be a critical concept for understanding motivated cognitive control as it suggests boundary conditions on when reward incentives should be expected to mobilize effortful cognitive processing.

Learning an efficient effort allocation should require exposure to its marginal utility on the basis of direct experience. This idea is in line with reinforcement learning accounts of control allocation, which posit that people learn the expected value of control allocation based on features of the task environment (Abrahamse et al., 2016; Bustamante et al., 2021; Chiu & Egner, 2019; Krebs et al., 2010; Silvetti et al., 2018; Verguts et al., 2015). Extending this line of inquiry, here we consider whether individuals learn to modulate reward-guided effort, over time, in accordance with the marginal utility of effort allocation in different contexts. We investigated the hypothesis that differentially experienced marginal benefits of effort intensification can drive context-specific learning of effort allocation.

To illustrate the idea of marginal benefit of control allocation, consider the case of task switching for which flexible responding in the face of changes in stimulus-response rules imposes considerable executive control demands (Monsell, 2003). These task set reconfigurations result in the pervasive "switch costs," which are reduced by monetary incentives (Kleinsorge & Rinkenauer, 2012; Otto & Vassena, 2021; Sandra & Otto, 2018). As in other cognitive control tasks, larger incentives are thought to bring about increases in effort investment, compared to small or no incentives (Botvinick & Braver, 2015; Shenhav et al., 2017).

At the same time, the frequency with which task switches occur modulates switch costs—such that more frequent task switches engender smaller switch costs—presumably due to the increased demand for effortful control processes required relative to when switches are rare (Bogdanov et al., 2021; Dreisbach & Haider, 2006; Liu & Yeung, 2020; Monsell & Mizon, 2006). Here, we propose that the marginal benefit of increasing effort allocation depends on context and, specifically, on the "default" control level, vis-à-vis the task switch rate. Hence, boosting the benefits of additional effort exertion—for example by increasing reward incentives—should lead, over time, to different magnitudes of reward-induced effort modulations in contexts with different default control levels. In contexts with already-high control demand, people should learn that the marginal benefit of additional effort exertion is negligible and over time, as a consequence, reduce reward-induced effort modulations. Conversely, in contexts with a lower default level of control demand, the marginal benefit of increased control allocation in switch trials remains substantial.

We formalized these predictions with simulations of an established computational model of task switching (Yeung & Monsell, 2003). This model allows making ex ante predictions about the relationship between cognitive control (i.e., effort) allocation and task switch costs, depending on the task switch rate. We used this approach to demonstrate, computationally, that enacting switch cost reductions is more effortful under high-demand (i.e., high task switch rate) versus low-demand (i.e., low switch rate) conditions. In doing so, we are able to quantify, from first principles, the expected marginal benefit (in terms of task switch cost reductions) of increasing effort allocation under different environment demand levels. In short, Yeung and Monsell's simple response competition model assumes that a top-down "control input" signal—of which the application is effortful and modulated by reward incentives (Liu & Yeung, 2020; Otto & Vassena, 2021)—reduces response interference between tasks. By simulating two different task switch rates, instantiated as task-level priming from the previous trial, the model yields predicted task switch and task repetition response times (RTs) as a function of control input (Figures 1A and B) and, consequently, task switch costs (Figure 1C).

In this model (see "Method" below for simulation details), an important RT pattern emerges, irrespective of control input: A higher task switch probability engenders smaller task switch costs (Figures 1A, B, and C), corroborating previous observations (Dreisbach & Haider, 2006; Liu & Yeung, 2020; Monsell & Mizon, 2006). Of note, the slope of this curve relating control input to switch costs (Figure 1C) is steeper in the low-switch-rate context than the high-switch-rate case, which means that a 25-ms switch cost reduction, for example, would require a markedly larger proportional increase in control input under a high-switch-rate (dark curve) than a low-switch-rate environment (light curve). To better illustrate this point, the marginal decrease (or derivative; Figure 1D) in switch costs with respect to increases in control input is more negative in the low, versus high, switch rate context at all values of control input considered.

Having demonstrated, computationally, that the marginal benefit of additional effort exertion—in terms of switch cost reduction—should be greater in low-demand versus high-demand contexts, we experimentally probed a key question: Do people learn, over time, efficient reward-guided control allocation in a task-switching paradigm (as illustrated in Figure 2)? More specifically, does this (learned) efficient allocation reflect both (a) varying reward incentives and (b) the differing marginal benefits of increasing effort investment across demand contexts?

To formalize our ex ante predictions concerning how reward-induced switch cost modulations might change over time as a function of the (presumably learned) marginal value of effort, Figure 1E depicts predicted switch costs for a model that applies control input directly in proportion to reward incentives—indiscriminately across demand levels—over and above the default control input determined by the task switch rate. Thus, we predicted that individuals would increase effort in accordance with reward incentives, irrespective of demand level, which would manifest as smaller switch costs in higher-reward conditions (Hall-McMaster et al., 2019; Otto & Vassena, 2021). Over time, we hypothesized that participants learn that reward-induced effort investment is more beneficial in the low-demand context, by virtue of the larger marginal decrease in switch cost per unit of additional control. We can formalize this idea with a model that modulates control input in accordance with reward incentives but, critically, scales these control input modulations by the marginal value of effort—that is, the slope of the relationship between control input changes and switch cost changes

Figure 1

Simulated Task Switch Costs Resulting From Yeung and Monsell's (2003) Model of Task Switching

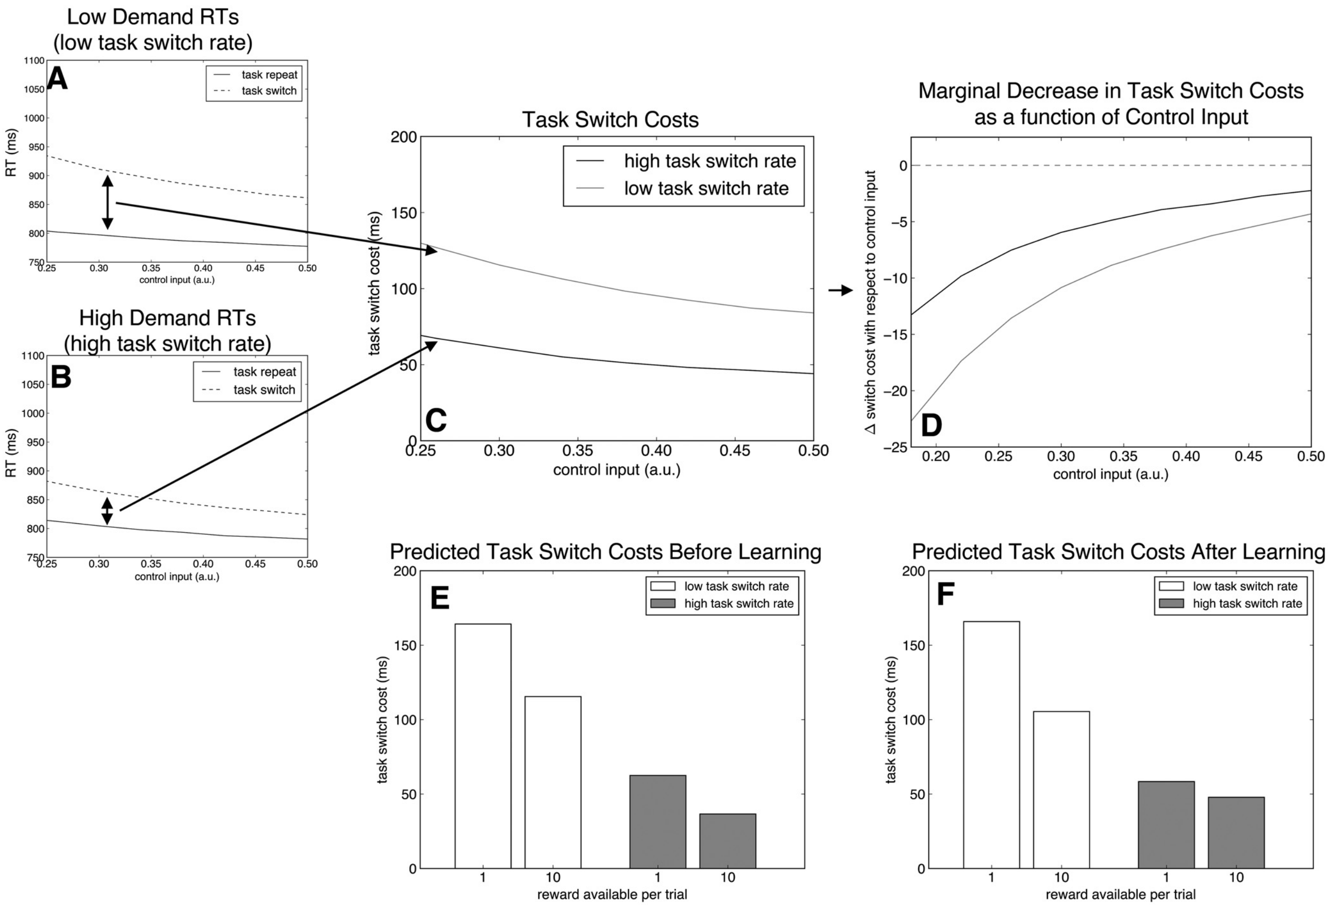

Note. Panels A and B: The relationships between task repetition and task switch response times (RTs) as a function of control input for low-switch-rate (i.e., low-demand) contexts and high-switch-rate (i.e., high-demand) contexts. Panel C: Computed task switch costs (task switch minus task repetition RTs) as a function of control input. Panel D: Change in switch cost with respect to increases in control input for high- and low-switch-rate contexts. Panel E: Predicted task switch costs, as a function of reward incentive level, for a model that indiscriminately increases control input with reward level. Panel F: Predicted task switch costs for a model that increases control with reward levels in accordance with the marginal value of effort in each demand context.

(Figure 1D)—for which the slope of the high-demand context is roughly one half of that of the low-demand context. Accordingly, we predicted that we would observe, as an end point of learning this marginal value, reward-induced modulations of switch costs *only* in the low-demand context, as depicted in Figure 1F.

To demonstrate the generalizability of this adaptive learning mechanism across diverse cognitive control paradigms, we also examined whether individuals learn this same efficient allocation of control in the classic Eriksen flanker task (Eriksen & Eriksen, 1974; Yu et al., 2009). Indeed, as it has been observed in variants of this simple conflict task that reward incentives reduce incongruence effects (Burton et al., 2021; Hübner & Schlösser, 2010; Yamaguchi & Nishimura, 2019), we should also expect to find evidence for learned sensitivity to the marginal value of effort exertion in reward-induced reductions in control of response interference.

Finally, across both experiments, we examined, in exploratory analyses, whether self-reported individual differences in effort costs, operationalized by the Need for Cognition scale (NFC; Cacioppo et al., 1984)—a trait measure of an individual's intrinsic motivation to exert cognitive effort—could predict participants' sensitivity to the marginal value of effort investment. In past work, we have observed that lower-NFC individuals are more inclined to increase cognitive effort investment in accordance with reward incentive, manifesting as larger reward-induced switch cost modulations (Sandra & Otto, 2018). Taking this result as evidence that perceived effort costs, measured by the NFC, also bear upon cost-benefit effort allocation decisions, we might expect here that low-NFC individuals should also be more sensitive to the marginal benefits of effort investment, compared to high-NFC individuals, for which effort costs might be negligible.

Experiment 1: Task Switching

In Experiment 1, we used a task-switching paradigm, depicted in Figure 2, to examine whether participants would learn to selectively

Figure 2
Task-Switching Paradigm Used in Experiment 1

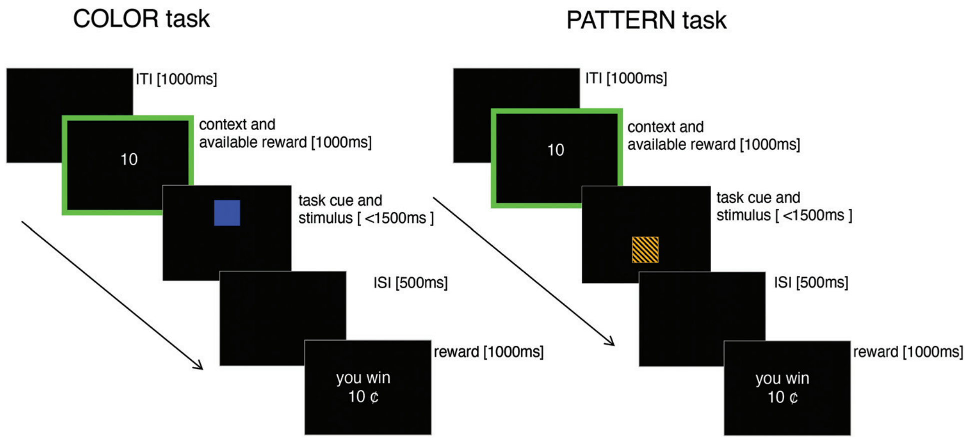

Note. Depending on the stimulus location (top or bottom), participants either indicated the color (blue or orange; dark versus light gray, respectively, as printed) or the pattern (stripes or solid) of a square. The reward available for making a correct response was displayed before the stimulus, and the switch rate (i.e., demand level) was signaled by a green or red border, light and dark gray, respectively, in print, indicating a low- or high-switch-rate (i.e., demand) context, respectively. ITI = intertrial interval; ISI = interstimulus interval. See the online article for the color version of this figure.

modulate their effort in accordance with reward incentive levels. Specifically, we sought to examine whether participants would take into consideration the differing marginal utility of increasing effort expenditure across two different demand contexts.

Method

Participants

We recruited 104 U.S. participants on Amazon's Mechanical Turk (MTurk; Crump et al., 2013), who were paid a fixed amount ($3 U.S.) plus a bonus contingent on their task performance, ranging from $1 to $2. Participants provided informed consent in accordance with the McGill University Research Ethics Board. Following the criteria used in our previous work (Otto & Vassena, 2021), we excluded the data of 13 participants who failed to perform either task with an accuracy of at least 75% on task repetitions and seven participants who missed 10% or more response deadlines on either the preliminary or reward phase of the task, leaving 83 participants in the final analyses.

Task-Switching Paradigm

In a preliminary phase, participants completed 100 trials of a task-switching paradigm in the absence of incentives to gain familiarity with the task. On each trial of this paradigm, which followed previous work (Otto & Daw, 2019; Otto & Vassena, 2021), a colored square appeared onscreen and participants needed to report either whether the square was blue or orange (the "COLOR" task) or whether the square's fill was solid or striped (the "PATTERN" task). The position of the square on the screen (lower half vs. upper half; counterbalanced across subjects) indicated which task the subject was to perform (see Figure 2). Across both tasks, responses were either associated with a left- or right-hand button press (e.g., blue = left, orange = right; solid = left, striped = right), using the "E" or "I" buttons on the keyboard. Mappings of stimuli features to keys were counterbalanced across participants. Fifty of these preliminary trials had a low task switch probability (10%; meaning that 90% of the trials repeated the previous task, the "low demand" context), while 50 had a high task switch rate (50%; the "high demand" context). Participants were given 1,500 ms to respond, at which point they received feedback onscreen indicating they made a correct response.

Following this preliminary phase, subjects began the reward phase, where a number at the beginning of each trial signaled the reward available for correct responses (either 1 or 10 cents; see Figure 2). The demand context was indicated by a colored frame around the available reward. In the low (green) and high (red) difficulty contexts, the task switch probabilities were 10% and 50%, respectively. Each difficulty context block was 40 trials long and consisted of two reward "miniblocks" composed of 20 trials of either 1- or 10-cent incentives contingent on making a correct response within the response deadline. Thus, each difficulty context contained two levels of reward, pseudorandomized within difficulty contexts. Participants completed six demand context blocks, the orders of which were pseudorandomized across participants, totaling 240 trials.

Prior to the main task, participants completed the NFC scale, an 18-item questionnaire that measures participants' intrinsic motivation to engage in cognitively demanding activities (for example, "I prefer complex to simple problems" and "I prefer my life to be filled with puzzles I must solve"; Cacioppo et al., 1984). Participants also completed the Behavioral Inhibition System/Behavioral Activation System scales (BIS/BAS; Carver & White, 1994).

Data Analysis

To ensure that task-switching behavior reflected participants' experienced demand context—that is, the high or low task switch rate—our analyses omitted the first 10 trials of each demand block and, following previous work, excluded outlier trials with RTs greater than 3 standard deviations from each subject's mean RT (Otto & Daw, 2019; $<$ 1% of trials per participant; note that the key patterns of significance hold without this exclusion applied) and excluded trials following errors (Liu & Yeung, 2020; $<$ 10% of trials per participant). We estimated mixed-effects regressions using the lme4 package for the R programming language (Bates & Maechler, 2009), taking each miniblock's task switch cost (mean log-transformed correct task switch RT minus the mean log-transformed correct task repetition RT) as the outcome variable and the demand level (that is, task switch probability) and available reward for that miniblock, task block (representing whether the miniblock occurred in the first or last 120 trials of the experiment) as predictor variables with all possible interactions between these predictors. This model specification took all predictor variables as fixed and random effects (see online supplemental materials for regression equation). Significance tests for each coefficient estimate were performed using Satterthwaite's method in the lmerTest package (Kuznetsova et al., 2017). In the model examining individual differences in NFC, we *z*-scored these NFC scores across participants before entering them as a predictor variable. In performing post hoc pairwise comparisons, we corrected for false discovery rate using the Benjamini and Hochberg (1995) procedure.

Computational Model of Task Switching

We implemented an established model of task switching, described in detail by Yeung and Monsell (2003), previously used to model the effects of reward incentives (Otto & Vassena, 2021) and task strength (Spitzer et al., 2019) on task switch RT costs. This model assumes that task responses result from competition between the two tasks (here, the "shape" vs. "pattern" task), each activated in accordance with the default "strength" of the task as well as possible priming from task activation on the previous trial. Note that this model, as originally described, makes predictions about RTs but not response accuracy per se. In particular, the task performed on the previous trial receives additional input via priming—thereby increasing its activation—while the task that was not performed on the previous trial receives zero additional input from priming. As a result, on task switches, competition between the two tasks is heightened—due to task priming—and the consequent activation levels of the two tasks result in slower RTs. Importantly, providing additional "control input"—which increases the activation level of the to-be-completed task, instantiating a form of top-down control—can counteract the task priming responsible for task-switching costs.

On each trial, the net input for a given task i (that is, shape or pattern) is determined by a linear combination of four input sources:

$$\text{input}_i = \text{strength}_i + \text{priming}_i + \text{control}_i + \text{noise} \tag{1}$$

where $strength_i$ corresponds to the strength of the task, $priming_i$ corresponds to previous activation (or "inertia") from the previously executed task (zero if the task was not executed on the previous trial), $control_i$ represents control input (varied parametrically in our simulations), and *noise* is simply zero-mean Gaussian noise (σ = .1). The activation of each task set is computed by the following negatively accelerated function:

$$\text{activation}_i = 1 - \exp(-1.5 \times \text{input}_i) \tag{2}$$

The response generation time is computed by dividing a threshold by the generation rate, calculated as the normalized activation of both tasks:

$$\text{generation rate}_i = \text{activation}_i / \Sigma\text{activation} \tag{3}$$

$$\text{generation time}_i = \text{THRESHOLD}/\text{generation rate}_i \tag{4}$$

The threshold parameter was arbitrarily set to 350. Finally, the model outputs a simulated RT for the task set that first produces a response:

$$\text{RT} = 150 + \text{generation time}_i + \text{resolution time} \tag{5}$$

where the resolution time is computed by sampling an ex-Gaussian distribution (μ = 150, σ = 10, τ = 40).

In the present simulations, we instantiated the task switch rate (10% versus 50%) by altering the priming parameter: A low task switch rate corresponds to a larger value of the priming parameter (.1) and more task priming should occur with frequent task repetitions, whereas a high task switch rate is instantiated with a smaller value of the task priming parameter (.05) because task priming (with respect to the previous trial's task) is less beneficial in the face of less frequent task repetitions. Further, we assume that the strength parameters for the two tasks are equal (.1) on the basis of similar performance (in both accuracies and RTs) we observe between the shape and pattern tasks (see "Overall Task Performance" below).

We simulated 500,000 trials of the model under each of the two demand levels (that is, task switch rates) for each control input parameter value, ranging from .20 to .50, and plotted task switch and task repeat RTs as a function of control input (Figure 1A and B for low and high demand, respectively). Task switch costs were then computed as the difference between task switch and task repeat, as a function of control input for each task switch rate (Figure 1C). Finally, for each demand level, we computed the marginal decrease of task switch costs with respect to control input by taking the backward difference of switch costs (with control input Δ = .04) yielding an approximate derivative of switch costs with respect to control input (Figure 1D).

Incentivized Model Simulations

We simulated the effect of trial-level reward incentives on task switch costs in Experiment 1 (Figure 1E and F), which we operationalized as an increase in control input over and above a default control level associated with the demand context, following our previous work (Otto & Vassena, 2021):

$$\text{control}_{\text{context,reward}} = \text{control}_{\text{context}} + (\text{reward} \times \text{marginal\_value}) \times 0.05$$

We set the default level of control for each demand context, $control_{context}$, to .05 and .25 in the low- and high-demand contexts, respectively. Our simulations assumed that initially, the model naively and indiscriminately applies control as a function of reward level across the two demand contexts, which is instantiated by setting *marginal_value* to .5 in both reward contexts, yielding the patterns of reward-induced switch cost reductions seen in

Figure 1E. We assume that when the differing marginal values of effort for the two demand contexts have been learned (Figure 1D), reward-induced control in the high-demand context is scaled by a factor of one half, mirroring the ratio of the average slopes of the value function relating control input to switch costs in the high- versus low-demand contexts. Accordingly, the end-of-learning model assumes that the *marginal_value* takes a value of .5 in the low-demand context and .25 in the high-demand context, resulting in the pattern of switch cost reductions depicted in Figure 1F. In all cases, the resultant control input is scaled by a constant (.05) to ensure input is provided in units appropriate for the task-switching model.

Results and Discussion

Overall Task Performance

We observed significant task switch costs, expressed as longer RTs on task switch versus task repetition RTs, both in low-demand blocks, which required task switches on 10% of trials (mixed-effects regression on RTs: $\beta = .205$, $SE = .018$, $p < .0001$), and in high-demand blocks, which required task switches on 50% of trials ($\beta = .094$, $SE = .007$, $p < .0001$; see Table 1). Mirroring previous work examining task switch costs as a function of task switch rates (Duthoo et al., 2012; Fröber & Dreisbach, 2017; Mayr, 2006; Monsell & Mizon, 2006), these task switch costs—expressed as the difference between task switch and repetition RTs—were significantly smaller in high-demand ($M = 82.34$, $SD = 79.30$) as compared to low-demand blocks ($M = 250.12$, $SD = 127.79$; $\beta = -.20$, $SE = .03$, $p < .0001$). Similarly, task switches were significantly less accurate than task repetitions in both low-demand blocks (mixed-effects logistic regression; task switch effect $\beta = -1.598$, $SE = .154$, $p < .0001$) and high-demand blocks ($\beta = -.443$, $SE = .094$, $p < .0001$; Table 1), and this accuracy effect was smaller in high-demand blocks compared to low-demand blocks (interaction $\beta = .546$, $SE = .0393$, $p < .0001$).

Reward Incentives and Switch Costs

We next examined, over time, how reward incentives modulated switch costs in two demand levels. Figure 3 depicts task switch costs as a function of demand level (low vs. high), reward incentive (1 cent vs. 10 cents), and trial block (first half vs. second half, depicted in Panels A and B, respectively). While it is apparent that switch costs were smaller overall with larger incentive—consistent with previously observed reward-induced switch cost reductions (Fröber & Dreisbach, 2016; Kleinsorge & Rinkenauer, 2012; Otto & Vassena, 2021)—the effect of incentives in each of the two demand levels changed over time. In the first half of the experiment, reward incentives reduced switch costs both in high- and low-demand blocks, suggesting that initially, participants were willing to increase control allocation in accordance with reward regardless of the cost of increasing control—that is, even if marginal increases in effort have lower net benefit in high-demand blocks. However, in the second half of the experiment, reward incentives reduced switch costs only in low-demand blocks, where marginal increases in effort expenditure carry greater net benefit, but not in high-demand blocks, where marginal increases in effort expenditure carry a smaller benefit.

Statistically, in a mixed-effects model examining switch costs as a function of reward, demand level, and trial block (see full coefficient estimates in Table 2), we observed a significant interaction between reward, demand level, and trial block upon switch costs ($\beta = .1634$, $SE = .0589$, $p = .006$), indicating that the observed reward-induced switch cost reductions jointly depended on demand level and trial block. We also observed a negative but nonsignificant main effect of reward ($\beta = -.0227$, $SE = .0326$, $p = .441$), suggesting that this reward incentive did not uniformly operate over both demand levels and over time. Further supporting this observation that these switch cost modulations were time dependent, we did not find a significant two-way interaction between reward incentives and demand level ($\beta = -.0363$, $SE = .0411$, $p = .378$).

We further probed the specificity of the three-way interaction between reward incentive level, demand context, and trial block in a series of post hoc, pairwise tests. In the first half of the experiment, we observed that reward incentive level exerted a significant effect of reward incentive level in the high-demand context ($t = 2.596$, $p = .0415$; all tests corrected for false discovery rate) but not in the low-demand context ($t = 1.2728$, $p = .274$). In the second half, we observed a marginally significant effect of reward incentive level in the low-demand context ($t = 2.050$, $p = .0870$) but no significant effect in the high-demand context ($t = .0525$, $p = .9581$). Together, these comparisons suggest that participants initially exhibited reward-induced control modulation in the high-demand condition, but later, these reward-induced control modulations were only apparent in the low-demand condition.

We also probed whether participants' diminishing sensitivity to incentive levels in the high-demand condition resulted in lower overall levels of reward receipt, rather than merely modulating the extent of reward-guided effort allocation between reward amounts.

Table 1
Average Median Response Times (RTs) and Error Rates for Task Repeat and Task Switch Trials Across the Reward Incentive and Demand Levels in the Task-Switching Paradigm (Experiment 1)

Reward amount	Repeat RT (*SD*)	Switch RT (*SD*)	Repeat accuracy (*SD*)	Switch accuracy (*SD*)
Low-demand blocks				
1 cent	658.59 (95.46)	923.39 (158.42)	0.95 (0.09)	0.87 (0.14)
10 cents	668.12 (94.14)	903.01 (151.14)	0.96 (0.05)	0.89 (0.14)
High-demand blocks				
1 cent	752.19 (110.2)	832.47 (119.17)	0.9 (0.1)	0.85 (0.13)
10 cents	753.66 (119.08)	838.06 (114.28)	0.92 (0.09)	0.88 (0.11)

Figure 3

Task Switch Costs in Experiment 1, Expressed as the Difference Between Median Task Switch RTs and Task Repetition RTs, as a Function of Available Reward and Demand Context (i.e., Low Versus High Task Switch Rate) and Trial Block (First Versus Second Half, Depicted in Panels A and B, Respectively)

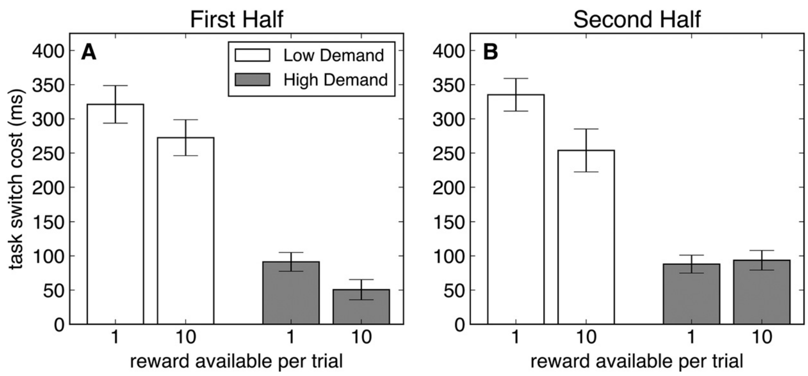

Note. Initially, participants exhibit reward-induced switch cost reductions in both demand contexts, but after learning, reward-induced switch cost reductions are only apparent in the low-demand context.

As reward incentives were tied to correct responses, we could probe this question by analyzing trial accuracy on high-demand blocks as a function of trial block (early vs. late), finding no significant effect of trial block upon accuracy on high-demand blocks (β = .0430, *SE* = .0558, *p* = .441). This lack of a trial block effect suggests against the possibility that the above-mentioned changes in strategy reflected decreased overall performance (and reward receipt) in the high-demand condition, but rather indicates that people learned to exert less (costly) effort without losing reward, in accordance with our learned marginal benefit of effort hypothesis. Finally, we probed whether fatigue—manifesting as a general slowing effect as a function of time on task (Lorist et al., 2005)—might explain the abolition of a reward effect in the high-demand context over time. However, we observed no apparent effect of trial number upon high-demand RTs, controlling for task switches (main effect of trial number: $\beta = -.0142$, *SE* = .0134, *p* = .295; interaction between trial number and trial type: β = .0073, *SE* = .006, *p* = .268), suggesting against the possibility of a fatigue effect selective to the high-demand context. We also found no evidence for more general fatigue effects (see online supplemental materials).

Table 2

Mixed-Effects Regression Coefficients Indicating the Influence of Demand Level (Task Switch Probability), Reward Incentive Level, and Experimental Block (Early Versus Late) Upon Task Switch Costs in Experiment 1

Coefficient	Estimate (*SE*)	*p* value
(Intercept)	0.3311 (0.0248)	<.0001*
Reward	−0.0227 (0.0326)	.487
Demand level	−0.212 (0.0299)	<.0001*
Trial block	0.0517 (0.0327)	.115
Reward × Demand Level	−0.035 (0.0415)	.399
Reward × Trial Block	−0.0815 (0.0472)	.085
Demand Level × Trial Block	−0.077 (0.0413)	.063
Reward × Demand Level × Trial Block	0.1634 (0.0589)	.006*

Note. *SE* = standard error.
* Denotes significance at $p < .05$ level.

Individual Differences in Need for Cognition (NFC) and Reward Sensitivity (BAS)

We also examined, in exploratory analyses, how the dynamics of reward-induced switch cost modulations differed as a function of individual participants' NFC scores (*M* = 62.03, *SD* = 17.90). While we intuited that low- and high-NFC participants might be differentially sensitive to the marginal value of effort, we did not have strong predictions about the locus of the possible predictive NFC effect vis-à-vis main effects or interactions. Figure 4 depicts the same analysis as above, separately considering low-NFC and high-NFC participants (defined by a median split). Interestingly, we found that these time-dependent changes in reward-induced switch cost modulations were especially pronounced in low-NFC individuals. By contrast, high-NFC participants, particularly on later trials, exhibited switch cost reductions in accordance with reward levels in both demand levels. Statistically, this observation was supported by a significant four-way interaction between NFC (taken continuously), reward, demand level, and trial block ($\beta = -.1266$, *SE* = .0539, *p* = .033; see Table 3 for full coefficient estimates), indicating that the observed group-level interaction between demand level, reward, and trial block was moderated by participants' individual NFC levels. Using the same regression approach, we did not observe that reward responsiveness, as measured by the BAS subscale, exerted any predictive effect on reward-induced switch cost reductions overall and as a function of demand level, trial block, or their interactions (all interactions *ps* > .381). In summary, and dovetailing with previous observations that high-NFC individuals tend to exert more effort overall (Sandra & Otto, 2018; Westbrook et al., 2013), this pattern of learned effort allocation suggests that high-NFC individuals may be less inclined to optimize efficiency of effort allocation.

Figure 4

Task Switch Costs in Experiment 1 Plotted as a Function of Reward Incentive Amount, Demand Context, and Trial Block for Participants Low in Need for Cognition (Panels A and B) and Participants High in Need for Cognition (Panels C and D)

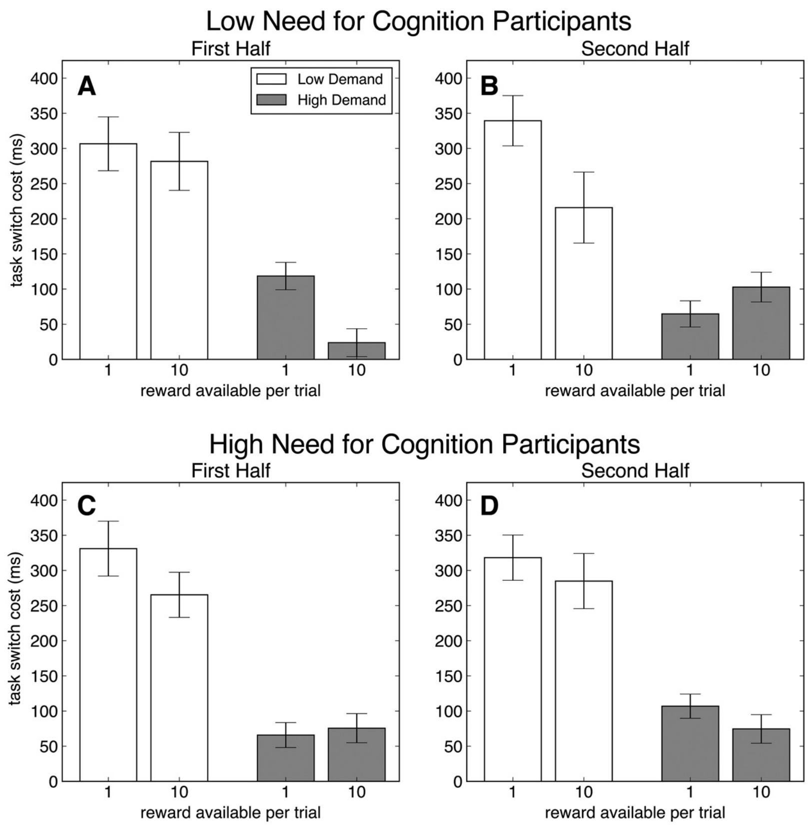

Note. Of note, the learning effect observed in the entire sample (Figure 3) is most pronounced in participants low in need for cognition.

Experiment 2: The Flanker Task

Experiment 2 examined the generality of this learning effect to a wholly different cognitive control paradigm: the arrow flanker task (Eriksen & Eriksen, 1974). On each trial of the task (see Figure 5), participants indicate the direction of a central target stimulus (< or >) presented between either congruent (>>>>>) or incongruent (> > < > >) "flankers." Typically, responses are slower on incongruent trials as compared to congruent trials, which we term the flanker incongruence effect. Importantly, past studies have observed modulations of the flanker incongruence effect observed in accordance with the proportions of incongruent stimuli, typically manipulated in a block-wise fashion: The incongruence effect is reduced in contexts with a larger proportion of incongruent trials (for example, 75% incongruent trials) compared to contexts with a smaller proportion of incongruent trials (for example, 25% incongruent trials), which has been attributed to strategic modulations of control (Gratton et al., 1992; Yu et al., 2009) or contingency learning (Braem et al., 2019; Schmidt, 2019).

Following the results of the demand level (that is, task switch rate) manipulation in Experiment 1, we intuited that the marginal benefit of increased control allocation—manifesting as reductions in incongruence costs under large incentive amounts—would also differ across demand contexts, as defined by the proportion of incongruent trials. To illustrate the marginal utility of effort allocation across these two demand contexts, we modified the task-switching model of Yeung and Monsell (2003), which in its original formulation predicts that flanker-like incongruence effects arise from "task sets" of unequal strength (see "Method" below for details). Our modified flanker model predicts smaller incongruence costs in contexts with larger proportions of incongruent trials (Figure 6A). These incongruence costs decrease in both the 50% and 80% incongruence contexts with increasing control input but, importantly, at different rates. Echoing the task-switching model simulations, the marginal utility of increasing control allocation, which can be seen as the rate of decrease of incongruence costs with respect to control input (Figure 6B), is consistently larger (that is, more negative) in the 50% (low-demand) context than in the 80% (high-demand) context.

On the basis of the differing marginal utility of effort investment across demand contexts, we reasoned that in the 50% incongruent trial context, participants would consistently modulate their control allocation in accordance with reward incentives. However, in the 80%

Table 3
Mixed-Effects Regression Coefficients Indicating the Influence of Demand Level (Task Switch Probability), Reward Incentive Level, Experimental Block (Early Versus Late), and Need for Cognition Upon Task Switch Costs in Experiment 1

Coefficient	Estimate (*SE*)	*p* value
(Intercept)	0.3314 (0.0248)	<.0001*
Reward	−0.0213 (0.0324)	.513
Demand level	−0.2122 (0.0297)	<.0001*
Trial block	0.0522 (0.0329)	.113
NFC	−0.0022 (0.0235)	.925
Reward × Demand Level	−0.0363 (0.0411)	.378
Reward × Trial Block	−0.0854 (0.0473)	.072
Demand Level × Trial Block	−0.0778 (0.0414)	.061
Reward × NFC	−0.0198 (0.0326)	.543
Demand Level × NFC	−0.0185 (0.0285)	.516
Trial Block × NFC	0.0077 (0.031)	.803
Reward × Demand Level × Trial Block	0.1673 (0.0591)	.005*
Reward × Demand Level × NFC	0.0669 (0.0411)	.104
Reward × Trial Block × NFC	0.0359 (0.0469)	.444
Demand Level × Trial Block × NFC	0.0302 (0.0397)	.448
Reward × Demand Level × Trial Block × NFC	−0.1252 (0.0585)	.033*

Note. NFC = need for cognition; *SE* = standard error.
* Denotes significance at $p < .05$ level.

incongruent condition, we would expect—as in Experiment 1—that participants would learn the negligible marginal value of effort investment and would accordingly cease to modulate effort in accordance with incentives over time. As in Experiment 1, we formalized these learning predictions with two different models. First, we simulated a "naïve" model that allocates control input directly in proportion to reward incentives (Figure 6C) irrespective of demand context. Then, we considered a model that has learned the differing marginal value of effort for each context (Figure 6B), which scales these control input increases by the marginal value of effort for each context, yielding the predicted pattern of incongruence costs in depicted in Figure 6D.

Method

Participants

We recruited 100 U.S. participants on MTurk, who were paid a fixed amount ($4 U.S.) plus a bonus contingent on their task performance, ranging from $1 to $2. Participants provided informed consent in accordance with the McGill University Research Ethics Board. We excluded the data of 20 participants who failed to perform with an accuracy of at least 75% on congruent trials and 14 participants who missed 10% or more response deadlines on either the preliminary or reward phase of the task, leaving 66 participants

Figure 5
The Eriksen Flanker Task, Used in Experiment 2

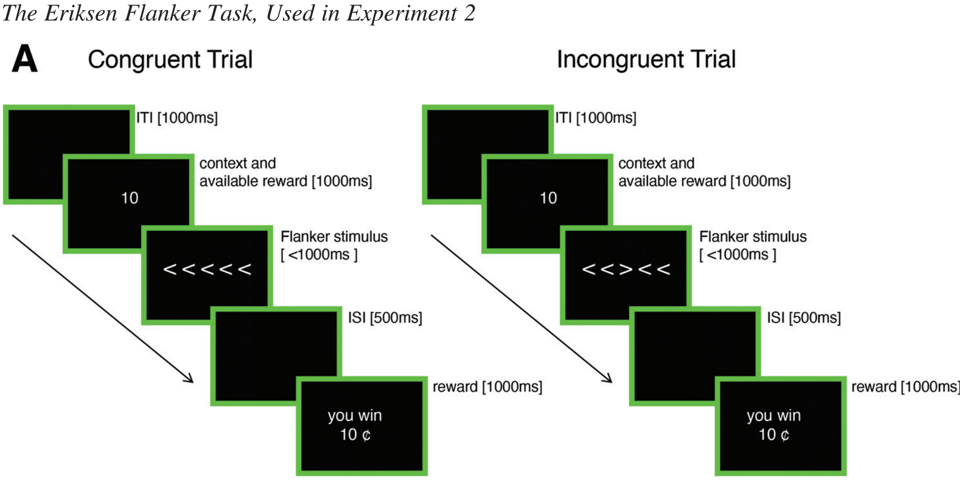

Note. The reward available for making a correct response was displayed before the stimulus, and proportion of incongruent trials (i.e., demand level) was signaled by a green or red (light and dark gray, respectively, in print) border around the screen. ITI = intertrial interval; ISI = interstimulus interval. See the online article for the color version of this figure.

Figure 6

Simulated Flanker Incongruence Costs Resulting From Model Simulations

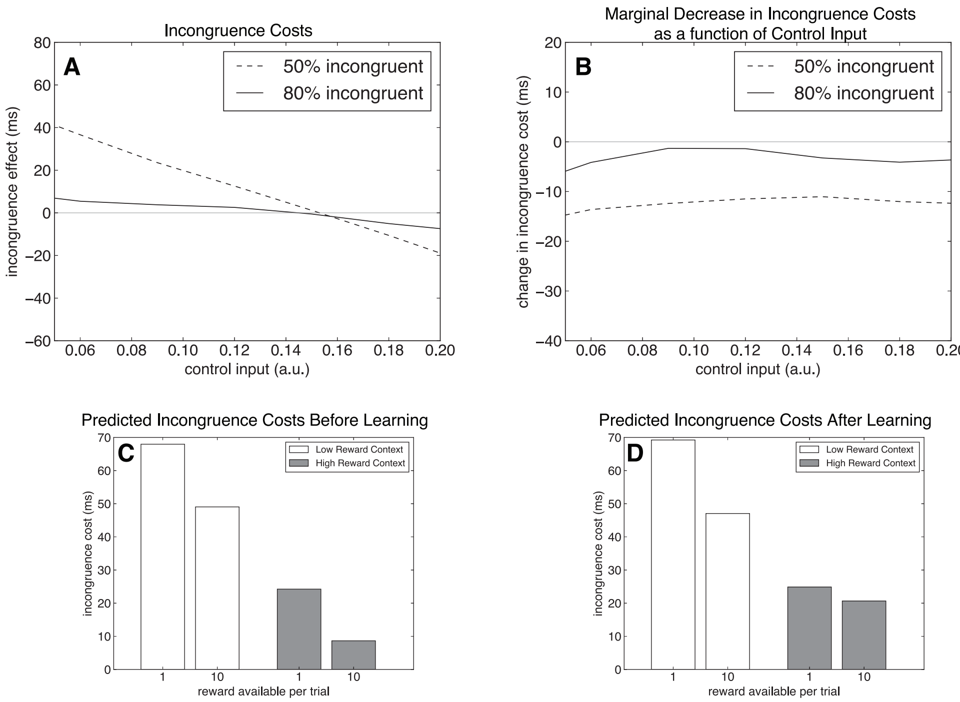

Note. Panel A: Simulated flanker incongruence effects costs (expressed as the difference between median incongruent and congruent RTs) as a function of control input, plotted for the 50% incongruent (low-demand) context and 80% incongruent (high-demand) context. Panel B: Change in incongruence cost with respect to increases in control input for 50% incongruent (low-demand) context and 80% incongruent (high-demand) contexts. Panel C: Predicted incongruence costs, as a function of reward incentive level, for a model that indiscriminately increases control input with reward level. Panel D: Predicted incongruence costs for a model that increases control with reward levels in accordance with the marginal value of effort in each demand context.

in the final analyses.[1] Prior to the flanker task, participants completed the NFC and BIS/BAS questionnaires.

Flanker Task

Participants performed a standard version of the flanker task (see Figure 5) in which they were required to identify the directionality of a central target ($<$ or $>$) presented between four congruent or incongruent flankers on either side of the stimulus (Eriksen & Eriksen, 1974), and the stimulus remained on the screen until the participant responded or 1,000 ms had elapsed. The timing and presentation of feedback for both the preliminary and reward phase of the task mirrored Experiment 1. In a preliminary (no reward) phase, participants first completed 50 low-demand trials (50% proportion incongruent) and 50 high-demand trials (80% proportion incongruent) to familiarize themselves with the flanker task and each demand level. The order of these demand levels was counterbalanced across participants. On all trials, the demand level context was signaled by a colored border (green vs. red, signaling low and high demand, respectively).

Following the preliminary phase, subjects began the reward phase, where, following Experiment 1, a number at the beginning of each trial signaled the reward available (1 or 10 cents) for making a correct response. Each demand level block was 20 trials long, comprised of two reward miniblocks of 10 trials of 1- or 10-cent incentives. Participants completed 12 demand context blocks, the orders of which were pseudorandomized across participants, totaling 240 trials.

Data Analysis

Following Experiment 1, we omitted the first five trials of each demand block to ensure that flanker behavior reflected the demand level of the current context and excluded outlier trials with RTs greater than 3 standard deviations (note that the exclusion of these outlier RTs does not affect the key patterns of significance described below). Our mixed-effects regression approach mirrored that of Experiment 1 but took each miniblock's incongruence costs—computed as the mean log-transformed correct incongruent RT minus the mean log-transformed correct congruent RT—as the outcome variable. We took all predictor variables as fixed and random effects (see online supplemental materials for regression equation).

[1] We should note this sample was collected during the global coronavirus pandemic (November 2020). Participant samples collected during this period have been characterized previously as less attentive MTurk than previous MTurk samples (Arechar & Rand, 2021).

Flanker Task Adaptation of Yeung and Monsell's (2003) Model

As Yeung and Monsell's response competition model of task switching can also account for flanker- or Stroop-like incongruence effects in the case of tasks with unequal strength (irrespective of repetitions or switches of the to-be-completed task), we adapted the model to make predictions about incongruence costs in the flanker task, which, importantly, varies as a function of the proportion of incongruent trials. To do this, we assume that each flanker trial involves two (possibly) competing "tasks"—identifying the directionality of the flanker stimuli versus identifying the directionality of the center stimulus—and that flanker stimuli exert more influence over responses than the central stimulus (Yu et al., 2009). Specifically, we assume that the "flanker" task and "center" task have respective strength parameters of .6 and .2. We further assume that control input is only applied to the center task, and accordingly, the flanker task receives zero control input on every trial.

In line with the literature on congruency sequence effects (Braem et al., 2019; Gratton et al., 1992), we assume trial-by-trial priming effects apply that depend on the congruency status of the previous trial. Namely, we assume that both the center and flanker tasks were primed when the previous trial was congruent (as both led to the correct response) and that task priming only takes effect for the center task (but not the flanker task) when the previous trial type was incongruent, to simulate the idea that people focus more on the relevant dimension following incongruent trials (e.g., Botvinick et al., 2001) or carry over control settings from the previous trial more generally (e.g., Braem et al., 2019; Egner, 2014). Following Yeung and Monsell's original model, the resolution time (in Equation 5 above) is governed by:

$$\text{resolution time} = r + f\left[r - \left(\text{generation time}_i - \text{generation time}_j\right)\right] \qquad (6)$$

where r was sampled from an ex-Gaussian distribution ($\mu = 150$, $\sigma = 10$, $\tau = 40$) and f is a function that dictates whether the difference in generation time between the two tasks causes interference or facilitation. Following Yeung and Monsell, f takes a value of .5 for incongruent stimuli and 0 for congruent stimuli.

We simulated 500,000 trials under each demand level (50% vs. 80% proportion incongruent) at control input parameter values ranging from .05 to .20. The task priming and threshold parameters were set arbitrarily to .4 and 300, respectively. The flanker incongruence effect was computed as the difference between incongruent and congruent RTs (Figure 6A). As in the task-switching model, we also computed an approximate derivative of incongruence effects with respect to control input (Figure 6B) via backward differencing. We applied the same procedure as described above to simulate the effect of reward incentives in the flanker task model, initially setting marginal_value to .1 in both reward contexts (Figure 6C) and, after learning, setting marginal_value to .1 and .025 in the low- and high-demand contexts (Figure 6D), respectively (mirroring the approximate ratio of the slopes of the incongruence costs depicted in Figure 6B).

Results and Discussion

Overall Task Performance

We observed significant incongruence effects in the flanker task—expressed as longer RTs on incongruent versus congruent RTs—in low-demand (50% congruent) blocks (mixed-effects regression on RTs: incongruence effect $\beta = .0843$, $SE = .00488$, $p < .0001$) and in high-demand blocks (80% incongruent; $\beta = .0729$, $SE = .005247$, $p < .0001$; see Table 4). Jointly examining the effects of trial type (congruent vs. incongruent) and proportion incongruence upon RTs, we observed a significant interaction between proportion of incongruent trials and trial type (congruent vs. incongruent), indicating that incongruence costs were significantly smaller in the 20% congruent (high-demand) condition ($\beta = -.0113$, $SE = .00482$, $p = .0225$), but there was no significant main effect of proportion congruence ($\beta = .003535$, $SE = .00399$, $p = .379$).

Examining accuracy, we observed incongruent trials were significantly less accurate in the low-demand condition (mixed-effects logistic regression on accuracy: incongruence effect $\beta = -3.498$, $SE = 1.442$, $p = .0153$) and in the high-demand condition ($\beta = -.6439$, $SE = .2311$, $p = .00534$; see Table 4). We did not observe a significant effect of proportion congruence (main effect $\beta = .3304$, $p = .2393$, $p = .167$) nor a significant modulation of incongruence costs by proportion incongruent trials (interaction $\beta = -.3332$, $SE = .2369$, $p = .160$).

Reward Incentives and Incongruence Effects

We then examined flanker incongruence costs (incongruent − congruent correct RTs) as a function of demand level (low vs. high), reward incentive (1 cent vs. 10 cents), and trial block (first half vs. second half). In the early blocks of the experiment (Figure 7A), reward incentive level appeared to decrease flanker effects only in the high-demand (80% incongruent) context, but in later blocks (Figure 7B), the locus of reward incentive effects shifted from the high-demand to the low-demand (50% incongruent)

Table 4
Average Median Response Times (RTs) and Error Rates for Congruent and Incongruent Trials Across the Reward Incentive and Demand Levels in the Flanker Task (Experiment 2)

Reward amount	Congruent RT (*SD*)	Incongruent RT (*SD*)	Congruent accuracy (*SD*)	Incongruent accuracy (*SD*)
Low-demand blocks				
1 cent	522.12 (81.68)	572.39 (82.06)	0.98 (0.05)	0.96 (0.07)
10 cents	519.84 (78.89)	565.48 (83.01)	0.99 (0.03)	0.97 (0.06)
High-demand blocks				
1 cent	520.7 (82.84)	567.52 (80.48)	0.99 (0.04)	0.97 (0.04)
10 cents	523.33 (82.56)	559.27 (82.99)	0.99 (0.04)	0.97 (0.05)

Figure 7

Flanker Incongruence Costs in Experiment 2, Expressed as the Difference Between Median Incongruent RTs and Median Congruent RTs, as a Function of Available Reward and Demand Context (i.e., Low Versus High Task Switch Rate) and Trial Block (First Versus Second Half, Depicted in Panels A and B, Respectively)

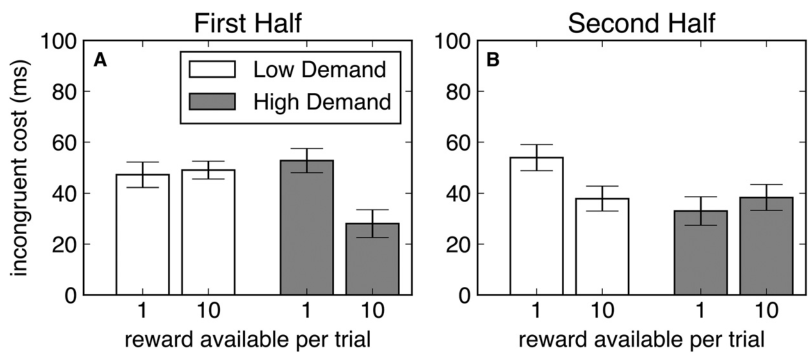

context. This time course of reward-induced flanker effect modulations is consistent with the notion that participants shifted the locus of reward-guided control to contexts with the largest marginal benefit for control allocation—in other words, from the high-demand condition (small marginal benefit) to the low-demand condition (large marginal benefit).

We estimated a mixed-effects regression examining incongruence costs as a function of reward, demand level, and trial block (early vs. late; see fixed-effect coefficient estimates in Table 5), finding significant interactions between reward and demand level ($\beta = -.0427$, $SE = .0128$, $p = .001$), between demand level and trial block ($\beta = -.035$, $SE = .0128$, $p = .006$), and, critically, between reward, demand level, and trial block ($\beta = .0701$, $SE = .018$, $p = .0001$). That is, in accordance with a learning account, the relationship between reward incentive and demand level depended on trial block. We further probed this interaction with post hoc pairwise tests and found that in the first half of the experiment, reward incentive level significantly decreased incongruence costs in the high-demand ($t = 3.40257$, $p = .0035$; all tests corrected) but not the low-demand context ($t = .29873$, $p = .7656$). In the second half, we found that reward incentive level significantly decreased incongruence costs in the low-demand context ($t = 2.272$, $p = .0494$) but not in the high-demand context ($t = -.7015$, $p = .645$). Mirroring the pattern of results in Experiment 1, these comparisons suggest that reward incentives initially influenced control allocation in the high-demand context, but later, the locus of these incentive effects shifted to the low-demand context.

As in Experiment 1, we also probed whether the apparent withdrawal of effort in the high-demand condition resulted in lower overall levels of reward receipt. To do this, we analyzed trial accuracy on high-demand blocks as a function of trial block (early vs. late) and found no significant effect of trial block upon accuracy on high-demand blocks ($\beta = -.1085$, $SE = .1418$, $p = .444$). Again, the absence of trial block effect in the high-demand condition indicates that effort modulations became less sensitive to varying incentives levels over time, rather than a general decrease over time in overall performance. Finally, examining the possibility of a fatigue explanation for the time-dependent disappearance of a reward effect in the high-demand condition, we probed whether RTs slowed over time in the high-demand context, finding little support for a slowing effect (effect of trial number: $\beta = -.002317$, $SE = .004$, $p = .1598$; interaction between trial number and trial type $\beta = -.005250$, $SE = .00375$, $p = .167$). We also found no evidence for more general fatigue effects (see online supplemental materials).

Individual Differences in Need for Cognition (NFC) and Reward Sensitivity (BAS)

Finally, following the exploratory analyses in Experiment 1, we probed whether individual differences in NFC moderated the learning effects observed in the flanker task. Adding individual NFC scores to the incongruence cost predicting model reported above, we observed no significant interactions between NFC and reward level, demand level, trial block, and the possible interactions between these variables ($ps > .355$; see Table 6 for full coefficients). The lack of predictive effect of NFC suggests that, unlike in task switching, individual differences in intrinsic motivation to exert cognitive effort did not exert any predictive bearing on reward-guided control allocation (and its learning) in the flanker task. Similarly, we did not observe any interactions between BAS reward responsivity and reward effects on incongruence costs, or between reward responsivity, demand level, and trial block (all interactions $ps > .228$). Put another way, our observed learning effects in the flanker task were similar for both low- and high-NFC individuals, possibly due to the considerably different structure of the flanker task.

General Discussion

The notion that reward incentives can mobilize cognitive processing resources has been influential and finds broad empirical support (Kool & Botvinick, 2018; Westbrook & Braver, 2015). Up to now, investigations of cost-benefit effort decision-making have typically treated costs and benefits as factors that exert a time-invariant influence on effort allocation decisions, presuming that given reward incentive level will evoke the same change in effort allocation at different times. Here, we considered the

Table 5

Mixed-Effects Regression Coefficients Indicating the Influence of Demand Level (Proportion Incongruent Trials), Reward Incentive Level, and Experimental Block (Early Versus Late) Upon Flanker Incongruence Costs in Experiment 2

Coefficient	Estimate (*SE*)	*p* value
(Intercept)	0.0761 (0.0076)	<.0001*
Reward	0.0202 (0.009)	.026*
Demand level	0.0125 (0.0092)	.177
Trial block	0.017 (0.0094)	.073
Reward × Demand Level	−0.0427 (0.0128)	.001*
Reward × Trial Block	−0.0391 (0.0128)	.002*
Demand Level × Trial Block	−0.035 (0.0128)	.006*
Reward × Demand Level × Trial Block	0.0701 (0.018)	.0001*

Note. *SE* = standard error.
* Denotes significance at $p < .05$ level.

possibility that people learn the marginal utility of allocating additional cognitive effort—that is, the benefit of increasing control allocation in accordance with incentives—over time.

Across two different cognitive control paradigms, we examined how reward-induced effort modulations vary over time across contexts with different marginal utilities of effort investment. We found compelling evidence that people learn to efficiently allocate effort in accordance with both the incentives and the context-dependent, marginal utility of increasing effort allocation, illustrated by simulations of a simple computational model (Figures 2 and 6). In task switching (Experiment 1), participants initially exhibited reward-induced effort increases in both high- and low-switch-rate contexts—as evidenced by decreases in task switch costs—but over time ceased to modulate effort investment in the high-switch-rate context. Similarly, we examined the same apparent learning in the flanker task (Experiment 2) by manipulating demand contexts by altering the proportion of incongruent trials. We found that participants learned to increase effort investment in low-demand contexts, where increasing control allocation results in appreciable reductions of incongruence costs, and at the same time ceased investing effort in accordance with incentives in the high-demand contexts, where increasing control allocation results in appreciable reductions of incongruence costs.

These results highlight the importance of marginal utility calculation in cost-benefit decision-making about cognitive effort allocation. Further, they speak to influential accounts of motivation, explaining the relationship between task demand and effort investment (Brehm & Self, 1989). The classic motivational intensity theory (MIT) postulates that individuals invest effort only when it yields a tangible benefit, withdrawing resources when this is not the case. This basic tenet is fully in line with our results. Second, MIT postulates then when difficulty is unknown, effort investment fully depends on "success importance" (i.e., the influence of reward incentive in the case of our task). This is in line with participants' initial behavior (where switch cost reduction is overall driven by reward) and with the idea that after learning of the marginal benefit, effort investment changes (resources are withdrawn when not worth it). While MIT theory's predictions by and large concern physiological markers of effort exertion (e.g., cardiovascular reactivity; Gendolla et al., 2012), our results show how they conceptually align with learning efficient effort allocation over time across two established cognitive control paradigms, in terms of switch costs (Experiment 1) or flanker incongruence costs (Experiment 2).

An important question concerns the sort of learning underlying the changes in effortful behavior observed here: What aspect(s) of the task environment are participants learning in order to produce the apparent changes in reward-induced control allocation between early and late blocks of the experiments? One possibility is that the notion

Table 6

Mixed-Effects Regression Coefficients Indicating the Influence of Demand Level (Proportion Incongruent Trials), Reward Incentive Level, Experimental Block (Early Versus Late), and Need for Cognition Upon Flanker Incongruence Costs in Experiment 2

Coefficient	Estimate (*SE*)	*p* value
(Intercept)	0.0761 (0.0077)	<.0001*
Reward	0.0202 (0.0091)	.027*
Demand level	0.0125 (0.0093)	.18
Trial block	0.017 (0.0095)	.074
NFC	−0.0072 (0.0077)	.346
Reward × Demand Level	−0.0427 (0.0128)	.001*
Reward × Trial Block	−0.0391 (0.0128)	.002*
Demand Level × Trial Block	−0.035 (0.0128)	.007*
Reward × NFC	0.0084 (0.0091)	.355
Demand Level × NFC	0.0048 (0.0093)	.606
Trial Block × NFC	0.0006 (0.0095)	.952
Reward × Demand Level × Trial Block	0.0701 (0.0182)	<.0001*
Reward × Demand Level × NFC	−0.0092 (0.0128)	.473
Reward × Trial Block × NFC	−0.0099 (0.0128)	.442
Demand Level × Trial Block × NFC	−0.0042 (0.0128)	.744
Reward × Demand Level × Trial Block × NFC	0.0027 (0.0182)	.881

Note. NFC = need for cognition; *SE* = standard error.
* Denotes significance at $p < .05$ level.

of the marginal value of effort is a general principle (or rule of thumb) that humans and animals abide by, to varying extents (Hsee et al., 2003; Reinagel, 2021), in making real-world effort allocation decisions, but in each task domain, the relationship governing effort increases and performance benefits is learned experientially. Accordingly, we believe the learned insensitivity of control allocation levels to reward incentives (in accordance with demand level) observed here reflects participants' application of this general principle to the specific task domains and operationalizations of demand examined here, which requires that participants have had enough experience to learn the appropriate effort-performance relationships in different demand contexts. On this view, individuals are capable of adjusting effort expenditure in accordance with its marginal utility, but we would only expect to observe that participants tune reward-induced effort increases in accordance with demand context after having learned the specific effort-performance relationship for each demand context in our experiments. Future research should endeavor to understand (a) the extent to which peoples' understanding and application of marginal value of effort is a general phenomenon divorced from specific task domains and (b) how and when an individual's learning (and application) of this general principle unfolds over the life span (Insel et al., 2017; Rodman et al., 2021).

Beyond learning an efficient demand-level and reward-guided effort allocation strategy, we also observed that participants' task switch costs were smaller in high-switch-rate (i.e., high-demand) contexts and their flanker incongruence costs were smaller in contexts with a high proportion of incongruent stimuli, suggesting that participants adjusted their levels of control in accordance with the environmental demand level. These shifts in control allocation mirror previous findings in task-switching (Liu & Yeung, 2020) and flanker (Aben et al., 2017) paradigms. A tacit assumption in these lines of work is that participants must somehow learn to adapt control allocation to the environmental demand level—irrespective of reward incentives—over time. Indeed, we found varying degrees of evidence for this sort of learning in the two experiments, vis-à-vis interactions between trial block and switch/incongruence costs (Tables 2 and 5). We assume that these presumably strategic shifts in reward-induced effort allocation are operating over and above these more general shifts in control mode prompted by environment demand level (and the regression modeling approach we take supports that interpretation). Nonetheless, as little previous work has addressed how individuals learn this sort of control allocation strategy, understanding the nature of learning of these broader control adaptations would also be a fruitful avenue for future research.

Another key question concerns how the brain computes and represents the marginal utility of effort and, furthermore, how this is used to modulate attention and decision-making. One proposed neurocomputational model, the reinforcement metalearner (Silvetti et al., 2018), suggests that cognitive control optimization results from the interplay between mesolimbic dopamine (coding for effort costs and reward-related information), norepinephrine (implementing control), and the medial prefrontal cortex (operating performance monitoring and deciding control signal intensity). This cortical-subcortical circuit directs cognitive control allocation by minimizing the cost of control, maximizing reward, and weighting the control-dependent performance improvement—that is, the marginal utility of effort. From this perspective, cognitive control allocation is boosted only when its resultant performance improvements overcome its intrinsic cost. Marginal utility computation is thus part of an optimization process in which the brain learns to efficiently allocate control, which in turn influences reward-based decision-making (metalearning). Similarly, our experimental result suggests refinements to other reinforcement-learning-based accounts of effort allocation (Holroyd & McClure, 2015; Lieder et al., 2018; Shenhav et al., 2017) to incorporate, over and above the value of effort exertion, peoples' apparent sensitivity to the (presumably learned) relationship between changes in effort investment and changes in performance.

It is also worth noting that in high-demand conditions in the second half of both experiments, participants did not appear to modulate their control levels in accordance with the incentive levels available. In our account, this lack of differentiation between reward levels is a result of learning that exerting additional effort to either reduce switch costs (in Experiment 1) or flanker incongruence costs (in Experiment 2) confers little performance benefits in high-demand conditions. Indeed, inspecting the high-demand condition in Experiment 1, we observed that high-incentive (10-cent) switch costs actually increased over time—resembling the low-incentive switch costs observed in the first half of the experiment— suggesting that participants simply ceased to intensify their control levels in accordance with incentives. However, in Experiment 2, we observed a different pattern of control modulation in the high-demand condition: Participants began to up-regulate their control on low-incentive conditions such that the low- and high-incentive incongruence costs both began to resemble the high-incentive incongruence costs observed early in the experiment. It is possible that, here, the lack of apparent differentiation between these incentive levels resulted from participants intensifying their control levels in the high-demand context over time, irrespective of incentive levels (resulting in lower incongruence costs across both incentive levels), rather than learning to withhold additional control in the high-incentive condition, possibly because maintaining a high level of control with in a demand context is less effortful than continually reallocating control on a trial-by-trial basis. At the same time, flanker-like effects—unlike task switch costs (Rogers & Monsell, 1995)—have been observed to decrease incrementally over time with repeated task exposure (Kelley & Yantis, 2009), suggesting that these changes in control modulations could be operating over and above decreases in "baseline" incongruence costs in the flanker task. Future work should aim to address (a) the generality of these observed effects across task domains and (b) the possibility that individuals might find effortful—and consequently, avoid—frequent reallocation of cognitive control in response to changing incentive levels in favor of maintaining a steady control level.

We also found, in Experiment 1, that individuals low in NFC—who have little intrinsic motivation to expend cognitive effort—appeared to be more sensitive to the marginal utility of effort investment. In other words, the learning effects we observed with respect to reward-modulated switch cost reductions appeared strongest in low-NFC individuals. We previously found that NFC predicts the extent to which individuals exhibit reward-modulated task switch cost reductions (Sandra & Otto, 2018). The present result adds nuance to the idea that subjective effort costs (operationalized by low NFC levels) bear upon not only benefit sensitivity but also sensitivity to the marginal utility of effort investment. Curiously, we did not observe this predictive effect in the flanker task (Experiment 2), which was considerably easier (see overall task performance in Table 1 vs. Table 4). This could suggest that NFC might only exhibit predictive bearing on effort outlay (or changes in effort outlay) in sufficiently difficult tasks. Supporting this idea, a recent study found no evidence for a relationship between NFC levels and flanker effects (Gärtner et al., 2021). At the same time, it is possible that the effective 66-participant sample in Experiment 2 could have been underpowered to detect a statistically meaningful predictive relationship between NFC levels and changes in reward-induced control allocation over time.

One open question concerns whether an individual's NFC level —which has also been interpreted as reflecting an individual's subjective effort cost (Inzlicht et al., 2018)—might have predictive bearing on the individual's default control level in the present tasks. Through the lens of the simple computational model considered here (Figure 1C), it might be the case that low- versus high-NFC individuals employ higher default control levels. Future work could leverage physiological measures like pupillometry to elucidate both default levels and reward-induced changes in effort allocation (van der Wel & van Steenbergen, 2018). To this point, we have also found suggestive evidence that the relationship between task-evoked pupillary responses and task switch costs is stronger in low-NFC individuals (da Silva Castanheira et al., 2021).

Interestingly, our results also dovetail with the recent observation that the efficacy of cognitive effort—the environmental contingency between performance and reward receipt, holding demand constant—is a critical determinant of individuals' reward-induced effort investment in a Stroop-like task (Frömer et al., 2021). When efficacy, which was explicitly signaled to participants before each trial, was high, participants were more inclined to invest effort in the task in accordance with varying incentive level (as evidenced by faster and more accurate responses) compared to when efficacy was low. In contrast to this design, the present experiments did not instruct participants about the marginal utility of effort expenditure, but nonetheless, both our results and Frömer et al.'s result suggest that participants appeared to find reward-guided effort allocation strategies efficient.

It is also worth noting that the task switch rate manipulation in Experiment 1—used to define demand contexts—simultaneously alters both the control demands of task switching (Mayr et al., 2013; Monsell & Mizon, 2006) and the trial-to-trial predictability of the subtask. That is, the 50% (vs. 10%) switch rate in the high-demand condition is more difficult not only because participants need to maintain flexible control in the face of increased task switches but also because the task to be completed in each trial is less predictable. While recent studies have found that individuals both avoid roughly equiprobable task switch rates (Sayalı & Badre, 2019) and experience them as more demanding (Devine & Otto, 2021), the present design does not allow us to disentangle the joint contributions of control demands and unpredictability to the increased demand level of the high-demand context. Nonetheless, our overarching hypotheses about flexible, reward-guided allocation of control over time concern experienced demand levels associated with different contexts but are agnostic to specific source(s) of the demand level differences and presuppose only that one context need be more difficult than another. An open question, then, concerns the relative contributions of control demands and task unpredictability to cost-benefit effort decision-making.

Another plausible alternative explanation for the pattern of results observed in Experiment 2 (the flanker task) could be that participants did not adjust "control parameters" over time but adjusted the extent to which they chose to rely on learned contingencies between task features and responses (Schmidt, 2019). For example, in the high-demand (80% incongruent) context of the flanker task, the flanking arrows were predictive of the response (i.e., with 80% accuracy). Accordingly, participants may exhibit smaller incongruence costs in the high-demand context not because they allocate more control but simply because they benefited more from the mostly predictive flanker feature—that is, they learned that they can simply make the opposite response as indicated by the visually dominating flanker arrows. By the same token, in the high-demand context, participants were also more likely to experience incongruent stimuli and benefit from exact trial repetitions. However, while these (arguably) low-level forms of learning have been demonstrated to contribute substantially to the modulation of congruence effects by proportion congruence in flanker and flanker-like tasks (Braem et al., 2019; Schmidt, 2019), it is unclear why such contingency learning effects would further interact with reward and experiment half in the manner we observed here. Supporting the differential demand interpretation of the two flanker task contexts examined here, a body of literature suggests that individuals rate flanker-like tasks with higher proportions of incongruent trials as more demanding than settings with smaller proportions of incongruent trials (Desender et al., 2017) and prefer to avoid these contexts if given the choice (Schouppe et al., 2014). Nonetheless, useful follow-up experiments would be to investigate whether these learning effects are observed using paradigms with equal opportunities for contingency learning across both demand conditions (e.g., 20% incongruent vs. 80% incongruent) or paradigms without opportunities for contingency learning (Braem et al., 2019; Schmidt, 2019).

It is also worth pointing out that in these two experiments, participants were incentivized to make accurate responses (within a certain response deadline), but we operationalized exertion of cognitive effort using RT-based measures. We should note here that a number of studies examining motivated cognitive control also index control/effort allocation using RT-based measures but, importantly, tie reward incentives—cued before each trial or block of the to-be-completed task—to response accuracy. For example, past studies by other groups (Braem et al., 2012; Capa et al., 2013; Umemoto & Holroyd, 2015) and our own group (da Silva Castanheira et al., 2021; Otto & Vassena, 2021; Vassena et al., 2019) have repeatedly found that even when reward incentives are tied solely to response accuracy, these incentives shape RTs in demanding cognitive tasks. In our view, this body of work suggests that individuals engage in some form of effort-reward calculus even when reward outcomes are not necessarily tied to the performance consequences of these effort modulations. One possible interpretation of these ubiquitous RT incentive effects is that in many of these studies (including the present experiments), participants must nonetheless expend effort to ensure their responses are fast enough to meet RT deadlines (1,500 ms and 1,000 ms in Experiments 1 and 2, respectively) in order to obtain the offered reward amount. On this view, large incentives justify response speeding, while at the same time, engaging in such speeding under smaller incentives would constitute an unjustified, overexertion of control. Moreover, speeding (or, alternatively, invigoration of responses) while either maintaining or improving accuracy is often interpreted as a signature of effort investment (Hübner & Schlösser, 2010; Manohar et al., 2015; Otto & Daw, 2019).

Finally, it is worth noting that the demand contexts were explicitly signaled with border colors (and minimal instructions) in both experiments (Figures 1 and 5), which we believe was necessary here for participants to cumulatively learn the demand contexts over blocks and, accordingly, the marginal utility of effort. It remains an open question whether participants would be able to learn the demand context without explicit cues and evidence the type of learning exhibited in the two experiments reported here. On the basis of prior flanker work finding evidence for implicit learning of control requirements (Ghinescu et al., 2010)—that is, without any instructions describing the relationship between cues signaling proportion of incongruent trials—we might expect that the learning effects observed here might also be possible in environments without clear cues signaling demand contexts.

References

Aben, B., Verguts, T., & Van den Bussche, E. (2017). Beyond trial-by-trial adaptation: A quantification of the time scale of cognitive control. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(3), 509–517. https://doi.org/10.1037/xhp0000324

Abrahamse, E., Braem, S., Notebaert, W., & Verguts, T. (2016). Grounding cognitive control in associative learning. *Psychological Bulletin*, *142*(7), 693–728. https://doi.org/10.1037/bul0000047

Arechar, A. A., & Rand, D. G. (2021). Turking in the time of COVID. *Behavior Research Methods*. Advance online publication. https://doi.org/10.3758/s13428-021-01588-4

Bates, D., & Maechler, M. (2009). lme4: Linear mixed-effects models using S4 classes (Version 1.1-28) [Computer software]. http://CRAN.R-project.org/package=lme4

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, *57*(1), 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

Bijleveld, E., Custers, R., & Aarts, H. (2010). Unconscious reward cues increase invested effort, but do not change speed-accuracy tradeoffs. *Cognition*, *115*(2), 330–335. https://doi.org/10.1016/j.cognition.2009.12.012

Bogdanov, M., Nitschke, J. P., LoParco, S., Bartz, J. A., & Otto, A. R. (2021). Acute psychosocial stress increases cognitive-effort avoidance. *Psychological Science*, *32*, 1463–1475. https://doi.org/10.1177/09567976211005465

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

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*(3), 624–652. https://doi.org/10.1037/0033-295X.108.3.624

Braem, S. (2017). Conditioning task switching behavior. *Cognition*, *166*, 272–276. https://doi.org/10.1016/j.cognition.2017.05.037

Braem, S., Bugg, J. M., Schmidt, J. R., Crump, M. J. C., Weissman, D. H., Notebaert, W., & Egner, T. (2019). Measuring adaptive control in conflict tasks. *Trends in Cognitive Sciences*, *23*(9), 769–783. https://doi.org/10.1016/j.tics.2019.07.002

Braem, S., Verguts, T., Roggeman, C., & Notebaert, W. (2012). Reward modulates adaptations to conflict. *Cognition*, *125*(2), 324–332. https://doi.org/10.1016/j.cognition.2012.07.015

Brehm, J. W., & Self, E. A. (1989). The intensity of motivation. *Annual Review of Psychology*, *40*, 109–131. https://doi.org/10.1146/annurev.ps.40.020189.000545

Burton, S., Knibb, G., & Jones, A. (2021). A meta-analytic investigation of the role of reward on inhibitory control. *The Quarterly Journal of Experimental Psychology*, *74*(10), 1818–1828. https://doi.org/10.1177/17470218211008895

Bustamante, L., Lieder, F., Musslick, S., Shenhav, A., & Cohen, J. (2021). Learning to overexert cognitive control in a Stroop task. *Cognitive, Affective & Behavioral Neuroscience*, *21*(3), 453–471. https://doi.org/10.3758/s13415-020-00845-x

Cacioppo, J. T., Petty, R. E., & Kao, C. F. (1984). The efficient assessment of need for cognition. *Journal of Personality Assessment*, *48*(3), 306–307. https://doi.org/10.1207/s15327752jpa4803_13

Capa, R. L., Bouquet, C. A., Dreher, J.-C., & Dufour, A. (2013). Long-lasting effects of performance-contingent unconscious and conscious reward incentives during cued task-switching. *Cortex*, *49*(7), 1943–1954. https://doi.org/10.1016/j.cortex.2012.05.018

Carver, C. S., & White, T. L. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS Scales. *Journal of Personality and Social Psychology*, *67*(2), 319–333. https://doi.org/10.1037/0022-3514.67.2.319

Chiu, Y.-C., & Egner, T. (2019). Cortical and subcortical contributions to context-control learning. *Neuroscience and Biobehavioral Reviews*, *99*, 33–41. https://doi.org/10.1016/j.neubiorev.2019.01.019

Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS ONE*, *8*(3), Article e57410. https://doi.org/10.1371/journal.pone.0057410

da Silva Castanheira, K., LoParco, S., & Otto, A. R. (2021). Task-evoked pupillary responses track effort exertion: Evidence from task-switching. *Cognitive, Affective & Behavioral Neuroscience*, *21*(3), 592–606. https://doi.org/10.3758/s13415-020-00843-z

Desender, K., Buc Calderon, C., Van Opstal, F., & Van den Bussche, E. (2017). Avoiding the conflict: Metacognitive awareness drives the selection of low-demand contexts. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(7), 1397–1410. https://doi.org/10.1037/xhp0000391

Devine, S., & Otto, A. R. (2021). *Information about task progress modulates cognitive demand avoidance*. PsyArXiv. https://doi.org/10.31234/osf.io/kpy8a

Dreisbach, G., & Haider, H. (2006). Preparatory adjustment of cognitive control in the task switching paradigm. *Psychonomic Bulletin & Review*, *13*(2), 334–338. https://doi.org/10.3758/BF03193853

Duthoo, W., De Baene, W., Wühr, P., & Notebaert, W. (2012). When predictions take control: The effect of task predictions on task switching performance. *Frontiers in Psychology*, *3*, Article 282. https://doi.org/10.3389/fpsyg.2012.00282

Egner, T. (2014). Creatures of habit (and control): A multi-level learning perspective on the modulation of congruency effects, *Frontiers in Psychology*, *5*, Article 1247. https://www.frontiersin.org/article/10.3389/fpsyg.2014.01247

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, *16*(1), 143–149. https://doi.org/10.3758/BF03203267

Fröber, K., & Dreisbach, G. (2016). How sequential changes in reward magnitude modulate cognitive flexibility: Evidence from voluntary task switching. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(2), 285–295. https://doi.org/10.1037/xlm0000166

Fröber, K., & Dreisbach, G. (2017). Keep flexible – Keep switching! The influence of forced task switching on voluntary task switching. *Cognition*, *162*, 48–53. https://doi.org/10.1016/j.cognition.2017.01.024

Frömer, R., Lin, H., Dean Wolf, C. K., Inzlicht, M., & Shenhav, A. (2021). Expectations of reward and efficacy guide cognitive control allocation. *Nature Communications*, *12*(1), Article 1030. https://doi.org/10.1038/s41467-021-21315-z

Gärtner, A., Grass, J., Wolff, M., Goschke, T., Strobel, A., & Strobel, A. (2021). No relation of need for cognition to basic executive functions.

Journal of Personality, *89*(6), 1113–1125. https://doi.org/10.1111/jopy.12639

Gendolla, G. H. E., Wright, R. A., & Richter, M. (2012). Effort intensity: Some insights from the cardiovascular system. In R. M. Ryan (Ed.), *The Oxford handbook of human motivation* (pp. 420–438). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780195399820.013.0024

Ghinescu, R., Schachtman, T. R., Stadler, M. A., Fabiani, M., & Gratton, G. (2010). Strategic behavior without awareness? Effects of implicit learning in the Eriksen flanker paradigm. *Memory & Cognition*, *38*(2), 197–205. https://doi.org/10.3758/MC.38.2.197

Gratton, G., Coles, M. G., & Donchin, E. (1992). Optimizing the use of information: Strategic control of activation of responses. *Journal of Experimental Psychology: General*, *121*(4), 480–506. https://doi.org/10.1037/0096-3445.121.4.480

Hall-McMaster, S., Muhle-Karbe, P. S., Myers, N. E., & Stokes, M. G. (2019). Reward boosts neural coding of task rules to optimize cognitive flexibility. *The Journal of Neuroscience*, *39*(43), 8549–8561. https://doi.org/10.1523/JNEUROSCI.0631-19.2019

Holroyd, C. B., & McClure, S. M. (2015). Hierarchical control over effortful behavior by rodent medial frontal cortex: A computational model. *Psychological Review*, *122*(1), 54–83. https://doi.org/10.1037/a0038339

Hsee, C. K., Yu, F., Zhang, J., & Zhang, Y. (2003). Medium maximization. *Journal of Consumer Research*, *30*(1), 1–14. https://doi.org/10.1086/374702

Hübner, R., & Schlösser, J. (2010). Monetary reward increases attentional effort in the flanker task. *Psychonomic Bulletin & Review*, *17*(6), 821–826. https://doi.org/10.3758/PBR.17.6.821

Hull, C. L. (1943). *Principles of behavior: An introduction to behavior theory*. Appleton-Century.

Insel, C., Kastman, E. K., Glenn, C. R., & Somerville, L. H. (2017). Development of corticostriatal connectivity constrains goal-directed behavior during adolescence. *Nature Communications*, *8*(1), Article 1605. https://doi.org/10.1038/s41467-017-01369-8

Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The effort paradox: Effort is both costly and valued. *Trends in Cognitive Sciences*, *22*(4), 337–349. https://doi.org/10.1016/j.tics.2018.01.007

Kelley, T. A., & Yantis, S. (2009). Learning to attend: Effects of practice on information selection. *Journal of Vision*, *9*(7), Article 16. https://doi.org/10.1167/9.7.16

Kleinsorge, T., & Rinkenauer, G. (2012). Effects of monetary incentives on task switching. *Experimental Psychology*, *59*(4), 216–226. https://doi.org/10.1027/1618-3169/a000146

Kool, W., & Botvinick, M. (2018). Mental labour. *Nature Human Behaviour*, *2*(12), 899–908. https://doi.org/10.1038/s41562-018-0401-9

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General*, *139*(4), 665–682. https://doi.org/10.1037/a0020198

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

Krebs, R. M., Boehler, C. N., & Woldorff, M. G. (2010). The influence of reward associations on conflict processing in the Stroop task. *Cognition*, *117*(3), 341–347. https://doi.org/10.1016/j.cognition.2010.08.018

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13), 1–26. https://doi.org/10.18637/jss.v082.i13

Li, X., & Hsee, C. K. (2021). The psychology of marginal utility. *Journal of Consumer Research, ucaa064*, *48*(1), 169–188. https://doi.org/10.1093/jcr/ucaa064

Lieder, F., Shenhav, A., Musslick, S., & Griffiths, T. L. (2018). Rational metareasoning and the plasticity of cognitive control. *PLoS Computational Biology*, *14*(4), Article e1006043. https://doi.org/10.1371/journal.pcbi.1006043

Liu, C., & Yeung, N. (2020). Dissociating expectancy-based and experience-based control in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, *46*(2), 131–154. https://doi.org/10.1037/xhp0000704

Lorist, M. M., Boksem, M. A. S., & Ridderinkhof, K. R. (2005). Impaired cognitive control and reduced cingulate activity during mental fatigue. *Cognitive Brain Research*, *24*(2), 199–205. https://doi.org/10.1016/j.cogbrainres.2005.01.018

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

Mayr, U. (2006). What matters in the cued task-switching paradigm: Tasks or cues? *Psychonomic Bulletin & Review*, *13*(5), 794–799. https://doi.org/10.3758/BF03193999

Mayr, U., Kuhns, D., & Rieter, M. (2013). Eye movements reveal dynamics of task control. *Journal of Experimental Psychology: General*, *142*(2), 489–509. https://doi.org/10.1037/a0029353

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140. https://doi.org/10.1016/S1364-6613(03)00028-7

Monsell, S., & Mizon, G. A. (2006). Can the task-cuing paradigm measure an endogenous task-set reconfiguration process? *Journal of Experimental Psychology: Human Perception and Performance*, *32*(3), 493–516. https://doi.org/10.1037/0096-1523.32.3.493

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

Otto, A. R., & Vassena, E. (2021). It's all relative: Reward-induced cognitive control modulation depends on context. *Journal of Experimental Psychology: General*, *150*(2), 306–313. https://doi.org/10.1037/xge0000842

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

Reinagel, P. (2021). Rational regulation of water-seeking effort in rodents. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(48), Article e2111742118. https://doi.org/10.1073/pnas.2111742118

Rodman, A. M., Powers, K. E., Insel, C., Kastman, E. K., Kabotyanski, K. E., Stark, A. M., Worthington, S., & Somerville, L. H. (2021). How adolescents and adults translate motivational value to action: Age-related shifts in strategic physical effort exertion for monetary rewards. *Journal of Experimental Psychology: General*, *150*(1), 103–113. https://doi.org/10.1037/xge0000769

Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, *124*(2), 207–231. https://doi.org/10.1037/0096-3445.124.2.207

Sandra, D. A., & Otto, A. R. (2018). Cognitive capacity limitations and need for cognition differentially predict reward-induced cognitive effort expenditure. *Cognition*, *172*, 101–106. https://doi.org/10.1016/j.cognition.2017.12.004

Sayalı, C., & Badre, D. (2019). Neural systems of cognitive demand avoidance. *Neuropsychologia*, *123*, 41–54. https://doi.org/10.1016/j.neuropsychologia.2018.06.016

Schmidt, J. R. (2019). Evidence against conflict monitoring and adaptation: An updated review. *Psychonomic Bulletin & Review*, *26*(3), 753–771. https://doi.org/10.3758/s13423-018-1520-z

Schouppe, N., Demanet, J., Boehler, C. N., Ridderinkhof, K. R., & Notebaert, W. (2014). The role of the striatum in effort-based decision-making in the absence of reward. *The Journal of Neuroscience*, *34*(6), 2148–2154. https://doi.org/10.1523/JNEUROSCI.1214-13.2014

Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a rational and mechanistic account of mental effort. *Annual Review of Neuroscience*, *40*(1), 99–124. https://doi.org/10.1146/annurev-neuro-072116-031526

Silvetti, M., Vassena, E., Abrahamse, E., & Verguts, T. (2018). Dorsal anterior cingulate-brainstem ensemble as a reinforcement meta-learner. *PLoS Computational Biology*, *14*(8), Article e1006370. https://doi.org/10.1371/journal.pcbi.1006370

Spitzer, M., Musslick, S., Shvartsman, M., Shenhav, A., & Cohen, J. D. (2019). Asymmetric switch costs as a function of task strength. *41st Annual Meeting of the Cognitive Science Society, Montreal, Canada* (pp. 1070–1076).

Umemoto, A., & Holroyd, C. B. (2015). Task-specific effects of reward on task switching. *Psychological Research*, *79*(4), 698–707. https://doi.org/10.1007/s00426-014-0595-z

van der Wel, P., & van Steenbergen, H. (2018). Pupil dilation as an index of effort in cognitive control tasks: A review. *Psychonomic Bulletin & Review*, *25*(6), 2005–2015. https://doi.org/10.3758/s13423-018-1432-y

Vassena, E., Deraeve, J., & Alexander, W. H. (2019). Task-specific prioritization of reward and effort information: Novel insights from behavior and computational modeling. *Cognitive, Affective & Behavioral Neuroscience*, *19*(3), 619–636. https://doi.org/10.3758/s13415-018-00685-w

Verguts, T., Vassena, E., & Silvetti, M. (2015). Adaptive effort investment in cognitive and physical tasks: A neurocomputational model. *Frontiers in Behavioral Neuroscience*, *9*, Article 57. https://doi.org/10.3389/fnbeh.2015.00057

Vogel, T. A., Savelson, Z. M., Otto, A. R., & Roy, M. (2020). Forced choices reveal a trade-off between cognitive effort and physical pain. *eLife*, *9*, Article e59410. https://doi.org/10.7554/eLife.59410

Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive, Affective & Behavioral Neuroscience*, *15*(2), 395–415. https://doi.org/10.3758/s13415-015-0334-y

Westbrook, A., Kester, D., & Braver, T. S. (2013). What is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. *PLoS ONE*, *8*(7), Article e68210. https://doi.org/10.1371/journal.pone.0068210

Yamaguchi, M., & Nishimura, A. (2019). Modulating proactive cognitive control by reward: Differential anticipatory effects of performance-contingent and non-contingent rewards. *Psychological Research*, *83*(2), 258–274. https://doi.org/10.1007/s00426-018-1027-2

Yeung, N., & Monsell, S. (2003). Switching between tasks of unequal familiarity: The role of stimulus-attribute and response-set selection. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(2), 455–469. https://doi.org/10.1037/0096-1523.29.2.455

Yu, A. J., Dayan, P., & Cohen, J. D. (2009). Dynamics of attentional selection under conflict: Toward a rational Bayesian account. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(3), 700–717. https://doi.org/10.1037/a0013553

Received May 24, 2021
Revision received February 4, 2022
Accepted February 4, 2022 ■