AMERICAN PSYCHOLOGICAL ASSOCIATION

Journal of Experimental Psychology: General

© 2020 American Psychological Association
ISSN: 0096-3445

2021, Vol. 150, No. 2, 306–313
http://dx.doi.org/10.1037/xge0000842

It's All Relative: Reward-Induced Cognitive Control Modulation Depends on Context

A. Ross Otto
McGill University

Eliana Vassena
Radboud University and Radboud University Medical Center, Nijmegen, the Netherlands

Although people seek to avoid expenditure of cognitive effort, reward incentives can increase investment of processing resources in challenging situations that require cognitive control, resulting in improved performance. At the same time, subjective value is relative, rather than absolute: The value of a reward is increased if the local context is reward-poor versus reward-rich. Although this notion is supported by work in economics and psychology, we propose that reward relativity should also play a critical role in the cost–benefit computations that inform cognitive effort allocation. Here we demonstrate that reward-induced cognitive effort allocation in a task-switching paradigm is sensitive to reward context, consistent with the notion of relative value. Informed by predictions of a computational model of divisive reward normalization, we demonstrate that reward-induced switch cost reductions depend critically upon reward context, such that the same reward amount engenders greater control allocation in impoverished versus rich reward context. Succinctly, these results confirm that reward relativity factors into the value computation driving effort allocation, revealing that motivated cognitive control, like choice, is all relative.

Keywords: cognitive control, reward, switch costs, task-switching, value normalization

Supplemental materials: http://dx.doi.org/10.1037/xge0000842.supp

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

We expend to cognitive effort when it is worth our while—that is, the resource-limited nature of cognitive processing dictates that we invest mental effort in a task only when its benefits outweigh its costs (Kool & Botvinick, 2018). A classic situation requiring effort is task-switching: Repeating the same task is easy, switching to a different task is hard. This task set reconfiguration process slows people down, producing the pervasive *switch cost* (Monsell, 2003). As in other tasks requiring cognitive control, monetary incentives reduce switch costs (Kleinsorge & Rinkenauer, 2012; Sandra & Otto, 2018). This reward-driven facilitation is seen as the consequence of enhanced control allocation (Silvetti, Vassena, Abrahamse, & Verguts, 2018) presumably resulting from a cost-benefit tradeoff (Shenhav et al., 2017; Westbrook & Braver, 2015).

At the same time, a large body of work examining how economic value is represented psychologically and neurally emphasizes the context-dependent nature of valuation—that is, the perceived value of a reward strongly depends on the context in which it is evaluated (Rangel & Clithero, 2012; Tversky & Simonson, 1993). Accordingly, the subjective value of an action or good is increased in a low-value context and decreased in a high-value context (Nieuwenhuis et al., 2005; Rigoli, Friston, & Dolan, 2016; Vlaev, Seymour, Dolan, & Chater, 2009). This sort of relative, versus absolute, valuation explains a number of interesting patterns of choice in value-based decision-making (Bavard, Lebreton, Khamassi, Coricelli, & Palminteri, 2018; Klein, Ullsperger, & Jocham, 2017; Louie, Khaw, & Glimcher, 2013).

Although the idea that a reward's value is considered relative to its context has been considered across the psychology, neuroscience, and economics literatures (Abeler, Falk, Goette, & Huffman, 2011; Elliott, Agnew, & Deakin, 2008; Kahneman & Tversky, 1979; Rangel & Clithero, 2012; Seymour & McClure, 2008), reward relativity effects on higher-order cognitive processes such as cognitive control is surprisingly understudied. Yet, the interaction between motivational effects of incentives upon cognitive control allocation is a central theme in cognitive neuroscience (Botvinick & Braver, 2015; Braver et al., 2014). To this end, a number of computational models make specific proposals concerning how reward information is integrated with the costs of effortful control required to obtain these rewards (Shenhav, Botvinick, & Cohen, 2013; Silvetti et al., 2018; Verguts, Vassena, & Silvetti,

This article was published Online First August 13, 2020.

A. Ross Otto, Department of Psychology, McGill University; Eliana Vassena, Department of Experimental Psychopathology and Treatment, Behavioral Science Institute, Radboud University, and Donders Institute for Brain, Cognition and Behaviour, Radboud University Medical Center, Nijmegen, the Netherlands.

This research was supported by a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, a New Researchers Grant from the Fonds de Recherche du Québec – Nature et Technologies, and a Canada Foundation for Innovation (CFI) Infrastructure Grant to A. Ross Otto. 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/gwrj9/.

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. E-mail: ross.otto@mcgill.ca

2015). However, none of these models considers the relativity of reward value as a critical factor in these cost–benefit computations.

This yields an important but unexplored question in understanding motivated cognitive control: Is reward-driven allocation of cognitive effort similarly dependent on reward context? To examine this question, we manipulate performance-dependent monetary incentives in a task-switching paradigm, in which incentives are offered within broader reward contexts: a low-reward context, containing 1- and 10-cent incentives, and a high-reward context, containing 10- and 19-cent incentives (see Figure 1B). If reward-induced cognitive effort modulations are indeed context-dependent, then we should observe different levels of control allocation—and accordingly, different switch costs—between the two 10-cent reward conditions—because this reward amount is large in a low-reward context, but small in a high-reward context.

Computationally, we can draw predictions for context-dependent reward effects on control allocation using a simple divisive normalization model—a canonical value computation scaling the immediately available reward by the average of available rewards in the current context (Khaw, Glimcher, & Louie, 2017). The consequences of divisive reward normalization are intuitive in the low- and high-reward contexts examined here: the contextual value of 10 cents is markedly smaller in the high-reward context than the low-reward context because it is normalized by a large value—the average of 10 and 19 cents in the high-reward context versus the average of 1 and 10 cents in the low-reward context.

We combined value normalization with an established task-switching model (Yeung & Monsell, 2003), wherein a top-down control input signal reduces response interference between tasks, yielding predicted response time (RT) switch costs in each of the reward conditions and contexts (see Figure 2). These predictions are generated from fits of this task-switching model to participant task-switching behavior under conditions of no reward incentives, and with the assumption that additional control input is provided in proportion either to absolute rewards (in the case of the No Value Normalization model) or normalized rewards (in the case of the Divisive Normalization model; see Experiment 1 Methods for details).

Intuitively, the baseline (No Value Normalization) model, which applies control input in accordance with absolute rewards (Figure 2A), predicts that increasing rewards should result in monotonically smaller task switch costs, with equal switch costs predicted across the two locally identical available rewards (10 cents). However, across the two locally identical available rewards (10 cents), the Divisive Normalization model predicts markedly

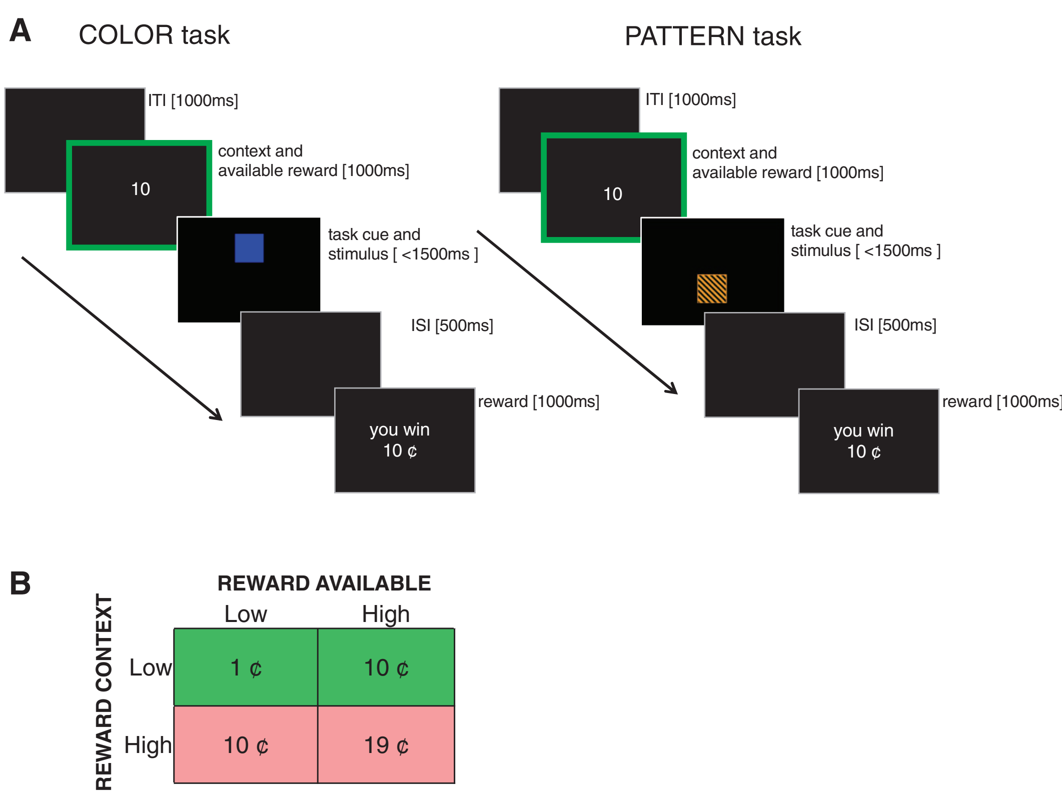

Figure 1. (A) Task-switching paradigm. Subjects either indicated the color (blue or orange) or the pattern (stripes or solid) of a square, depending on its location (top or bottom) on the display. The reward available for making a correct response was displayed before the stimulus and was accompanied by a green (dark gray) or red (light gray) border, indicating a low- or high-reward context respectively. Note that these contexts were only explicitly signaled in Experiment 1. (B) Definition of reward contexts as a function of immediately available rewards. See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

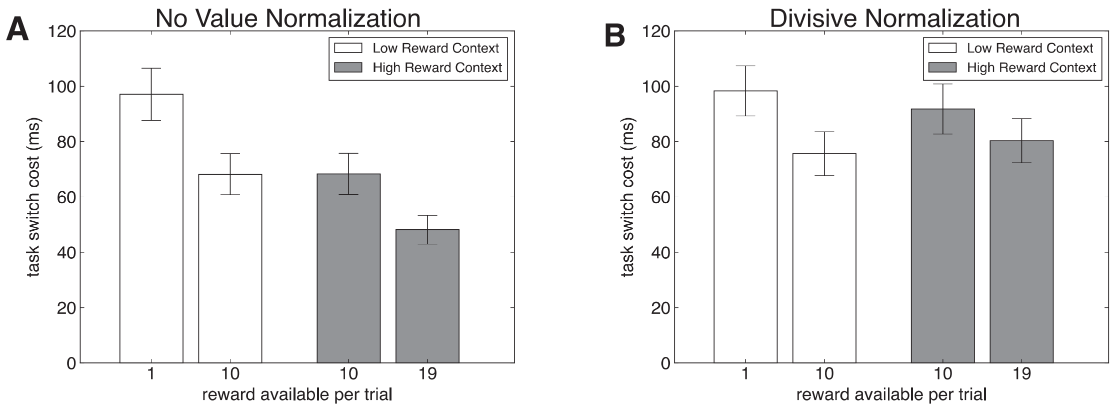

Figure 2. Model-predicted task-switch costs (expressed as the difference between task switch and task repeat RTs) as a function of reward incentive and reward context condition in the case of No Value Normalization (Panel A) and Divisive Normalization (Panel B).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

smaller switch costs in the low-reward context (1 and 10 cents) than in the high-reward context (10 and 19 cents). Further, this account predicts that switch costs in the 19-cent condition—the largest possible incentive in absolute terms—should not result in the lowest switch costs observed across the four conditions because, again, value normalization holds that its subjective value is modulated by its (high) reward context.

Informed by these model predictions, we test the reward-relativity hypothesis of control allocation in two experiments. First, as an initial demonstration, we examine whether reward-induced cognitive control modulations are sensitive to reward contexts that are explicitly signaled to participants (see Figure 1A). Second, we examine whether this context-dependence extends to more ecologically valid circumstances where the environment's underlying reward context can only be inferred from the recent history of experienced available rewards (Khaw et al., 2017). To foreshadow, we find that in both situations, reward-induced cognitive effort engagement is strongly modulated by reward context.

Experiment 1

In Experiment 1, we examined whether the effects of local reward incentives upon effort mobilization are modulated by explicitly signaled reward contexts.

Method

Participants. We recruited 100 U.S. participants on Amazon's Mechanical Turk (Crump, McDonnell, & Gureckis, 2013), who were paid a fixed amount ($3 USD) plus a bonus contingent on their task performance, ranging from $1–3 USD. This sample size, which we have adopted as a standard for online studies employing within-subject designs, ensures adequate statistical power to detect meaningful differences across conditions, while at the same time protecting against false positives results or inflated effect sizes which can result from small samples. Participants provided informed consent in accordance with the McGill University Research Ethics Board. We excluded the data of 15 participants who failed to perform either task with an accuracy of at least 75% on task repetitions and three participants who missed 10 or more response deadlines, leaving 82 participants in the final analyses. These criteria were developed on the basis of a series of earlier pilot studies using an identical task-switching paradigm. The effect of reward remained significant with the inclusion of these participants.

Task-switching paradigm. In a preliminary phase, participants completed 80 trials of a task-switching paradigm in the absence of incentives to gain familiarity with the task. On each trial of this paradigm, following Otto and Daw (2019), a box appeared on screen and participants needed either to report whether the box that appeared on screen was blue or orange (the "COLOR" task) or whether the box's fill was solid or striped (the "PATTERN" task). Critically, the position of the box on the screen (lower half vs. upper half, counterbalanced across subjects) indicated which subtask the subject was to perform (Figure 1A). Across both subtasks, 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. 50% of trials repeated the previous subtask. Participants had 1,500 ms to respond, at which point they received feedback indicating they made a correct response.

Following this preliminary phase, subjects began a reward phase, where a number at the beginning of each trial signaled the reward available for correct responses. The reward context was indicated by a colored frame around the available reward. In the low-reward context (green), possible rewards were 1 or 10 cents, and in the high-reward context (red), possible rewards were 10 or 19 cents (Figure 1B). Each reward context block was 32 trials long, and consisted of 4 'miniblocks' composed of 8 consecutive trials of high (10 or 19 cents, for the low- and high-reward contexts, respectively) or low reward (1 or 10 cents, for the low- and high-reward contexts, respectively) amounts. The order of reward miniblocks was randomized within reward contexts, whose orders were randomized across participants. Participants completed eight reward context blocks, totaling 256 trials. To ensure that task-switching behavior reflected the reward context, our analyses omitted the first four trials of each reward context block, and to ensure switch costs were computed with respect to both task switches and repetitions for the same reward level, the first trial of each reward miniblock was also omitted. RTs were log-transformed to remove skew before being submitted to ANOVAs. Repeated-measures ANOVAs were computed, taking the context-relative available reward—*low*, corresponding to either 1 (low context) or 10 cents (high context), or *high*, corresponding to either 10 (low context) or 19 cents (high context), as the reward factor, and reward context (low vs. high) as within-subjects factors.

Computational model of task-switching and value normalization. We implemented an established model of task switching, described in detail by Yeung and Monsell (2003), previously used to model the effects of task priming and control upon RTs in task-switching paradigms. In short, the model assumes that task responses are the result of competition between the two subtasks, which are activated in accordance with the default 'task strength' as well as task priming levels. When competition between the two subtasks is strong, owing in part to task priming from the previous trial, the consequent activation levels of the two subtasks are similar, in turn, resulting in long resolution times (which translate to RTs) for task switches, compared with task repetitions, for which activation of the most recently performed task, via priming, more strongly activates the current, to-be-performed subtask. Critically, providing additional control input—a parameter that increases the activation level of the relevant task, instantiating a form of top-down control—can counteract the task priming responsible for task switching costs.

We fit this model to participants' task switch and task repetition RTs in the preliminary phase of the experiment—in the absence of rewards—so that we could make predictions about reward incentive effects on the basis of participants' presumed 'default' level of control input. To do this, we fit the model's four parameters (task strength, task priming, control input, and threshold) to each individual's RT distributions for task repetitions and switches using the Weighted Least Squares Fitting Method, with five RT quantiles, as detailed by Ratcliff and Tuerlinckx (2002), originally used to fit drift diffusion models to experimental data. The mean best-fitting parameter values (across participants) were 0.042, 0.031, 385.68, and 0.058 for task strength, task priming, threshold, and control input, respectively, which resulted in mean per-participant sum of squared error of 94.77 (SD = 85.41). See Figure S1 in online supplemental materials for a comparison of model-predicted and observed switch costs.

The resultant parameter values for each participant were then used to simulate the effect of trial-level reward incentives, which we operationalized as a marginal increase in control input over and above the participant's default control level, $control_{prelim}$. Following common formulations of divisive normalization (Khaw et al., 2017), we calculate the control input on a particular trial as:

$$control_{reward,context} = control_{prelim} + \frac{reward}{1 + context} \cdot 0.05$$

where *reward* denotes the reward available on the current trial and *context* denotes the mean of the reward available for the current context (5.5 and 14.5 in the low and high contexts, respectively). This normalized control input (Figure 2A) was multiplied by a constant (0.05) to scale the units of normalized reward to appropriate units of control input in the task-switching model. To generate predictions for a baseline model with no reward normalization (Figure 2A), we simply used the raw reward amount, scaled by a constant, as control input to the model. We report model-predicted task switch costs for the two models, as a function of reward amount and context, averaging across the 82 simulated participants, in Figure 2A and 2B.

Results and Discussion

Task performance. As expected, task switches (M = 830.50, SD = 202.43) engendered significantly slower RTs than task repetitions (M = 764.61, SD = 202.35), $F(1, 81) = 482.3$, $p < .0001$, $\eta_p^2 = 0.026$, mirroring the RT costs typically observed in task-switching paradigms (Monsell, 2003). Overall accuracy was quite high across task switches (M = 0.901, SD = 0.085) and repetitions (M = 0.932, SD = 0.070; see Table 1), though switches were significantly less accurate, $F(1, 81) = 50.36$, $p < .0001$, $\eta_p^2 = 0.383$. Following previous work, our analyses focus on RTs for correct responses (Sandra & Otto, 2018; Yeung & Monsell, 2003).

Available reward and reward context effects on performance. Figure 3 depicts task switch costs, calculated as the difference of each subjects' median correct RT for task switch and task repetitions, averaged according to reward available (1 vs. 10 cents) and the reward context (low- vs. high-reward context). Overall, across reward contexts, we observed a significant main effect of reward amount on switch costs, $F(1, 81) = 4.27$, $p = .039$, $\eta_p^2 = 0.017$. Post-hoc tests revealed that this reward effect was significant within the low-reward context (i.e., smaller switch costs in the 10- vs. 1-cent conditions), ($t(81) = 2.14$, $p = .0354$) but did not reach statistical significance within the high-reward context (i.e., smaller switch costs in the 19- vs. 10-cent conditions), ($t(81) = 1.892$, $p = .0619$). This reward-induced switch cost reduction is in line with the finding that reward incentives increase allocation of cognitive control resources (Westbrook & Braver, 2015). We found no significant interaction between reward amount (relative to the context), $F(1, 81) = 0.01$, $p = .922$, $\eta_p^2 = 0.000$) nor main effect of reward context itself, $F(1, 81) = 0.081$, $p = .776$. Examining RTs overall, irrespective of task repetitions or switches, we did not observe significant effects of reward context, $F(1, 81) = 1.537$, $p = .216$, $\eta_p^2 = 0.006$, or reward amount, $F(1, 81) = 3.522$, $p = .0620$, $\eta_p^2 = 0.014$.

In light of the divergent predictions made by the baseline model and the divisive normalization models (see Figure 2), we were particularly interested in the comparison of switch costs observed when the available rewards were 10 cents across the low-reward

Table 1
Average Median RTs for Correct Repeat and Switch Trials as a Function of Reward Context and Reward Amount on Offer for Experiment 1

Reward amount	Repeat RT (SD)	Switch RT (SD)
Low reward context		
1 cent	737.86 (119.34)	812.92 (125.51)
10 cents	758.53 (109.60)	817.38 (124.16)
High reward context		
10 cents	742.48 (108.88)	815.45 (132.76)
19 cents	741.43 (102.033)	806.34 (124.02)

Note. RT = response time (ms).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

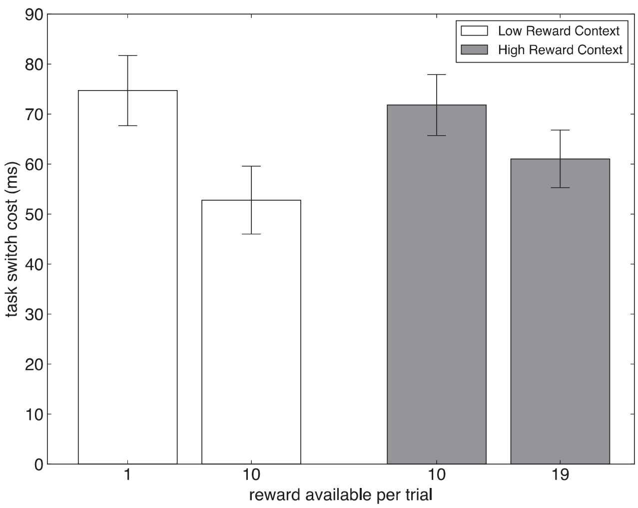

Figure 3. Task-switch costs (expressed as the difference between median task switch RTs and task repetition RTs) as a function of available reward and reward context condition in Experiment 1. Of note, switch costs are smaller in the 10-cent incentive condition within the low-reward context than the 10-cent incentive condition within the high-reward context. Error bars denote standard error of the mean.

and high-reward contexts: If locally available rewards are normalized by the larger reward context in which they occur, then we should observe lower switch costs in the low-reward context than in the high-reward context. Accordingly, examining switch costs on only these 10-cent trials, we observed a significant effect of context upon switch costs, $F(1, 81) = 7.89$, $p = .016$, $\eta_p^2 = 0.069$, suggesting that effort allocation engendered by this reward amount is dependent on reward context. Further in line with the normalization account, the switch cost observed in the 19-cent condition—the largest reward amount available in absolute terms—was not significantly smaller than either of the 10-cent conditions (contrast $ps > 0.10$). That is, the observed patterns of reward-induced switch cost modulations favor the predictions of a divisive reward normalization account (Figure 2B) over a baseline model with no reward normalization (Figure 2A).

Experiment 2

Experiment 2 examined the generality of this context-dependence where, rather than being explicitly signaled, the reward context is defined by the participant's recently experienced history of available rewards.

Method

Participants and design. With the exception of the unsignaled reward contexts, described below, this experiment closely followed the design of Experiment 1. We recruited 100 AMT participants, who were paid a fixed amount ($3 USD) plus a bonus contingent on their decision task performance, ranging from $1–3 USD. Participants provided informed consent in accordance with the McGill University Research Ethics Board. Because of a technical issue, the data of seven participants were lost. Applying Experiment 1's exclusion criteria, we excluded 12 participants for missing response deadlines and another 10 for low accuracy, leaving 71 participants in the final analysis. These exclusions did not alter the significance of the effect of reward on task switch costs.

Experiment 2 only presented the available reward before each stimulus (1, 10, and 19 cents), without explicit indication of the reward contexts (i.e., colored frames in Figure 1A). The unsignaled reward contexts (high vs. low; as in Experiment 1) were 64 trials long, divided into 8-trial-long reward miniblocks (Figure 4A; dashed line). We pseudorandomized the reward block orders to ensure that participants experienced roughly equal distributions of reward levels and reward contexts.

Data analysis. Following previous work (Guitart-Masip, Beierholm, Dolan, Duzel, & Dayan, 2011; Khaw et al., 2017), we estimated trial-by-trial reward context as a recency-weighted av-

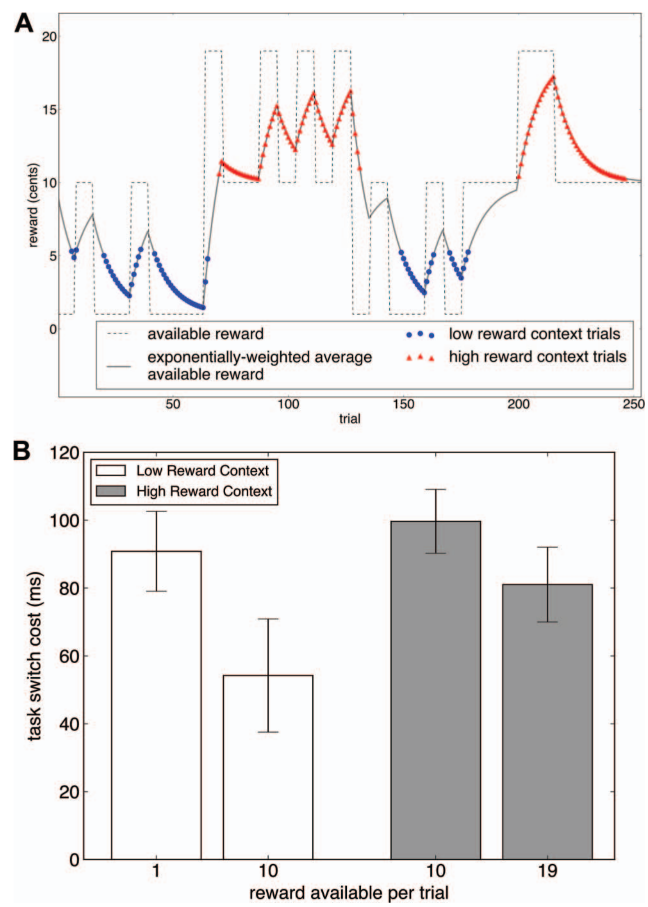

Figure 4. (A) Trial-to-trial reward incentives used in Experiment 2. A participant's experienced-based reward context (solid gray line) is calculated as a recency-weighted average of past available rewards, which were manipulated blockwise (dashed line). We analyzed the switch costs for low- and high- context trials (blue circles and red triangles, respectively), defined by a tertile split upon this recency-weighted average. (B) Task-switch costs as a function of the available reward and unsignaled reward context in Experiment 2. Here the contextual value of the 10 cent incentive brings about larger switch cost reductions in the low-reward versus high-reward contexts. Error bars denote standard error of the mean. See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

erage of past available rewards (solid line in Figure 4A), updating the average reward rate $\bar{r}_{t+1}$ in accordance with the current reward r_t and the previous average reward estimate:

$$\bar{r}_{t+1} = \bar{r}_t + \alpha(r_t - \bar{r}_t)$$

where α is a learning rate parameter set to a value of 0.1, a learning rate for which there is strong evidence in analyses of choice (Eldar & Niv, 2015; Otto, Fleming, & Glimcher, 2016). We then classified trials as belonging to the high- or low- reward context based on a tertile split performed on this calculated average reward (Figure 4A) context, analyzing only the top and bottom tertiles to mitigate against uncertainty about the current reward context in average reward contexts close to the global central tendency (10¢). To ensure that task-switching behavior was stable within contexts and primary reflected the participant's learned reward context, we omitted the first 16 trials of each 64-trial-long reward context block.

Results and Discussion

Mirroring the task switch costs observed in Experiment 1 for correct responses, we observed significantly slower RTs on task switches (M = 817.95, SD = 212.59) than on task repetitions (M = 742.75, SD = 211.35; see Table 2), $F(1, 75) = 374.7$, $p < 0.0001$, $\eta_p^2 = 0.051$. Accuracy was significantly higher for task repetitions (M = 0.939, SD = 0.239) than for task switches (M = 0.907, SD = 0.290), $F(1, 75) = 35.61$, $p < .0001$, $\eta_p^2 = 0.322$. We observed a significant main effect of reward available, $F(1, 75) = 4.562$, $p = .0237$, $\eta_p^2 = 0.018$, upon RT switch costs, but no significant main effect of experienced reward context, $F(1, 75) = 0.014$, $p = .904$, $\eta_p^2 = 0.002$—calculated as the upper and lower tertiles of the recency-weighted average of past rewards (Figure 4A)—nor interaction between reward incentive and experienced reward context, $F(1, 75) = 1.716$ $p = .198$; see Figure 4B.

Of particular interest were the reward-induced switch-cost modulations in locally identical reward incentive conditions across the two experienced reward contexts. Directly comparing the two 10-cent conditions, we observed significantly smaller task switch costs in the low reward context (M = 55.18, SD = 104.86) than in the high reward context (M = 99.34, SD = 79.50), $F(1, 75) = 14.63$, $p < .0001$, 0.147, $\eta_p^2 = 0.147$. And again, following Experiment 1 and the predictions of the normalization account, the switch costs observed in the 19-cent condition were no smaller than either 10-cent conditions (contrast ps $> .255$). Finally, an analysis of overall RTs—irrespective of task switches or repetitions—revealed no significant effect of reward context, $F(1, 75) = 0.871$, $p = .781$; $\eta_p^2 = 0.000$, nor reward amount, $F(1, 75) = 0.831$, $p = .363$, $\eta_p^2 = 0.002$. These results indicate that even when reward contexts are learned, rather than explicitly signaled, reward-induced control modulations operate in a context-dependent manner.

Table 2
Average Median RTs for Correct Repeat and Switch Trials as a Function of Reward Context and Reward Amount on Offer for Experiment 2

Reward amount	Repeat RT (SD)	Switch RT (SD)
Low reward context		
1 cent	723.05 (134.41)	825.07 (169.28)
10 cents	744.15 (148.18)	811.17 (131.22)
High reward context		
10 cents	728.05 (137.27)	829.99 (147.04)
19 cents	737.55 (139.58)	829.41 (158.47)

Note. RT = response time (ms).

General Discussion

Up to now, investigations of reward-induced cognitive effort allocation have treated reward incentives exclusively as an absolute quantity—either by examining the effect of incentives versus the absence of incentives, or by parametrically manipulating incentive values (Chiew & Braver, 2013; Chong et al., 2017; Vassena, Deraeve, & Alexander, 2019). Our results bridge previous work from economics and psychology literatures by revealing a striking reward context relativity of effort allocation: the increase in control—measured by task switch cost reductions—brought about by large (vs. small) reward incentives critically depends on the wider context in which the rewards are situated. These contextual modulations of value occur both when the reward context is explicitly signaled (Experiment 1) and in a more ecologically valid setting, where reward context was covertly manipulated block-wise, and needed to be learned experientially from the environment (Experiment 2). In both cases, the value of a reward—and its potential to incentivize control allocation—hinged on the current reward context.

A consequence of contextual reward relativity is that, across two locally equivalent incentive situations, effort investment levels—as operationalized by task switch costs[1]—differ strikingly depending on the context. Similar value relativity is pervasive in economic choices in which explicit choices are made between options with described value (Rigoli et al., 2016; Vlaev et al., 2009). Here we demonstrate that this mechanism generalizes to reward-incentivized control allocation (Kool & Botvinick, 2018), revealing that values in cognitive control, like in decision-making, are not stable quantities considered in isolation but rather, depend on the reward context.

It is worth noting that, across both experiments, we did not observe general speeding effects brought about by the reward context or the immediate reward amount, suggesting against the possibility that these reward contexts engendered motivational vigor effects as a result of increased average reward rates (e.g., Guitart-Masip et al., 2011; Otto & Daw, 2019), or that immediately available rewards engendered nonspecific performance or speeding effects. Instead, these reward (and context) effects manifested specifically in task switch costs, suggesting that the apparent reward normalization effects observed here stem from modulations in in cognitive control levels.

[1] It is worth noting that, in this paradigm, because the to-be-performed subtask was denoted by location of the stimulus (i.e., top versus bottom of the display), it is possible that the observed switch costs might reflect, in addition to task set reconfiguration processes, a spatial attention reallocation, because every task switch also requires subjects to attend to a different location than the location cued on the previous trial. Nonetheless, the reduction of attentional reallocation costs is also believed to be effortful (Sarter, Gehring, & Kozak, 2006).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Our findings speak to an interesting line of work revealing that the relative balance between different control modes—stability versus flexibility—are 66 as well as by contextual factors (Fröber, Raith, & Dreisbach, 2018). For example, rewards generally increase stability (Fröber & Dreisbach, 2014), but the prospect of increased reward can also promote flexibility. Interestingly, most of these studies focus on varying the frequency of task switches in a forced choice context (where participants perform all tasks they receive in a fixed order), voluntary choice context (where participants can choose whether to switch tasks), and a combination of forced and voluntary switches. These experiments show that when participants perform many forced switch trials, they also tend to switch frequently in voluntary trials (when allowed to choose to switch or repeat the same task). This work highlights the importance of contextual effects on control *balance*. Taking a complementary perspective, our study highlights the general importance of contextual effects on control *allocation*. Critically, we demonstrate that reward incentive itself (rather than varying control demand) can produce contextual effects, particularly showing that trial-by-trial reward incentive do not impact control depending on absolute value. Rather, the incentive impact is contextualized by average reward context. This indicates that reward context is a crucial quantity that is taken into account when computing the value of allocating control resources. The trade-off between stable and flexible modes of control may result from the integrative computation of value per each control mode, calculated based on locally and globally varying features of the task at hand such as trial-by-trial reward, local reward context, and demand. Similar to a simple cost-benefit integration (Shenhav et al., 2013; Silvetti et al., 2018; Verguts et al., 2015), this computation may drive control allocation toward one or the other mode in an emergent fashion.

To further quantify the observed effects, the patterns of reward-induced effort modulations observed in Experiment 1 were predicted, qualitatively, by a simple computational model that assumes that subjective values of available reward amounts are contextually modulated using a divisive normalization scheme which scales value by the average reward available in a given context. Indeed, divisive normalization explains a number of behavioral and neural phenomena across scales of sensory processing (Carandini & Heeger, 2012), as well as subjective valuation in decision-making (Khaw et al., 2017; Louie et al., 2013). Here we provide an initial demonstration that a reward normalization mechanism, more generally, could provide a satisfactory explanation of these observed context effects in reward-induced cognitive control. It is worth noting that other mathematical forms of contextual value normalization—for example, subtractive normalization (Rigoli et al., 2016) and range normalization (Padoa-Schioppa, 2009)—have been put forth to explain context dependence in neural representations of value and choice behavior. Future work should be devoted to adjudicating between these more specific forms of reward normalization as they apply to reward-induced cognitive control modulations.

More generally, these findings both corroborate and extend a large body of behavioral and neural findings supporting the notion that values in decision making are not stable quantities considered in isolation, but rather, depend on value of available outcomes (Rangel & Clithero, 2012; Rigoli et al., 2016). To this point, a central tenet in behavioral economics is prospect theory, which posits that the subjective utility of an outcome is computed with respect to a reference (Kahneman & Tversky, 1979). Although this sort of reference dependence in effort allocation can be viewed, locally, as irrational or inconsistent, the apparent value normalization seen in effort decision-making here could be adaptive in a dynamic environment (Rigoli, 2019; Tversky & Simonson, 1993).

References

Abeler, J., Falk, A., Goette, L., & Huffman, D. (2011). Reference points and effort provision. *The American Economic Review, 101,* 470–492. http://dx.doi.org/10.1257/aer.101.2.470

Bavard, S., Lebreton, M., Khamassi, M., Coricelli, G., & Palminteri, S. (2018). Reference-point centering and range-adaptation enhance human reinforcement learning at the cost of irrational preferences. *Nature Communications, 9,* 4503–4505.

Botvinick, M., & Braver, T. (2015). Motivation and cognitive control: From behavior to neural mechanism. *Annual Review of Psychology, 66,* 83–113. http://dx.doi.org/10.1146/annurev-psych-010814-015044

Braver, T. S., Krug, M. K., Chiew, K. S., Kool, W., Westbrook, J. A., Clement, N. J., . . . the MOMCAI group. (2014). Mechanisms of motivation-cognition interaction: Challenges and opportunities. *Cognitive, Affective & Behavioral Neuroscience, 14,* 443–472. http://dx.doi.org/10.3758/s13415-014-0300-0

Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience, 13,* 51–62. http://dx.doi.org/10.1038/nrn3136

Chiew, K. S., & Braver, T. S. (2013). Temporal dynamics of motivation-cognitive control interactions revealed by high-resolution pupillometry. *Frontiers in Psychology, 4,* 15. http://dx.doi.org/10.3389/fpsyg.2013.00015

Chong, T. T-J., Apps, M., Giehl, K., Sillence, A., Grima, L. L., & Husain, M. (2017). Neurocomputational mechanisms underlying subjective valuation of effort costs. *PLoS Biology, 15,* e1002598. http://dx.doi.org/10.1371/journal.pbio.1002598

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,* e57410. http://dx.doi.org/10.1371/journal.pone.0057410

Eldar, E., & Niv, Y. (2015). Interaction between emotional state and learning underlies mood instability. *Nature Communications, 6,* 6149. http://dx.doi.org/10.1038/ncomms7149

Elliott, R., Agnew, Z., & Deakin, J. F. W. (2008). Medial orbitofrontal cortex codes relative rather than absolute value of financial rewards in humans. *European Journal of Neuroscience, 27,* 2213–2218. http://dx.doi.org/10.1111/j.1460-9568.2008.06202.x

Fröber, K., & Dreisbach, G. (2014). The differential influences of positive affect, random reward, and performance-contingent reward on cognitive control. *Cognitive, Affective, & Behavioral Neuroscience, 14,* 530–547.

Fröber, K., Raith, L., & Dreisbach, G. (2018). The dynamic balance between cognitive flexibility and stability: The influence of local changes in reward expectation and global task context on voluntary switch rate. *Psychological Research, 82,* 65–77.

Guitart-Masip, M., Beierholm, U. R., Dolan, R., Duzel, E., & Dayan, P. (2011). Vigor in the face of fluctuating rates of reward: An experimental examination. *Journal of Cognitive Neuroscience, 23,* 3933–3938. http://dx.doi.org/10.1162/jocn_a_00090

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291. http://dx.doi.org/10.2307/1914185

Khaw, M. W., Glimcher, P. W., & Louie, K. (2017). Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences of the United States of America, 114,* 12696–12701. http://dx.doi.org/10.1073/pnas.1715293114

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Klein, T. A., Ullsperger, M., & Jocham, G. (2017). Learning relative values in the striatum induces violations of normative decision making. *Nature Communications, 8,* 16033.

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

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

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences of the United States of America, 110,* 6139–6144. http://dx.doi.org/10.1073/pnas.1217854110

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

Nieuwenhuis, S., Heslenfeld, D. J., Alting von Geusau, N. J., Mars, R. B., Holroyd, C. B., & Yeung, N. (2005). Activity in human reward-sensitive brain areas is strongly context dependent. *NeuroImage, 25,* 1302–1309.

Otto, A. R., & Daw, N. D. (2019). The opportunity cost of time modulates cognitive effort. *Neuropsychologia, 123,* 92–105. http://dx.doi.org/10.1016/j.neuropsychologia.2018.05.006

Otto, A. R., Fleming, S. M., & Glimcher, P. W. (2016). Unexpected but incidental positive outcomes predict real-world gambling. *Psychological Science, 27,* 299–311. http://dx.doi.org/10.1177/0956797615618366

Padoa-Schioppa, C. (2009). Range-adapting representation of economic value in the orbitofrontal cortex. *The Journal of Neuroscience, 29,* 14004–14014. http://dx.doi.org/10.1523/JNEUROSCI.3751-09.2009

Rangel, A., & Clithero, J. A. (2012). Value normalization in decision making: Theory and evidence. *Current Opinion in Neurobiology, 22,* 970–981. http://dx.doi.org/10.1016/j.conb.2012.07.011

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review, 9,* 438–481. http://dx.doi.org/10.3758/BF03196302

Rigoli, F. (2019). Reference effects on decision-making elicited by previous rewards. *Cognition, 192,* 104034. http://dx.doi.org/10.1016/j.cognition.2019.104034

Rigoli, F., Friston, K. J., & Dolan, R. J. (2016). Neural processes mediating contextual influences on human choice behaviour. *Nature Communications, 7,* 12416. http://dx.doi.org/10.1038/ncomms12416

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. http://dx.doi.org/10.1016/j.cognition.2017.12.004

Sarter, M., Gehring, W. J., & Kozak, R. (2006). More attention must be paid: The neurobiology of attentional effort. *Brain Research Reviews, 51,* 145–160. http://dx.doi.org/10.1016/j.brainresrev.2005.11.002

Seymour, B., & McClure, S. M. (2008). Anchors, scales and the relative coding of value in the brain. *Current Opinion in Neurobiology, 18,* 173–178. http://dx.doi.org/10.1016/j.conb.2008.07.010

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron, 79,* 217–240. http://dx.doi.org/10.1016/j.neuron.2013.07.007

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,* 99–124. http://dx.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,* e1006370. http://dx.doi.org/10.1371/journal.pcbi.1006370

Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science, 39,* 1179–1189. http://dx.doi.org/10.1287/mnsc.39.10.1179

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,* 619–636. http://dx.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,* 57. http://dx.doi.org/10.3389/fnbeh.2015.00057

Vlaev, I., Seymour, B., Dolan, R. J., & Chater, N. (2009). The price of pain and the value of suffering. *Psychological Science, 20,* 309–317. http://dx.doi.org/10.1111/j.1467-9280.2009.02304.x

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

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,* 455–469. http://dx.doi.org/10.1037/0096-1523.29.2.455

Received September 29, 2019
Revision received April 21, 2020
Accepted April 21, 2020 ■

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.