AMERICAN PSYCHOLOGICAL ASSOCIATION

Journal of Experimental Psychology: General

ISSN: 0096-3445

https://doi.org/10.1037/xge0001561

Proximity to Rewards Modulates Parameters of Effortful Control Exertion

Sean Devine[1], Mathieu Roy[1], Ulrik Beierholm[2], and A. Ross Otto[1]
[1] Department of Psychology, McGill University
[2] Department of Psychology, Durham University

The now-classic goal-gradient hypothesis posits that organisms increase effort expenditure as a function of their proximity to a goal. Despite nearly a century having passed since its original formulation, goal-gradient-like behavior in human cognitive performance remains poorly understood: Are we more willing to engage in costly cognitive processing when we are near, versus far, from a goal state? Moreover, the computational mechanisms underpinning these potential goal-gradient effects—for example, whether goal proximity affects fidelity of stimulus encoding, response caution, or other identifiable mechanisms governing speed and accuracy—are unclear. Here, in two experiments, we examine the effect of goal proximity, operationalized as progress toward the completion of a rewarded task block, upon task performance in an attentionally demanding oddball task. Supporting the goal-gradient hypothesis, we found that participants responded more quickly, but not less accurately, when rewards were proximal than when they were distal. Critically, this effect was only observed when participants were given information about goal proximity. Using hierarchical drift diffusion modeling, we found that these apparent goal-gradient performance effects were best explained by a collapsing bound model, in which proximity to a goal reduced response caution and increased information processing. Taken together, these results suggest that goal gradients could help explain the oft-observed fluctuations in engagement of cognitively effortful processing, extending the scope of the goal-gradient hypothesis to the domain of cognitive tasks.

Public Significance Statement

It is well known that humans and animals alike tend to work harder as they near a goal. Whether it be a hungry rat moving closer to a food reward or a runner sprinting the final kilometer of a race, organisms appear to intensify their effort as a function of their proximity to a goal. But does the same principle apply in purely mental tasks—for example, when writing an exam, or doing one's taxes? And if so, how does behavior change? In these studies, we examine whether proximity to a goal affects a person's willingness to exert mental effort in a simple, but cognitively demanding task. Consistent with the established goal-gradient hypothesis, we find that participants intensified their level of cognitive effort—as indexed by their response speed and ability to correctly respond—closer to a goal (vs. further away). Using computational modeling, we found that while participants processed information more efficiently near a goal, they were also less cautious in their decision making. Taken together, our results extend past findings about the effects of goal proximity on effort to the domain of purely cognitive tasks.

Keywords: goal-gradient hypothesis, cognitive effort, drift diffusion model

Supplemental materials: https://doi.org/10.1037/xge0001561.supp

Sean Devine https://orcid.org/0000-0002-0445-2763

Some elements of this work were previously presented as part of a poster and flash talk at the 2022 Mental Effort Workshop and at a Nanosymposium at the 2022 Society for Neuroscience Meeting. This work was supported by a grant (awarded to A. Ross Otto) and a fellowship (awarded to Sean Devine) from the Natural Sciences and Engineering Research Council of Canada. Work on the project was facilitated by a Royal Society International Exchange Fellowship (IES\R3\203148) to Ulrik Beierholm. This research was enabled in part by support provided by Calcul Québec (https://www.calculquebec.ca/), the Digital Research Alliance of Canada (https://alliancecan.ca/en), and Advanced Research Computing at Durham University (https://www.dur.ac.uk/arc/), which provided the computing clusters used to run the computational models reported in this article. The authors declare no conflicts of interest. All task codes, materials, raw data, analysis scripts, results, and compiled models are openly available at https://github.com/seandamiandevine/EffortProgress2.

Sean Devine served as lead for conceptualization, data curation, formal analysis, investigation, methodology, project administration, resources, software, visualization, writing–original draft, and writing–review and editing. Mathieu Roy contributed equally to supervision and served in a supporting role for methodology. Ulrik Beierholm served in a supporting role for conceptualization. A. Ross Otto served as lead for funding acquisition and supervision and served in a supporting role for conceptualization and methodology. Mathieu Roy, Ulrik Beierholm, and A. Ross Otto contributed equally to writing–original draft and writing–review and editing.

Correspondence concerning this article should be addressed to Sean Devine, Department of Psychology, McGill University, 2001 Av. McGill College, Room 739, Montréal, QC H3A 1G1, Canada. Email: seandamiandevine@gmail.com

Many goals in daily life require the sustained investment of cognitive effort to accomplish (e.g., working an 8-hr shift, filing one's taxes, or writing a lengthy research article). However, ample evidence suggests that people find it difficult to exert sustained cognitive effort over long periods of time (Inzlicht et al., 2014; Kurzban, 2016; Lin et al., 2020; Massar et al., 2016; Shenhav et al., 2017; Umemoto et al., 2019; Wiehler et al., 2022) and that, all else being equal, people tend to avoid engagement in cognitively demanding behaviors (Hull, 1943; Kool et al., 2010; Shenhav et al., 2017; Vogel et al., 2020; Westbrook et al., 2013).

Yet, while evidence suggests that the sensation of effort accrues over sustained exertion (Ackerman, 2011; Francis et al., 2018; Lorist et al., 2000; Wiehler et al., 2022), recent and classical work highlights the human and animal tendency to increase effort exertion near a goal. This idea was first formalized by Hull's (1932) goal-gradient hypothesis, which posits that organisms increase (physical) effort expenditure as their distance to a goal decreases. Supporting this idea, Brown (1948) observed that the force (i.e., vigor) with which rats ran down a straight alley was proportional to their proximity to a reward, such that the animals exerted more force near a reward versus further away. Extending this idea to human behavior, swimmers and runners "sprint to finish" in the final distances of a race (McGibbon et al., 2018; Tucker et al., 2006), students study more at the end of the semester than the middle (Brahm et al., 2017), and laboratory participants take fewer breaks near the end of a long experiment (Katzir et al., 2020).

However, despite nearly a century having passed since the goal-gradient hypothesis was first postulated—which was originally concerned with physical (i.e., motor) effort—little work has examined potential goal-gradient effects in the cognitive domain. Namely, do individuals intensify their cognitive effort investment as they approach a task goal—for example, the end of a block of trials? While many researchers have observed that a person's exertion of flexible control over behavior fluctuates considerably over time (Braver et al., 2003; Kahneman, 1973), a heretofore overlooked, but potentially important determinant of momentary effort exertion is goal proximity. For example, Emanuel et al. (2022) recently demonstrated that people make more frequent button responses as they near a time limit during a computer game task. While the key outcome variable in this study was motor execution speed, it is noteworthy that the task itself involved a degree of cognitive effort (e.g., aiming and predicting a spaceship's trajectory). These results thus hint that cognitive effort may follow a similar gradient-like pattern to motor effort, but importantly, the task used in this work was not designed to directly index cognitive effort exertion.

Moreover, while past work has demonstrated that goal proximity invigorates responding (Emanuel et al., 2022; Hull, 1932), it is unclear how goal proximity affects cognitive task performance—that is, the speed and fidelity of information processing. More specifically, is information processing fidelity enhanced near a goal—simultaneously resulting in faster and more accurate decisions—or do individuals shift their speed-accuracy trade-off to respond more quickly, but less accurately, near a goal? Here, we provide a direct examination of goal-gradient effects in cognitive task performance as individuals completed a simple cognitive task.

To do this, we employ an attentionally demanding oddball task (Beierholm et al., 2013; see Figure 1), which required participants to maintain active attention to make rapid (subsecond) and accurate judgments about which of three stimuli presented on the screen is the "odd one out." Critically, on half of the trials, participants were presented with visual information about their progress with respect to a goal (Figure 1A)—that is, the number of remaining correct responses needed to complete a block and receive a monetary reward—allowing us to measure subject-specific modulations in performance (response time [RT] and accuracy) as a function of goal proximity and, in turn, to probe the specific computational mechanisms underpinning potential goal-gradient effects using a drift diffusion model (DDM; da Silva Castanheira et al., 2022; Wiecki et al., 2013).

In an initial experiment ($N = 40$) and a replication sample ($N = 42$), we find that participants invested greater cognitive effort—manifesting in performance increases—when a goal state was proximal versus distant, and critically, these performance effects were only observed when participants had information about their progress through a block (i.e., a progress bar [see Figure 1A], vs. no-progress information [see Figure 1B]).

To test whether goal proximity enhances attentional processing manifesting as strength of evidence, or if it shifts response strategy via higher response caution—we used a DDM (Ratcliff & McKoon, 2008; Wiecki et al., 2013). In brief, the DDM assumes that the internal evidence used to make a decision accumulates over time until a response threshold—which controls the speed–accuracy trade-off—is reached. By jointly analyzing response accuracy and RTs, fitting a DDM affords additional understanding of whether the effects of goal proximity on cognitive effort exertion are attributable to increases in the strength of evidence accumulation (i.e., a faster drift rate; which would result in more accurate and faster responses and indicate better information encoding), or by increases in response caution (i.e., elevated response thresholds; which would result in more accurate but slower responses and indicate a shift in response strategy near a goal). Critically then, the DDM allows us to dissociate between the cognitive mechanisms underpinning effortful control near a goal—is speeding driven by heightened drift rates, reflecting an uptick in the rate of information processing, and/or by reduced response caution, reflecting a shift along a speed–accuracy trade-off?

Method

Below, we describe the procedure for two experiments: an initial experiment (Experiment 1) and a nearly identical replication study (Experiment 2). As described below, differences between these two experiments were minimal. Because these experiment designs and ensuing results were nearly identical, our Results section reports results from both experiments, though both data sets were analyzed separately.

Participants

To estimate an appropriate sample size, we conducted a simulation-based power analysis (Arend & Schäfer, 2019) in which participant behavior that was consistent with our hypotheses was simulated on the Oddball task (described) below. Owing in part to the large number of repeated trials individual participants completed in a session ($Mdn = 1{,}372$ trials/participant), this a priori power analysis revealed that 25 participants would yield 80% power to detect effect sizes of minimal interest (see the online supplemental materials for full details of the power analysis). Expecting some participants to be excluded (see below), in Experiment 1, we recruited 40 healthy adult participants (91% female; average age = 22.24, $SD = 3.64$) from McGill University's participant pool. All participants gave informed consent

Figure 1
Schematic of the Oddball Task

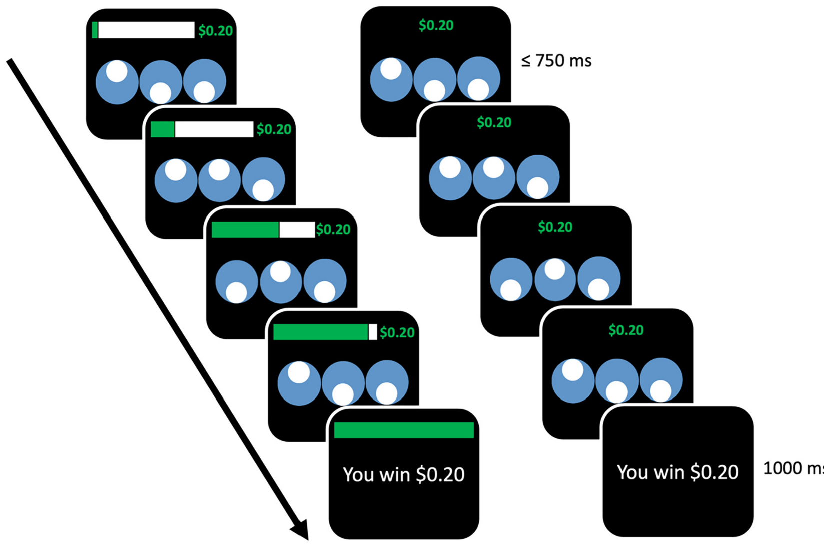

Note. Participants were asked to judge which of three shapes was the "odd-one-out," within 750 ms. The left path depicts an example block in the progress condition, with a green progress bar incrementing after each correct response, culminating in the advertised reward being won, assuming a minimum level of performance (at least 75% correct). The right path shows a no-progress block, where the progress bar is not shown, but the same number of correct trials (60) is needed to obtain the reward and progress in the task. In Experiment 2, a 100-ms ISI was added between each judgment. Rewards remained on the screen for 1,000 ms. ISI = Interstimulus interval. See the online article for the color version of this figure.

prior to testing and were compensated with course credit plus a cash bonus of $5 CAN. This procedure was approved by the McGill Research Ethics Board (REB 137-0816).

To ensure participants achieved a minimum level of performance, we excluded participants who failed to meet the following criteria in the final analysis: More than 25% of responses were missing, participants responded incorrectly, on average, less than 43% of the time (10% more than chance performance), and average RTs were smaller than 100 ms. These exclusion criteria unexpectedly resulted in the exclusion of 17 participants, leaving only 23 participants in the final analysis, falling short of our target sample size. Participant debriefings indicated that the lack of distinguishable boundaries between trials (i.e., interstimulus intervals [ISIs]; see Figure 1) may have contributed to some participants' poor accuracy rates in Experiment 1.

In Experiment 2, we attempted to mitigate this issue, recruiting an additional 42 participants (91% female; average age = 20.00, *SD* = 1.58) in a replication study that was identical to Experiment 1, except for the addition of a 100 ms ISI (see below). This addition improved performance such that, using the same criteria as described above, only five participants were excluded, leaving 37 participants in the final analysis—well above the target sample size.

Oddball Task

Participants completed a fast-paced oddball task, which we adapted from a previously used paradigm (Beierholm et al., 2013; Guitart-Masip et al., 2011). The task was programmed using the PsychoPy library for Python (Peirce, 2007). Each trial, participants were shown three blue circles that were evenly spaced horizontally across the screen. Two of these circles were identical, containing an inner white circle at the top (bottom) of the blue circles, while the third was different, containing a white circle at the bottom (top) of the blue circles (see Figure 1). Participants were asked to identify which shape was the "odd one out," using the Q (leftmost circle), W (middle circle), or E (rightmost circle) keys to indicate their choice. The position of the odd circle was randomized from trial to trial. Participants had 750 ms to indicate which stimulus was the "odd one out." If they did not respond in time, the shapes changed, and the trial was marked as a timeout. In other words, the experiment was self-paced, such that faster responses resulted in faster task completion. In Experiment 2, an ISI of 100 ms was included between each set of stimuli, such that a brief "flash" differentiated one set of three circles (i.e., one trial) from the next (the next trial).

Participants completed two types of oddball task blocks, each of which required 60 correct responses to complete: progress blocks and no-progress blocks (Figure 1). In progress blocks, a green progress bar was presented above the oddball stimuli (378 × 25 pixels). Every time participants correctly identified the odd stimulus, the progress bar would increment by 1/60th of its total size. In no-progress blocks, this progress bar was not shown. For each block, regardless of progress information, participants were informed that if they responded correctly and within the response deadline of 75% of trials during a block, they would receive a reward. The reward was indicated as either high ($0.20) or low ($0.02) and was displayed either beside the progress bar on progress blocks or in the center of the screen above the stimuli in no-progress blocks (Figure 1).

At the end of each block, a screen appeared reading "You won $X.XX" if participants met the performance criterion or "You made too many mistakes this round. No extra money" if they did not, for 1,000 ms.

Participants completed 16 blocks (~1,350 trials per participant) of the oddball task in a 2 (progress, no progress) × 2 (low reward, high reward) design repeated 4 times for each participant in a pseudorandomized order.

Inferential Statistics

To examine performance, we estimated mixed-effects regressions predicting participants (log) RTs on correct trials, and accuracy (using a logistic model) from progress information condition (deviance-coded, $-0.5 =$ no progress, $0.5 =$ progress), goal proximity (distance to the end a block, mean-centered), and reward magnitude (deviance-coded, $-0.5 =$ low reward, $0.5 =$ high reward). Importantly, the goal-gradient hypothesis predicts that performance should uptick sharply near a reward (Emanuel et al., 2022; Hull, 1932), and other work has suggested that this uptick near the end of task may be accompanied by a parallel decrease in performance near the start (Bonezzi et al., 2011). To capture this hypothesis in our model, we also included a quadratic term for proximity (hereafter, proximity^2), which is computed as the square of the goal proximity term. Larger coefficients on this term thus reflect a steeper increase in performance near the end of a task. All main effects and interactions were modeled.

Supporting classical goal-gradient-like behavior, we hypothesized that there would be a statistically significant interaction between progress condition and proximity and/or proximity^2 to a reward, such that participants' response speeds and accuracy would vary when they were aware that they were near a reward. Secondarily, we predicted that the strength of the goal gradient would differ between reward conditions (high vs. low), manifesting in a significant interaction between reward magnitude, reward proximity, and reward condition, such that (high vs. low).

These models were estimated using the *lme4* package for R (Bates et al., 2014). Likelihood-ratio tests were used to assess relative model fit, and reported confidence intervals are based on the likelihood profile (a.k.a., profile confidence intervals). Random slopes were included unless they caused convergence issues. To better convey the observed goal-gradient effects—and because reward effects were inconsistent across analyses and experiments—Figures 2 and 3 depict behavior collapsed across reward levels.

Hierarchical DDM (HDDM)

To better understand how goal gradients might modulate strategies of cognitive effort exertion, we fit a DDM to participants' responses and RTs across both experiments (Ratcliff & McKoon, 2008; Wiecki et al., 2013). The DDM is one of a family of sequential-sampling models which assume that people's decisions are the results of an iterative, noisy evidence accumulation process over time (Ratcliff & McKoon, 2008). The process begins at some starting point z and accumulates evidence over time to one of two decision boundaries at a constant rate v, known as the drift rate. This accumulation process is subject to random perturbations at each time step and continues until one of two boundaries are crossed, which corresponds to either option in the task: here, correct versus incorrect. The separation between these boundaries is defined by the parameter a, such that more evidence is required to reach a decision when a is larger. The direction the evidence accumulation process heads toward (correct/incorrect) depends on the sign of the drift rate, v, where positive values of v indicate evidence heading for one boundary (here correct responding) and negative values indicate evidence heading toward the other (incorrect responding). Once a boundary is crossed, a response is initiated, which takes some nonzero amount of time to encode and execute ($t0$; "nondecision time").

While the traditional DDM assumes that decision boundaries, a, other classes of DDMs allow for boundaries to vary dynamically over the course of a trial. For example, the collapsing bounds DDM assumes that decision boundaries get progressively narrower over the course of a trial, reflecting increased urgency, and reduced caution, as one spends more time on a task (Hawkins et al., 2015; Smith & Ratcliff, 2022). In these variants, the rate at which decision bounds linearly collapse is given by an additional parameter, θ. When θ is large, decision bounds collapse more quickly leading to reduced caution at longer RTs (i.e., a larger θ corresponds to a greater deviation from 90°, which would reflect stable bounds as in the standard DDM). Since the predictions of a collapsing bounds model are consistent with our predictions about goal gradients in cognitive effort exertion, we also consider this model in our analyses.[1]

We performed hierarchical Bayesian estimation of DDM parameters using the HDDM package for Python (Version 0.9.9; Fengler et al., 2022). We fit two broad classes of HDDMs to the present data: a standard DDM, and the collapsing bounds model described above. Posterior distributions for decision thresholds (a), drift rates (v), nondecision times ($t0$), and, in the case of the collapsing bounds model, collapse rate θ, were estimated on a trial-by-trial basis as a linear combination of progress condition and proximity to reward, with random intercepts taken per participant. All other parameters were assumed to be fixed to default values set by HDDM (see Fengler et al., 2022; Wiecki et al., 2013).

While analyses of task performance (described below) suggested a quadratic relationship between goal proximity and RTs, it was unclear which parameters of the HDDM would reflect this quadratic relationship. Moreover, multiple combinations of model structure (standard vs. collapsing bounds) and parameter values may readily capture the general pattern observed in the data. As such, it was important to explore the full parameterization space of these regression-based DDMs. Accordingly, we fit 2,744 separate HDDMs to the data, which covered the full model and parameter space, and compared model fit for each model using the deviance information criterion (DIC; Gelman et al., 2014). The specification of each model is included in the online supplemental materials, and goodness of fit for each model is visualized in Figure S4 in the online supplemental materials.

Five thousand samples were drawn from the posterior for each parameter, discarding the first 2,000 samples for burn-in, and no thinning was applied. Convergence for winning models was assessed via visual inspection of trace plots, Geweke's statistics (Geweke, 2005; reported in Table 3), and posterior predictive checking, presented in Figure S5 in the online supplemental materials. Geweke's statistic test for equality of the means of the first (10%)

[1] The authors thank an anonymous reviewer for suggesting we consider collapsing bounds variants of the DDM to the present data set.

Figure 2

Task Performance Over the Course of a Block in Experiments 1 (A and B) and 2 (C and D)

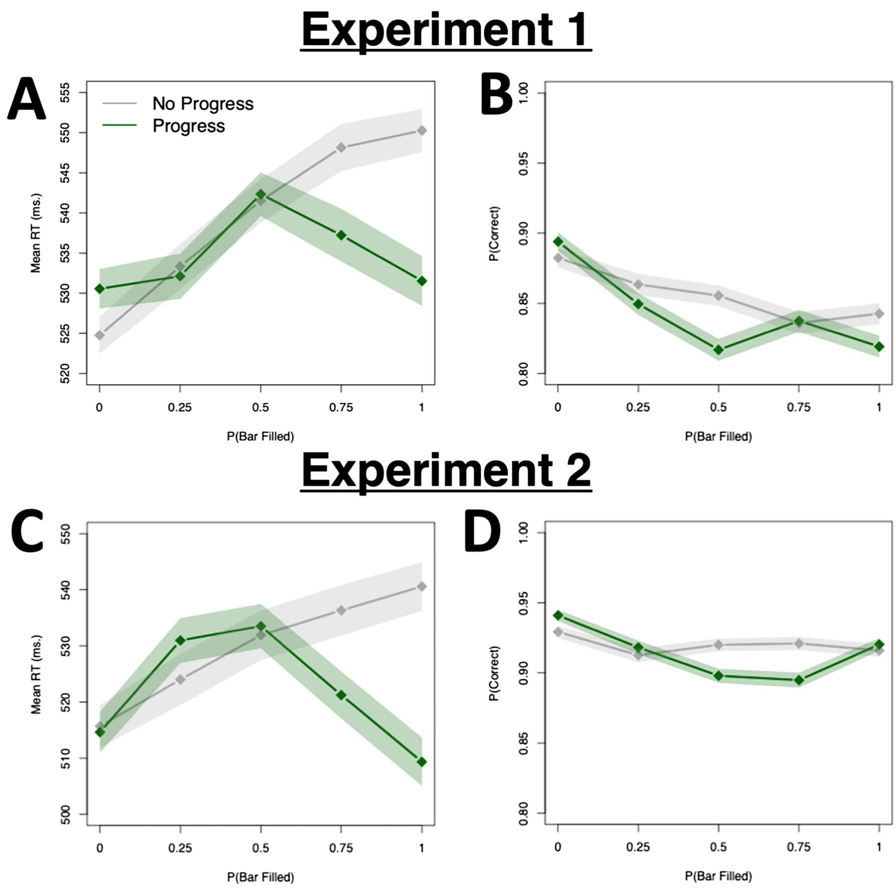

Note. Green lines (dark gray) represent performance during progress blocks and gray lines (light gray) represent performance during no-progress blocks. Each point represents an averaged bin of correct response times (A, C) or proportion of correct trials (C, D), represented on the *y*-axis. The *x*-axis shows the proportion that the progress bar was filled (i.e., proximity to the reward; whether (green/dark gray) or not (gray/light gray) it was shown to participants). Ribbons represent the standard error of the mean. RT = response time; P(Correct) = proportion of correct; P(Bar Filled) = progress bar was filled. See the online article for the color version of this figure.

and last part of a Markov chain (50%). If samples are drawn from a stationary distribution, the two means are statistically equivalent, and the associated absolute *Z*-score (known as Geweke's statistic) is below 1.96, which indicates that a chain has converged successfully. We also conducted a parameter recovery, which demonstrated that parameters from the winning model recovered well (Figure S6 in the online supplemental materials).

All reported coefficient estimates (b values) for the HDDM are mean posterior values and 95% highest posterior density intervals. Bayesian p values (p) represent one minus the proportion of the posterior that falls above or below zero (depending on the sign of the median posterior value: below zero if $b < 0$ and above if $b > 0$).

Transparency and Openness

All task codes, materials, raw data, analysis scripts, results, and computational models are openly available at https://github.com/seandamiandevine/EffortProgress2.

Results

Task Performance

Overall, participants' RTs were fast (Experiment 1: $M = 534.46$), taking on average 71% of the allotted time (750 ms) to respond as well as accurate—P(Correct) = 0.86. Critically, in the progress information condition ("progress" blocks), participants had a visual cue indicating their proximity to a reward whereas in "no-progress" blocks this information was absent. In accordance with goal-gradient hypothesis, we hypothesized that participants would exhibit speeded responses near a goal (i.e., the reward at the end of a block), but only when participants had information about task progress.

Figure 2A and 2C depict participants' RTs as a function of the proportion of correct responses made relative to the total number of correct responses needed to obtain a reward (60 per block) in Experiments 1 and 2, respectively. Consistent with classic goal-gradient effects—characterized by increased vigor near a reward

Figure 3

Predicted Parameter Values for the Best-Fitting DDM

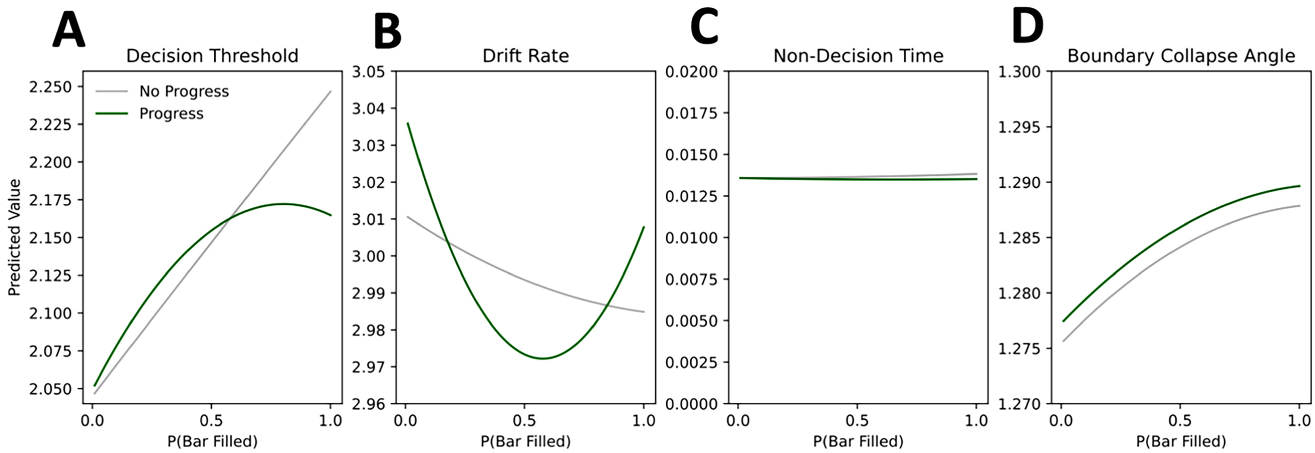

Note. The *y*-axis shows the posterior prediction for the parameter (referenced in each subfigure title) as a function of block progress. The *x*-axis shows the proportion that the progress bar was filled (i.e., proximity to the reward). Green lines (dark gray) represent predictions during progress blocks, where the progress bar was shown. Gray lines (light gray) represent predictions during no-progress blocks, where the progress bar was not shown. DDM = drift-diffusion model; P(Bar Filled) = progress bar was filled. See the online article for the color version of this figure.

(Brown, 1948; Hull, 1932)—participants' RTs hastened as they neared a reward, but only when information about their block progress was presented. In the absence of progress information, participants' RTs progressively slowed over the course of a block, consistent with an accruing mental fatigue or task disengagement (Lorist et al., 2005).

Examining (log) RTs with mixed-effect regressions, the best-fitting model in both experiments included both a linear and quadratic term representing proximity to reward—Experiment 1: $\chi^2(2) = 13.56$, $p = .001$; Experiment 2: $\chi^2(2) = 103.12$, $p < .0001$ (see Table 1 for coefficient estimates). More importantly, we found a significant interaction between progress condition and (linear) goal proximity (Experiment 1: $b = -0.08$, 95% confidence interval [CI] = [−0.12, −0.04], $p < .0001$; Experiment 2: $b = -0.08$, 95% CI = [−0.10, −0.05], $p < .0001$) as well as an interaction between progress condition and (quadratic) goal proximity (Experiment 1: $b = 0.26$, 95% CI = [−0.43, −0.09], $p = .002$; Experiment 2: $b = -0.37$, 95% CI = [−0.47, −0.28], $p < .0001$). These effects are visualized in Figure 3A and 3C.

Figure 3B and 3D depict response accuracy as a function of block progress. Again, we found that the best-fitting logistic regression model for both experiments included a linear and quadratic term for reward proximity—Experiment 1: $\chi^2(2) = 47.26$, $p < .0001$; Experiment 2: $\chi^2(2) = 28.10$, $p < .0001$. Overall, in both experiments, we observed that accuracy declined linearly over the course of a block (Experiment 1: $b = -0.83$, 95% CI = [−0.96, −0.69], $p < .0001$; Experiment 2: $b = -0.37$, 95% CI = [−0.50, −0.24], $p < .0001$). In Experiment 2, we found a slight improvement of accuracy as participants neared a reward, as reflected in an interaction between progress condition and (quadratic) reward proximity ($b = 1.02$, 95% CI = [0.00, 2.04], $p = .049$; see Table 2 for full coefficient estimates). This effect was not statistically significant in Experiment 1 ($b = 0.50$, 95% CI = [−0.50, 1.53], $p = .339$), though its positive directionality suggests against the possibility that accuracy worsened near a reward (Figure 2B and 2D).

As we manipulated reward amount in a blockwise fashion, holding goal proximity constant, reward incentives did not reliably shift task performance across experiments (Tables 1 and 2): Accuracy was higher in Experiment 1, and RTs were faster in Experiment 2, on high (vs. low) reward blocks, but these effects were not consistent across experiments. This is consistent with past work, showing that performance on the oddball task is particularly sensitive to trial-by-trial fluctuations in reward magnitudes (Beierholm et al., 2013; Guitart-Masip et al., 2011), but perhaps less sensitive to the blockwise manipulation used here, which was a necessary by-product of the goal proximity manipulation.

Finally, we examined whether the apparent goal-gradient effects differed across reward magnitudes, finding small and inconsistent effects of available reward for each block upon RTs across experiments. Specifically, we observed a three-way interaction between progress condition, proximity^2, and reward condition in both experiments. However, these observed effects were in the opposite direction across experiments: In Experiment 1, low-reward blocks engendered enhanced RT speeding when participants had block progress information ($b = 0.47$, 95% CI = 0.14–0.81, $p = .005$; Figure S2 in the online supplemental materials; see Tables 1 and 2), whereas in Experiment 2, we found enhanced speeding in high-reward blocks when progress information was present ($b = -0.22$, 95% CI = [−0.42, −0.02], $p = .03$; Figure S3 in the online supplemental materials). Similar inconsistent and weak reward effects on goal gradients have recently been reported in the literature (Emanuel, 2022; Emanuel et al., 2022).

DDM Fits

To examine whether the observed effects of goal proximity upon task performance are explained by an increased quality of sensory information from the stimuli (which would manifest as a goal proximity-induced changes in drift rate), a change in response caution (manifesting as proximity-induced changes in response thresholds), or a change in response execution time (manifesting as changes in

Table 1
Best-Fitting Mixed-Effects Linear Regressions Predicting (Log) Response Times in Experiments 1 and 2

Coefficient	Experiment 1			Experiment 2		
	b	95% CI	*p*	*b*	95% CI	*p*
Intercept	6.23	[6.20, 6.26]	<.001	6.23	[6.21, 6.25]	<.001
Prog. cond.	−0.00	[−0.02, 0.02]	.979	0.02	[0.01, 0.03]	<.001
Reward	−0.00	[−0.02, 0.02]	.761	−0.02	[−0.03, −0.00]	.007
Proximity	0.01	[−0.01, 0.04]	.21	0.01	[0.00, 0.03]	.047
Proximity^2	−0.09	[−0.17, −0.00]	.04	−0.18	[−0.23, −0.13]	<.001
Prog. Cond. × Reward	−0.02	[−0.06, 0.01]	.235	0.04	[0.02, 0.06]	.001
Prog. Cond. × Proximity	−0.08	[−0.12, −0.04]	<.001	−0.08	[−0.10, −0.05]	<.001
Prog. Cond. × Proximity^2	−0.26	[−0.43, −0.09]	.002	−0.37	[−0.47, −0.27]	<.001
Reward × Proximity	−0.03	[−0.07, 0.02]	.25	0.01	[−0.02, 0.03]	.666
Reward × Proximity^2	0.1	[−0.06, 0.27]	.232	0.09	[−0.01, 0.18]	.088
Prog. Cond. × Reward × Proximity	−0.07	[−0.15, 0.02]	.134	0.03	[−0.02, 0.08]	.292
Prog. Cond. × Reward × Proximity^2	0.47	[0.14, 0.81]	.005	−0.22	[−0.42, −0.02]	.03

Note. CI = confidence interval; prog. cond. = progress condition.

nondecision time), we jointly modeled participants' RTs and accuracies using two variants of a HDDM (Fengler et al., 2022; Wiecki et al., 2013): a standard three-parameter DDM (thresholds, drift rate, and nondecision times) and a collapsing bounds DDM (includes an additional linear threshold collapse parameter). Owing to the similarity in behavioral results between experiments, we fit a single model to the data from both experiments. Furthermore, given the observed quadratic relationship between reward proximity and RTs in the progress condition (Figure 2A and 2C), we intuited that some combination of trial-by-trial DDM parameters (drift rate, threshold, nondecision time, and, possibly, collapse rate) would depend on a quadratic term representing goal proximity (proximity^2).

As we did not have strong intuitions about how potential goal proximity effects might manifest in DDM parameters, we compared the goodness of fit of 2,477 different HDDM specifications—which exhaustively covered the space of possible combinations of linear and quadratic relationships between goal proximity and the DDM parameters of interest—to participants' behavior. Overall, we found that a collapsing bounds variant of the DDM fit the present data markedly better than a standard (fixed bounds) DDM (difference in median DICs = 22,719.63; Figure S4 in the online supplemental materials). The best-fitting model assumed that decision thresholds and drift rates depended (quadratically) on progress condition and goal proximity, where nondecision times and boundary collapse rate changed (linearly) over the course of a block. This model captured behavior well in both experiments, as evidenced by the posterior predictive checks visualized in Figure S5 in the online supplemental materials. Mean posterior values, credibility intervals, Bayesian *p*-values, and Geweke's statistics are presented in Table 3 for both experiments.

Figure 3 depicts predicted trial-to-trial parameter values of the winning collapsing bounds model over the course of a block, as a function of simulated values of progress condition (progress or no progress) and reward proximity. With respect to decision thresholds, we observed a robust interaction between progress condition and goal proximity ($b = -0.09$, 95% CI = [−0.11, −0.06], $p = 0$) and proximity^2 ($b = -0.19$, 95% CI = [−0.22, −0.16], $p = 0$), such that participants exhibited reduced decision thresholds as they approached a reward, when progress information was present compared to when it was not. In other words, when participants were aware of their proximity to a reward, participants responded less cautiously near a reward than when they were unaware. Moreover, boundaries collapsed more quickly (larger boundary collapse angle) overall during progress blocks as compared to no-progress blocks ($b = 0.002$, 95% CI = [0.0003, 0.003],

Table 2
Best-Fitting Mixed-Effects Logistic Regressions Predicting Correct Responses in Experiments 1 and 2

Coefficient	Experiment 1			Experiment 2		
	b	95% CI	*p*	*b*	95% CI	*p*
Intercept	1.92	[1.64, 2.20]	<.001	2.47	[2.31, 2.63]	<.001
Prog. cond.	−0.09	[−0.20, 0.01]	.089	−0.15	[−0.26, −0.04]	.006
Reward	0.17	[0.07, 0.28]	.002	−0.01	[−0.12, 0.10]	.882
Proximity	−0.83	[−0.97, −0.70]	<.001	−0.37	[−0.50, −0.24]	<.001
Proximity^2	1.75	[1.23, 2.26]	<.001	1.27	[0.76, 1.78]	<.001
Prog. Cond. × Reward	0.04	[−0.17, 0.25]	.714	0.15	[−0.07, 0.36]	.181
Prog. Cond. × Proximity	−0.16	[−0.43, 0.11]	.239	−0.28	[−0.54, −0.01]	.042
Prog. Cond. × Proximity^2	0.5	[−0.53, 1.53]	.339	1.02	[0.01, 2.04]	.049
Reward × Proximity	−0.11	[−0.38, 0.16]	.437	0.1	[−0.16, 0.37]	.45
Reward × Proximity^2	−1.21	[−2.24, −0.18]	.022	−0.07	[−1.10, 0.95]	.886
Prog. Cond. × Reward × Proximity	−0.16	[−0.70, 0.38]	.563	0.06	[−0.47, 0.60]	.819
Prog. Cond. × Reward × Proximity^2	0.41	[−1.66, 2.48]	.700	−0.09	[−2.09, 1.91]	.931

Note. CI = confidence interval; prog. cond. = progress condition.

Table 3
Parameter Estimates for the Best-Fitting DDM

Coefficient	*M*	95% CI	*p*	Geweke's statistic
Decision thresholds				
Intercept	2.151	[1.841, 2.523]	.000	−1.721
Proximity	0.159	[0.137, 0.179]	.000	−12.334
Proximity^2	−0.097	[−0.122, −0.074]	.000	9.411
Prog. cond.	0.008	[−0.002, 0.018]	.061	1.112
Prog. Cond. × Proximity	−0.086	[−0.11, −0.064]	.000	1.169
Prog. Cond. × Proximity^2	−0.187	[−0.215, −0.16]	.000	−2.707
Drift rate				
Intercept	2.983	[2.942,3.000]	.000	−0.253
Proximity	−0.028	[−0.048, −0.009]	.003	1.570
Proximity^2	0.108	[0.08, 0.142]	.000	−3.452
Prog. cond.	−0.020	[−0.04, −0.001]	.021	0.176
Prog. Cond. × Proximity	−0.004	[−0.014, 0.005]	.198	−0.881
Prog. Cond. × Proximity^2	0.181	[0.134, 0.248]	.000	−1.704
Nondecision times				
Intercept	0.014	[0.001, 0.043]	.000	−0.318
Proximity	0.000	[−0.001, 0.001]	.424	−0.676
Proximity^2	0.000	[−0.001, 0.001]	.282	1.610
Prog. Cond.	0.000	[−0.001, 0]	.286	1.492
Prog. Cond. × Proximity	0.000	[−0.002, 0.001]	.304	1.476
Boundary collapse rate				
Intercept	1.285	[1.25, 1.3]	.000	−1.330
Proximity	0.012	[0.008, 0.017]	.000	−9.691
Proximity^2	−0.010	[−0.017, −0.005]	.000	8.178
Prog. cond.	0.002	[0, 0.003]	.012	−0.579

Note. DDM = drift-diffusion model; CI = confidence interval; prog. cond. = progress condition.

$p = .01$), further suggesting reduced caution when progress information was available.

With respect to trial-to-trial drift rates, we observed an interaction between progress condition and proximity^2 ($b = 0.18$, 95% CI = [0.13, 0.25], $p = 0$), suggesting that drift rates initially decreased over the course of a block, but increased sharply near the end—and importantly, this uptick was more not present when progress information was not present (Figure 3B).

Finally, we did not observe any modulation of nondecision times by progress condition (see Table 3).

Taken together, the results of the HDDM analysis suggest that goal-proximity induced task performance shifts were best explained by two mechanisms: while the fidelity with which participants processed information (drift rates) increased near a reward, response caution diminished suggesting a shift from a slower, but accuracy-oriented, strategy in the middle of a task block to a faster, but incautious, response style during and near the end of a task block when progress information was known.

Discussion

Exerting sustained cognitive effort is taxing and aversive, and these effort costs are exacerbated with continued exertion (Kool et al., 2010; Shenhav et al., 2017). However, it has also long been observed that the investment of cognitive resources into a demanding task fluctuates considerably over time (Braver et al., 2003; Kahneman, 1973; Otto & Daw, 2019). In line with this view, both recent (Emanuel et al., 2022; Katzir et al., 2020) and classical (Hull, 1932) work on goal gradients posit that effort exertion should uptick near the end of a task—when the rewards associated with the successful completion of that task are proximal. Yet, empirical evidence for this claim in cognitive effort is scant and the computational mechanisms underpinning proximity-induced effort modulations of this type have not been investigated. Here—first in a novel experiment and then again in a replication sample—we examined goal-gradient-like behavior in a demanding attentional oddball task.

We found that participants engaged in speeded, but nevertheless accurate, responding as a function of goal state proximity and that this behavior was (a) present only when participants were aware of their proximity to a goal, and (b) consistent across levels of rewards on offer. Thus, participants were not only faster, but also equally—or even slightly more—accurate near the end of a block. Notably, the relationship between response speed and goal proximity was quadratic in nature, owing to slower performance in the middle of a block. This finding is consistent with past work demonstrating a negative relationship between time on task and sustained attention and motivation (Fortenbaugh et al., 2017) and highlights the relative nature of goal gradients as a means of enhancing effort exertion from motivational low points (e.g., during late-block trials in the no-progress condition in these experiments; Bonezzi et al., 2011), as well as rules out simple learning effects as an explanation for goal gradients. In this regard then, these results hint that goal gradients in the current task acted to restore motivation near a goal state, rather than to enhance it beyond a preestablished control point. Together, these shifts in performance reflect key indicators of increased cognitive effort exertion. Moreover, they extend past work on goal gradients which simultaneously measured effort exertion in both cognitive and physical domains (Rauch et al., 2013), and suggest that motor response vigor closer to the end of the task might be explained by a combination of adaptations of response strategy (above and beyond solely response speed).

In this respect, these results suggest a shift in cognitive strategies that cannot be explained by a simple speed–accuracy trade-off

(Heitz, 2014). Supporting this view, we demonstrate that these goal-gradient effects are well-characterized by a collapsing decision bounds variant of the DDM (Figures S4 and S5 in the online supplemental materials). Specifically, we find that decision thresholds decreased, but concurrently, drift rates rapidly upticked in proximity to a reward.

How might this constellation of DDM parameter changes be interpreted, psychologically? Previous research suggests that narrowing decision thresholds may be associated with reduced response caution, reflecting a failure, or disinclination to incorporate additional evidence into one's decision (Lin et al., 2020; Voss et al., 2004). At the same time, heightened drift rates have been interpreted as an increase in fidelity of stimulus encoding which manifests in an increased rate of information uptake per unit time (Voss et al., 2004), motivation to perform a task correctly (Bottemanne & Dreher, 2019), and increased cognitive control deployment (Cavanagh et al., 2014; Otto & Daw, 2019). More parsimoniously then, drift rates here should be understood to reflect a modulation in proximal cognitive control allocation—for example, momentary attention to the current task—which in turn can be up- and downregulated by motivation factors (Cavanagh et al., 2014; Leng et al., 2021). In line with the goal-gradient hypothesis, increased drift rates have also been theorized to reflect increased effort costs (Drugowitsch et al., 2012). Taken together, the pattern of DDM parameter changes observed here—and their contingent change over the course of a block based on the presence versus absence of progress information—suggests that proximity to rewards may reflect a global "urgency" signal, signaling a shift in behavioral strategy, from a slow and cautious approach to the task to a faster, completion-focused, but nevertheless cognitive effortful, one.

Connecting this account to our computational results more directly, recent work has highlighted how joint changes to drift rates and thresholds might reflect a simultaneous reconfiguration of control strategies in favor of maximizing reward rate. For example, Leng et al. (2021) observed that participants dynamically allocated cognitive control to a task to optimize the trial-by-trial reward rate in the task—captured by simultaneous drift rate increases and threshold decreases—suggesting an adaptive shift in response to shifts in task incentives (Ritz et al, 2022). Similarly, Otto and Daw (2019) found that moment-to-moment shifts in the environmental average reward rate engendered simultaneous decreases in drift rates and thresholds, interpreted as a withdrawal of cognitive effort with a concomitant reduction in response caution. While the present task was not specifically calibrated to detect nuanced changes in control parameters in response to changes in incentive (as rewards were fixed within a block of trials), participants here may have simply used goal proximity as a cue to reconfigure control strategies to favor reward rate maximization, without external change in the incentive structure of the task.

Finally, and in line with a growing body of evidence suggesting that such urgency signals guide neuronal activity and behavior within the course of a single trial (Cisek et al., 2009), one interpretation of the current results is that requirements on evidence accumulation over the course of a longer time period are modulated when rewards are proximal versus distal. Consistent with classical work on goal gradients in rats (Brown, 1948; Hull, 1932), similar urgency signals have been observed to guide low-level behavior in animals as well as humans at the trial level (Hanks & Summerfield, 2017; Hernández-Navarro et al., 2021; Thura & Cisek, 2017). Analogously then, goal gradients in this context seem to reflect a larger manifestation of urgency signals, governing behavior not only at the trial level, but also within the local context (here a block, or elsewhere an entire experiment; Katzir et al., 2020, a foot race; Tucker et al., 2006, a semester; Brahm et al., 2017, in consumer decisions; Zhu et al., 2018).

It is important to remark on limitations of the present results and future directions that they may suggest. First, the effects of reward on offer on goal-gradient behavior in the present data were inconsistent between experiments and, contrary to our initial hypothesis, goal-gradient behavior was largely invariant to changes in reward magnitude. On the one hand, it is possible that reward magnitudes were too low in the present set of experiments to evince larger differences in task performance, as evidenced by the relative lack of incentive effects, holding goal proximity constant. This would be consistent with past work, in which reward incentive effects can be highly variable and context-dependent (e.g., Otto & Daw, 2019). On the other hand, this finding is also consistent with recent work (Emanuel, 2022), showing that the effects of monetary reward on proximity-induced (motor) effort exertion were heterogeneous and, elsewhere (Emanuel et al., 2022), that changes in motor vigor may be less sensitive to goal proximity when a goal is based on performance-based metrics success.

Second, the present set of experiments only used rewards as a motivating stimulus. This is notable, as a key tenet of classical work is that goal gradients should be steeper for avoidance behaviors (Heilizer, 1977). Future work should therefore endeavor to test whether this hypothesis holds in the domain of cognitive effort. This could be accomplished by extending the present work to the avoidance domain, using, for instance, painful stimulation instead of (or in addition to) monetary reward.

Third, the present study was principally concerned with how goal gradients manifest in cognitive effort exertion, and the computational mechanisms that underpin these effects. Another important question is why these effects occur. In this respect, Emanuel et al. (2022) recently argued that goal-gradient effects in motivation can be explained by diminishing opportunity costs (of the to-be-completed task) near a goal (see also Beierholm et al., 2013; Dora et al., 2022; Kurzban et al., 2013). On this view, as an individual nears a goal, their motivation to complete the task increases because the value of engaging in an alternative activity (e.g., mind-wandering, adjusting one's chair) decreases, since these activities can be postponed until after a task is completed at a lower cost. We believe the present results support this interpretation, and also, importantly, shed light on the question of whether individuals are more motivated to "get it done" (Bonezzi et al., 2011) versus "to do it well" (Touré-Tillery & Fishbach, 2012) near the end of a task. Here, our results support both views, such that as opportunity costs decrease near the end of a task, participants invest increased attentional control into the task (as reflected by an uptick in drift rates), but also aim to finish the task quickly, by reducing evidentiary demands.

Taken together, our results extend past work on goal gradient on physical (motor) effort to the cognitive domain using a simple attentional task. Moreover, they highlight an important, but previously unconsidered, feature of goal gradients: While effort exertion increases near a goal (reflected here by heightened drift rates), caution tends to decrease (reflected by reduced, and collapsing, decision thresholds). This finding is informative not only for theories of goal gradients—relating them, in particular, to the concept of urgency—but also bears practical importance. Interventions utilizing progress indicators to boost performance have become increasingly commonplace in educational and workplace settings (Amabile & Kramer, 2011). These interventions rest on the view that information about one's progress will incrementally increase performance. Insofar as this higher order performance depends on constituent components of cognitive control—here, attention and inhibition—our results call for nuance, as gains in effort exertion should be weighed against potential losses in caution and deliberation.

Constraints on Generality

The current sample consisted principally of English-speaking Western students, recruited at a large Canadian research university. While we have no theoretical reason to suspect that the current results do not generalize to other cultural and demographic groups, this is an assumption which should be tested in future replication work.

References

Ackerman, P. L. (2011). *Cognitive fatigue: Multidisciplinary perspectives on current research and future applications*. American Psychological Association. https://doi.org/10.1037/12343-000

Amabile, T., & Kramer, S. (2011). *The progress principle: Using small wins to ignite joy, engagement, and creativity at work*. Harvard Business Press.

Arend, M. G., & Schäfer, T. (2019). Statistical power in two-level models: A tutorial based on Monte Carlo simulation. *Psychological Methods*, *24*(1), 1–19. https://doi.org/10.1037/met0000195

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. *arXiv preprint*, arXiv:1406.5823.

Beierholm, U., Guitart-Masip, M., Economides, M., Chowdhury, R., Düzel, E., Dolan, R., & Dayan, P. (2013). Dopamine modulates reward-related vigor. *Neuropsychopharmacology*, *38*(8), 1495–1503. https://doi.org/10.1038/npp.2013.48

Bonezzi, A., Brendl, C. M., & De Angelis, M. (2011). Stuck in the middle: The psychophysics of goal pursuit. *Psychological Science*, *22*(5), 607–612. https://doi.org/10.1177/0956797611404899

Bottemanne, L., & Dreher, J. C. (2019). Vicarious rewards modulate the drift rate of evidence accumulation from the drift diffusion model. *Frontiers in Behavioral Neuroscience*, *13*, 142. https://doi.org/10.3389/fnbeh.2019.00142

Brahm, T., Jenert, T., & Wagner, D. (2017). The crucial first year: A longitudinal study of students' motivational development at a Swiss Business School. *Higher Education*, *73*(3), 459–478. https://doi.org/10.1007/s10734-016-0095-8

Braver, T. S., Reynolds, J. R., & Donaldson, D. I. (2003). Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron*, *39*(4), 713–726. https://doi.org/10.1016/S0896-6273(03)00466-5

Brown, J. S. (1948). Gradients of approach and avoidance responses and their relation to level of motivation. *Journal of Comparative and Physiological Psychology*, *41*(6), 450–465. https://doi.org/10.1037/h0055463

Cavanagh, J. F., Wiecki, T. V., Kochar, A., & Frank, M. J. (2014). Eye tracking and pupillometry are indicators of dissociable latent decision processes. *Journal of Experimental Psychology: General*, *143*(4), 1476–1488. https://doi.org/10.1037/a0035813

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience*, *29*(37), 11560–11571. https://doi.org/10.1523/JNEUROSCI.1844-09.2009

da Silva Castanheira, K., Lalla, A., Ocampo, K., Otto, A. R., & Sheldon, S. (2022). Reward at encoding but not retrieval modulates memory for detailed events. *Cognition*, *219*, Article 104957. https://doi.org/10.1016/j.cognition.2021.104957

Dora, J., van Hooff, M. L. M., Geurts, S. A. E., Kompier, M. A. J., & Bijleveld, E. (2022). The effect of opportunity costs on mental fatigue in labor/leisure trade-offs. *Journal of Experimental Psychology: General*, *151*(3), 695–710. https://doi.org/10.1037/xge0001095

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience*, *32*(11), 3612–3628. https://doi.org/10.1523/JNEUROSCI.4010-11.2012

Emanuel, A. (2022). *How to stay motivated: The influence of agency and progress related incentives on sustained effort exertion*. PsyArXiv. https://doi.org/10.31234/osf.io/3bf8p

Emanuel, A., Katzir, M., & Liberman, N. (2022). Why do people increase effort near a deadline? An opportunity-cost model of goal gradients. *Journal of Experimental Psychology: General*, *151*(11), 2910–2926. https://doi.org/10.1037/xge0001218

Fengler, A., Bera, K., Pedersen, M. L., & Frank, M. J. (2022). Beyond drift diffusion models: Fitting a broad class of decision and reinforcement learning models with HDDM. *Journal of Cognitive Neuroscience*, *34*(10), 1780–1805. https://doi.org/10.1162/jocn_a_01902

Fortenbaugh, F. C., DeGutis, J., & Esterman, M. (2017). Recent theoretical, neural, and clinical advances in sustained attention research. *Annals of the New York Academy of Sciences*, *1396*(1), 70–91. https://doi.org/10.1111/nyas.13318

Francis, Z., Milyavskaya, M., Lin, H., & Inzlicht, M. (2018). Development of a within-subject, repeated-measures ego-depletion paradigm: Inconsistent results and future recommendations. *Social Psychology*, *49*(5), 271–286. https://doi.org/10.1027/1864-9335/a000348

Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, *24*(6), 997–1016. https://doi.org/10.1007/s11222-013-9416-2

Geweke, J. (2005). *Contemporary Bayesian econometrics and statistics*. John Wiley & Sons.

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*(12), 3933–3938. https://doi.org/10.1162/jocn_a_00090

Hanks, T. D., & Summerfield, C. (2017). Perceptual decision making in rodents, monkeys, and humans. *Neuron*, *93*(1), 15–31. https://doi.org/10.1016/j.neuron.2016.12.003

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *The Journal of Neuroscience*, *35*(6), 2476–2484. https://doi.org/10.1523/JNEUROSCI.2410-14.2015

Heilizer, F. (1977). A review of theory and research on the assumptions of Miller's response competition (conflict) models: Response gradients. *The Journal of General Psychology*, *97*(1), 17–71. https://doi.org/10.1080/00221309.1977.9918503

Heitz, R. P. (2014). The speed-accuracy tradeoff: History, physiology, methodology, and behavior. *Frontiers in Neuroscience*, *8*, Article 150. https://doi.org/10.3389/fnins.2014.00150

Hernández-Navarro, L., Hermoso-Mendizabal, A., Duque, D., de la Rocha, J., & Hyafil, A. (2021). Proactive and reactive accumulation-to-bound processes compete during perceptual decisions. *Nature Communications*, *12*(1), Article 1. https://doi.org/10.1038/s41467-021-27302-8

Hull, C. L. (1932). The goal-gradient hypothesis and maze learning. *Psychological Review*, *39*(1), 25–43. https://doi.org/10.1037/h0072640

Hull, C. L. (1943). *Principles of behavior: An introduction to behavior theory* (pp. x, 422). Appleton-Century.

Inzlicht, M., Legault, L., & Teper, R. (2014). Exploring the mechanisms of self-control improvement. *Current Directions in Psychological Science*, *23*(4), 302–307. https://doi.org/10.1177/0963721414534256

Kahneman, D. (1973). *Attention and effort*. Prentice-Hall.

Katzir, M., Emanuel, A., & Liberman, N. (2020). Cognitive performance is enhanced if one knows when the task will end. *Cognition*, *197*, Article 104189. https://doi.org/10.1016/j.cognition.2020.104189

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

Kurzban, R. (2016). The sense of effort. *Current Opinion in Psychology*, *7*, 67–70. https://doi.org/10.1016/j.copsyc.2015.08.003

Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, *36*(6), 661–679. https://doi.org/10.1017/S0140525X12003196

Leng, X., Yee, D., Ritz, H., & Shenhav, A. (2021). Dissociable influences of reward and punishment on adaptive cognitive control. *PLOS Computational Biology*, *17*(12), Article e1009737. https://doi.org/10.1371/journal.pcbi.1009737

Lin, H., Saunders, B., Friese, M., Evans, N. J., & Inzlicht, M. (2020). Strong effort manipulations reduce response caution: A preregistered reinvention of the Ego-Depletion Paradigm. *Psychological Science*, *31*(5), 531–547. https://doi.org/10.1177/0956797620904990

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

Lorist, M. M., Klein, M., Nieuwenhuis, S., de Jong, R., Mulder, G., & Meijman, T. F. (2000). Mental fatigue and task control: Planning and preparation. *Psychophysiology*, *37*(5), 614–625. https://doi.org/10.1111/1469-8986.3750614

Massar, S. A. A., Lim, J., Sasmita, K., & Chee, M. W. L. (2016). Rewards boost sustained attention through higher effort: A value-based decision making approach. *Biological Psychology*, *120*, 21–27. https://doi.org/10.1016/j.biopsycho.2016.07.019

McGibbon, K. E., Pyne, D. B., Shephard, M. E., & Thompson, K. G. (2018). Pacing in swimming: A systematic review. *Sports Medicine*, *48*(7), 1621–1633. https://doi.org/10.1007/s40279-018-0901-9

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

Peirce, J. W. (2007). Psychopy—psychophysics software in Python. *Journal of Neuroscience Methods*, *162*(1–2), 8–13. https://doi.org/10.1016/j.jneumeth.2006.11.017

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Rauch, H. G. L., Schönbächler, G., & Noakes, T. D. (2013). Neural correlates of motor vigour and motor urgency during exercise. *Sports Medicine*, *43*(4), 227–241. https://doi.org/10.1007/s40279-013-0025-1

Ritz, H., Leng, X., & Shenhav, A. (2022). Cognitive control as a multivariate optimization problem. *Journal of Cognitive Neuroscience*, *34*(4), 569–591. https://doi.org/10.1162/jocn_a_01822

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

Smith, P. L., & Ratcliff, R. (2022). Modeling evidence accumulation decision processes using integral equations: Urgency-gating and collapsing boundaries. *Psychological Review*, *129*(2), 235–267. https://doi.org/10.1037/rev0000301

Thura, D., & Cisek, P. (2017). The basal Ganglia do not select reach targets but control the urgency of commitment. *Neuron*, *95*(5), 1160–1170.e5. https://doi.org/10.1016/j.neuron.2017.07.039

Touré-Tillery, M., & Fishbach, A. (2012). The end justifies the means, but only in the middle. *Journal of Experimental Psychology: General*, *141*(3), 570–583. https://doi.org/10.1037/a0025928

Tucker, R., Lambert, M. I., & Noakes, T. D. (2006). An analysis of pacing strategies during men's world-record performances in track athletics. *International Journal of Sports Physiology and Performance*, *1*(3), 233–245. https://doi.org/10.1123/ijspp.1.3.233

Umemoto, A., Inzlicht, M., & Holroyd, C. B. (2019). Electrophysiological indices of anterior cingulate cortex function reveal changing levels of cognitive effort and reward valuation that sustain task performance. *Neuropsychologia*, *123*, 67–76. https://doi.org/10.1016/j.neuropsychologia.2018.06.010

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

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, *32*, 1206–1220. https://doi.org/10.3758/bf03196893

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

Wiecki, T., Sofer, I., & Frank, M. (2013). HDDM: Hierarchical Bayesian estimation of the drift-diffusion model in Python. *Frontiers in Neuroinformatics*, *7*, Article 14. https://doi.org/10.3389/fninf.2013.00014

Wiehler, A., Branzoli, F., Adanyeguh, I., Mochel, F., & Pessiglione, M. (2022). A neuro-metabolic account of why daylong cognitive work alters the control of economic decisions. *Current Biology*, *32*(16), 3564–3575.e5. https://doi.org/10.1016/j.cub.2022.07.010

Zhu, M., Yang, Y., & Hsee, C. K. (2018). The mere urgency effect. *Journal of Consumer Research*, *45*(3), 673–690. https://doi.org/10.1093/jcr/ucy008

Received March 30, 2023
Revision received January 17, 2024
Accepted January 23, 2024 ■