Instilling the value of effort

Nature Human Behaviour Stage 2 Registered Report (in-principle Stage 1 acceptance)

Hause Lin^{1,2*}, Andrew Westbrook^{3,4,5}, Frank Fan⁶, Michael Inzlicht^{6,7}

¹ University of Regina, Regina, SK, Canada

² Massachusetts Institute of Technology, Cambridge, MA, USA

³ Brown University, Providence, RI, USA

⁴ Radboud University, Nijmegen, Netherlands

⁵ Radboud University Medical Centre, Nijmegen, Netherlands

⁶ University of Toronto, Toronto, ON, Canada

⁷ Rotman School of Management, Toronto, ON, Canada

* Corresponding author: Hause Lin (hauselin@gmail.com)

Abstract

People who take on challenges and persevere longer are more likely to succeed in life. But individuals often avoid exerting effort, and there is limited experimental research investigating whether we can learn to value effort. Because existing research focuses on enhancing cognitive performance rather than increasing the value of effort, it also remains unclear whether individuals can learn to care more about challenging themselves than performing well. We developed a paradigm to test an intuitive idea: that people can learn to value effort and will seek effortful challenges if directly incentivized to do so. What's more, we dissociate the effects of rewarding people for choosing effortful challenges and performing well. Results revealed that rewarding effort increased people's willingness to choose harder tasks, even when rewards were no longer offered (near-transfer). Critically, the effects of this brief manipulation also carried over to an unrelated and unrewarded task (far-transfer). Our results suggest people can learn to value effort and that this valuation can generalise to unfamiliar and unrewarded tasks.

Introduction

Many people value effort. By exerting effort and persevering over time, we learn to approach challenges, become resilient, and develop self-regulatory skills. These qualities correlate not only with academic attainment ^{1–3}, but also outcomes as disparate as health, wealth, and criminal offending ^{4–6}. But despite growing awareness of the importance of these qualities for individuals and societies ^{7–11}, it remains unclear whether and how individuals can learn to value effort and approach challenges.

Although effort is costly ^{12–17}, much evidence suggests that humans and other animals can (learn to) value effort ^{18–21}: We sometimes willingly challenge ourselves with tasks that are difficult to perform (e.g., endurance sports) precisely because these tasks require effort and perseverance, contrary to the notion that we avoid tasks with high perceived effort or low success likelihood ^{14,22,23}. This so-called effort paradox ¹⁹ raises questions about how people (and other animals) come to value effort for its own sake.

Our goal is to develop a paradigm to investigate whether people can learn to seek out effortful tasks in an experimental context. Our approach contrasts with and addresses problems with existing interventions that aim to enhance cognitive skills and performance ^{24–26}. These interventions typically improve performance on trained tasks (i.e., near-transfer effects) but not untrained tasks even when the tasks are closely related ^{27–32}. Crucially, it is unclear whether these interventions benefit everyday behaviour and cognition ³³ (i.e., far-transfer), nor is it clear whether emphasising performance helps or hinders the learning of the value of effort and willingness to embrace effortful challenges.

Interventions that focus directly on improving these qualities have also provided mixed results. For example, training self-control regularly (by repeatedly practising overriding specific habitual responses for extended periods) does not reliably improve self-control performance in the laboratory or willingness to exert effort in everyday life ^{33–35}. Similarly, some work suggests that instilling grit or growth mindsets in students improves academic outcomes ^{36–39}; however, these effects might be less robust than previously thought ^{40–43}. Moreover, these interventions often target many skills concurrently (e.g., grit, goal setting, emotion regulation), making it difficult to evaluate how and why they have (not) worked, and thus, provide little insights into whether people can be trained to value effortful challenges.

Decades of research on learning suggests that rewards ^{44–46}—if they appropriately shape environmental contingencies—can change humans and non-human animals' preferences for challenge and exerting effort. For example, when infants observe adults persevere longer to achieve their goals, they subsequently work harder on a novel task ⁴⁷, suggesting that conducive environments can help individuals value effort ^{48,49}.

So why do people tend to avoid exerting effort even when they know that it is associated with positive, rewarding outcomes? We suggest that people have acquired associations that inadvertently encourage effort-avoidant behaviour. Rewards are often used to motivate good performance rather than instill the value of taking on effortful challenges. Many studies show that rewards improve immediate task performance ^{50–56}, but they often fail to consider how rewards could also reduce the intrinsic value of effort ^{57–59}, especially when rewards are later removed. Moreover, performance-based

rewards could make us avoid taking on challenges (and prefer easier tasks) because it is more difficult to perform well on harder tasks (and easier to perform well on easier tasks). Thus, performance-based rewards might make individuals avoid beneficial but effortful tasks (e.g., doing homework) that are challenging or invoke feelings of frustration^{60,61}.

A more promising approach is to directly reinforce behaviours that emphasise the value of engaging in effortful challenges^{62,63}. For example, when rewarded for performing harder tasks, children and rats persisted longer or worked harder on subsequent unrelated effortful tasks^{64–66}. These studies might not have teased apart or controlled for the value of performing well, but they nevertheless provided initial evidence for the idea that rewards can be used to reinforce behaviours that emphasise the value of taking on effortful challenges.

Moreover, when the act of engaging in effortful challenges is repeatedly associated with reward, effort itself can become a secondary reinforcer, making it rewarding in its own right^{19,45}. Effort becoming a secondary reinforcer might help explain why we often willingly engage in effortful tasks like endurance sports or crossword puzzles. This approach can not only potentially increase effort's value, but also provide insights into how to foster qualities like perseverance and conscientiousness⁶⁷, which might be particularly important for at-risk populations (e.g., low-income students^{68,69}).

What is emerging is a simple idea: People should become more willing to choose to perform effortful tasks if directly incentivized to do so. Crucially, this idea can be tested experimentally. To do so, we designed a three-section experimental paradigm (pre-training, training, post-training) that rewarded participants for either choosing to engage in hard tasks (*effort* condition) or performing well (*performance* condition) during training. We also included a *neutral* condition where participants received the same amount of rewards regardless of their choice and performance. This design allowed us to dissociate the effects of these three types of rewards. While the results of this experiment cannot offer strong conclusions regarding how to intervene in real-world settings⁷⁰, they could nevertheless establish that it is possible to reinforce choosing to engage in demanding tasks and provide insights into how to design future intervention studies.

We predict that rewarding willingness to choose and engage in hard tasks during training will instill in participants the value of effort, which will manifest in increased willingness to engage in hard tasks—even ones that have not been associated with rewards and in the absence of rewards. We test eight related hypotheses that provide different levels of evidence for our core idea: Participants who have been rewarded for engaging in effortful tasks (effort condition) will choose harder tasks over easier tasks more frequently than participants in the performance and neutral conditions. The outcome measures for the hypotheses are related measures of *effort preferences* (i.e., the proportion of choices whereby participants choose the hard task). Specifically, we predict the following (see Table 1 for details):

1. On *rewarded* trials during the *training* section (i.e., rewarded trials are interleaved with probe trials), participants in the effort condition will choose the hard task more frequently than those in the performance and neutral conditions (Table 1,

rows 1 and 2). This analysis mainly reveals whether participants have learned the task and reward contingencies.

2. On *probe* (i.e., unrewarded) trials during the *training* section (i.e., rewarded trials are interleaved with probe trials), participants in the effort condition will choose the hard task more frequently than those in the performance and neutral conditions (Table 1, rows 3 and 4), reflecting *within-block near-transfer* effects. Differences in effort preferences across conditions will likely reflect carryover effects (due to the interleaved reward and probe trials within a block) and actual changes in effort preferences (near-transfer).
3. On trials during the *post-training* section (i.e., completely unrewarded section), participants in the effort condition will choose the hard task more frequently than those in the performance and neutral condition (Table 1, rows 5 and 6). This analysis examines *across-block near-transfer* effects: Differences in effort preferences across conditions on these trials are less likely to be driven by carryover effects, since rewards are completely absent in this section.
4. In a separate block during the *post-training* section (i.e., completely unrewarded section), we will examine *across-block far-transfer* effects by presenting participants with a different cognitive task that had not been paired with rewards during the training section. We expect participants in the effort condition to choose the hard task more frequently than those in the performance and neutral conditions (Table 1, rows 7 and 8). Since this cognitive task has not been associated with rewards in the training section, this analysis provides evidence for whether the manipulations can lead to domain-general changes in effort preferences that transfer to unrelated tasks.

Two of the eight hypotheses in Table 1 (rows 6 and 8) most directly test our claims because they predict that in the absence of rewards (i.e., post-training section), effort preferences will be higher in the effort than neutral condition on the inhibition task (Hypothesis 6) and updating task (Hypothesis 8). If the Bayesian analyses provide positive evidence for only one of the two hypotheses, the results will not refute our overarching claims—instead, they will provide support for only across-block near-transfer (Hypothesis 5) or across-block far-transfer effects (Hypothesis 7). Nevertheless, any pattern of positive results for Hypotheses 3 to 8 should also provide some evidence for transfer effects and will be theoretically informative, given that transfer-effects are rarely reported in the literature. However, four broad patterns of results will be the most informative.

First, if we find positive results for Hypotheses 3 and 4 but not 5 to 8, the results will provide evidence for only within-block near-transfer effects. That is, changes in effort preferences are only observed during the training block but not during the post-training block. Although the lack of far-transfer effects (Hypotheses 5 to 8) is inconsistent with the core hypotheses, the within-block near-transfer effects (Hypothesis 3 and 4) suggest that longer or more intense interventions (e.g., longitudinal studies)

might be better suited to testing our hypotheses and examining the possibility of across-block near- and far-transfer effects.

Second, in addition to Hypotheses 3 and 4, we might also find support for Hypotheses 5 and 6. Positive results for these four hypotheses provide evidence for within-block (Hypotheses 3 and 4) and across-block (Hypotheses 5 and 6) effects that are limited to only the task that has been associated with rewards during training (i.e., only near- but not far-transfer). Thus, changes in effort preferences on one task carry over to the same task in a different block but do not generalise to a different task.

Third, we might obtain positive results only during the post-training block (Hypotheses 5 to 8) and not during the training block (Hypotheses 3 and 4), which provide evidence for across-block but not within-block transfer effects. If so, we speculate that consolidation processes that occur between the training and post-training sections might be necessary for changing effort preferences, and further analyses and studies will be necessary to examine this possibility.

Finally, we might find positive results only when contrasting the effort and performance conditions (Hypotheses 1, 3, 5, 7) but not the effort and neutral conditions (Hypotheses 2, 4, 6, 8). These results could suggest the presence of different processes: The reward-effort manipulation increases effort preferences but the reward-performance manipulation decreases effort preferences.

Figure 1 | Design. The experiment has pre-training, training, and post-training sections. Rewards are delivered only during the training section, but not the pre-training and post-training sections. a, The pre- and post-training sections consist of two demand selection task blocks each; these two blocks assess effort preferences on two unrelated cognitive tasks (to aid visualisation, fewer and larger dots are shown than in the actual dot-motion inhibition task). On each trial, participants choose the version of the task (easy or hard) they prefer and then perform their chosen task. Effort preference is defined as the proportion of choices whereby participants choose the hard task. b, The training section has only one block that consists of two distinct and pseudo-randomly interleaved trial types: rewarded trial ($n = 40$) and probe trial ($n = 20$). At the beginning of each trial, participants will see one of two cues (reward trial cue or probe trial cue) explicitly signalling the presence or absence of rewards on that trial. Next, they will be presented with two options (easy or hard); they will choose the version they

prefer and then perform their chosen task. If it is a rewarded trial (signalled by the reward-trial cue), participants can earn rewards or points; if it is a probe (i.e., unrewarded) trial (signalled by the probe-trial cue), participants will be fully aware that they will not receive any points (see also probe trial feedback panel), no matter their choice and performance. c, Value functions show how rewards on rewarded trials in the training section differ across the three experimental conditions. Choice (easy vs. hard) determines rewards in the effort condition; performance (correct reaction times) determines rewards in the performance condition; neither choice nor performance determines rewards in the neutral condition. Critically, the expected reward is identical across conditions. No points are given for incorrect or missed responses (reaction times deadlines are participant-specific). All conditions have identical task instructions, structure, sequence, and cues. The only difference is whether rewards are assigned based on the reward-effort value function (effort condition), reward-performance value function (performance condition), or a uniform distribution (neutral condition).

Results

Rewarded and probe trials during training (Hypotheses 1 to 4)

We found strong evidence for experimental Hypotheses 1 and 2 (Fig. 2), suggesting that participants successfully learned the task and reward contingencies. On rewarded trials in the training section, participants in the effort condition chose the hard inhibition task more frequently than those in the performance condition (H1), $b = 0.34$ [0.29, 0.38], $d = 1.22$, $BF > 100$ ($p < .001$). Effort preferences were also higher in the effort than neutral condition (H2), $b = 0.26$ [0.22, 0.31], $d = 0.95$, $BF > 100$ ($p < .001$). These effects were primarily driven by participants in the effort and performance conditions increasing and decreasing their effort preferences, respectively, relative to their preferences in the pre-training section (Fig. 3).

For Hypotheses 3 and 4 where we examined the effects of the experimental manipulation on the interleaved probe (i.e., unrewarded) trials in the training section, the Bayes factors suggest weaker evidence for our hypotheses. Effort preferences were slightly higher in the effort than performance condition (H3), $b = 0.04$ [0.00, 0.09], $d = 0.17$, $BF = 0.52$ ($p = .066$), and were also higher in the effort than neutral condition (H4), $b = 0.05$ [0.01, 0.09], $d = 0.19$, $BF = 0.83$ ($p = .036$). On these probe trials, participants in all three conditions were, on average, more likely to choose the hard inhibition task than in the pre-training section (i.e., mean difference scores were greater than 0; Fig. 3). The relatively weak evidence for these two hypotheses suggests weak carryover or “near-transfer” effects (i.e., from reward to probe trials) within the training section, but they do not preclude *across-block* transfer effects (Hypotheses 5 to 8), which also more directly test our hypotheses.

Post-training near-transfer (Hypotheses 5 to 6)

We then investigated the effects of the experimental manipulation on effort preferences in the post-training section whereby no rewards were provided for the

inhibition task (i.e., across-block “near-transfer”). Participants completed a demand selection task that was identical to the one they completed in the pre-training section (Fig. 1). Consistent with Hypotheses 5 and 6 (Fig. 2), we found that participants in the effort condition chose the hard inhibition task more frequently than those in the performance condition (H5), $b = 0.08$ [0.04, 0.12], $d = 0.35$, $BF > 100$ ($p < .001$), as well as those in the neutral condition (H6), $b = 0.04$ [0.00, 0.09], $d = 0.19$, $BF = 0.92$ ($p = .032$). As shown in Figure 3, the mean effort preference difference between the post- and pre-training sections was close to 0 for the effort condition but were negative in the neutral and performance conditions. That is, participants in the neutral and performance conditions were “depleted” or avoiding cognitive demand (see ^{14,71,72}), whereas those in effort condition had been buffered against this effect.

Post-training far-transfer (Hypotheses 7 and 8)

In the post-training section, participants also completed a separate block whereby no rewards were provided for performing a demand selection that required them to choose to perform either the easy or hard *updating* task (which they had also completed in one of the pre-training blocks, but not in the training block; see Fig. 1). Effort preferences in this block allowed us to evaluate whether there was “far-transfer” from the inhibition to updating task. The results were consistent with Hypotheses 7 and 8 (Fig. 2). Participants in the effort condition chose the hard updating task more frequently than those in the performance condition (H7), $b = 0.05$ [0.01, 0.08], $d = 0.26$, $BF = 6.19$ ($p = .004$); they also chose the hard task more frequently than those in the neutral condition (H8), $b = 0.04$ [0.01, 0.08], $d = 0.21$, $BF = 1.24$ ($p = .022$), though the effect size was smaller. Again, the pattern of results in Figure 3 suggests that participants in effort condition had been buffered from a decline in willingness to choose and engage in the hard updating task. Together, these results suggest that our manipulation led to across-task, domain-general, changes in effort preferences.

Figure 2 | Bayesian posterior densities for the effect of condition on effort preferences.

For the eight registered hypotheses (rows 1 and 2), the mass of the posterior distributions were mostly above zero, and five hypotheses had BFs greater than 1 (at least anecdotal evidence for the experimental hypothesis). Row 3 analyses were not registered and are included here only to facilitate interpretation of condition contrasts in rows 1 and 2. Positive Bayesian posterior estimates are hypothesis-consistent results (i.e., higher effort preference in the effort than the other condition). No result is consistent with the null hypothesis because no BF is smaller than the registered criteria of $BF \leq 0.3$.

Figure 3 | Effort preference relative to baseline preference in the pre-training section. Effort preference difference is computed by subtracting pre-training baseline effort preference from effort preference for rewarded trials (training section), probe trials (training section), post-training inhibition task, and post-training updating task. Zero indicates no change relative to pre-training baseline. Each dot is one participant. Mean and 95% CI are shown.

Exploratory analyses

Magnitude of rewards received influenced effort preference

We examined whether the magnitude of rewards received influenced effort preference (model specification: effort preference ~ condition * (pre-training baseline preference + mean rewards received)). Note that we only contrast the effort and performance conditions in this analysis because there was no variation in rewards

received in the neutral condition (300 points were delivered for all correct responses, regardless of choice and performance).

On rewarded trials, participants who received more rewards in the effort condition chose the hard task much more, $b = 0.66$ [0.62, 0.70], $d = 4.32$, $BF > 100$ ($p < .001$), and there was a strong interaction between condition and rewards received, $b = -0.73$ [-0.77, -0.69], $d = -4.81$, $BF > 100$ ($p < .001$), such that the positive relationship was attenuated (and became negative) for those in the performance condition. We found similar results on probe trials (positive relationship between rewards received and effort preference in the effort condition: $b = 0.21$ [0.14, 0.28], $d = 0.79$, $BF > 100$ ($p < .001$); interaction between condition and points received: $b = -0.23$ [-0.30, -0.15], $d = -0.85$, $BF > 100$ ($p < .001$)), post-training near-transfer trials (positive relationship in effort condition: $b = 0.14$ [0.08, 0.20], $d = 0.60$, $BF > 100$ ($p < .001$); interaction between condition and points received: $b = -0.16$ [-0.22, -0.10], $d = -0.71$, $BF > 100$ ($p < .001$)), and post-training far-transfer trials (positive relationship in effort condition: $b = 0.10$ [0.05, 0.14], $d = 0.55$, $BF = 29.78$ ($p < .001$); interaction between condition and points effect: $b = -0.11$ [-0.16, -0.06], $d = -0.59$, $BF > 100$ ($p < .001$)).

The consistent patterns of results above indicate opposing effects: In the effort condition, participants who, on average, received more rewards also chose the hard task *more* frequently across all sections of the task and trial types. However, in the performance condition, participants who received more rewards chose the hard task *less* frequently.

Task performance

For all eight registered hypotheses, the posterior distributions were mostly centered above zero, and five hypotheses had BFs greater than 1 (Fig. 2). These results support the idea that reward can be used to reinforce behaviours that emphasise the value of taking on effortful challenges. However, an important question is whether performance also differed across conditions (note that controlling for pre-training task accuracy led to the same conclusions regarding the efficacy of our manipulations; see Supplementary Fig. 2). Ideally, the manipulation should increase effort preferences without reducing accuracy, such that task accuracy should not differ across conditions. Alternatively, participants in the effort condition might perform worse because the manipulation could have caused them to seek rewards while sacrificing accuracy, and that they might continue to perform worse even when rewards were no longer available.

We fitted eight models (corresponding to the eight registered hypotheses; Fig. 4) to explore whether task accuracy differed between conditions (model specification: accuracy ~ condition + pre-training baseline accuracy + objective task difficulty; see Supplementary Fig. 3 for task reaction time results, and Supplementary Fig. 4 for accuracies for reaction time quartiles). For rewarded trials, accuracy was slightly higher in the effort than performance condition, $b = 0.02$ [0.00, 0.04], $d = 0.15$, $BF = 0.73$ ($p = .015$), but was similar to the neutral condition, $b = -0.01$ [-0.03, 0.00], $d = -0.09$, $BF = 0.18$ ($p = .151$). These results suggest that even though participants in the effort condition chose the hard task more frequently, they performed equally well or better than participants in the other two conditions. On probe trials, accuracy was also higher in the effort than performance condition, $b = 0.03$ [0.01, 0.05], $d = 0.16$, $BF = 0.90$ ($p = .010$), but not higher in the effort than neutral condition, $b = -0.00$ [-0.02, 0.02], $d = -$

0.01, $BF = 0.07$ ($p = .817$). These results for the probe trials were driven largely by a decrease in task accuracy in the performance condition (Fig. 4).

For the remaining four models (post-training near- and far-transfer blocks; Fig. 4), task accuracy did not differ between conditions ($bs < 0.2$, $ds < 0.11$, $BFs < 0.30$, $ps > .073$). Together with the registered Hypotheses 5 to 8 (Fig. 2), these strong null exploratory findings provide further support for the idea that reward can be used to reinforce taking on effortful challenges. Crucially, when rewards were no longer available (i.e., probe trials during training section, post-training section), participants in the effort condition did not perform worse than those in the performance and neutral conditions, even though they were performing the hard task more frequently.

Figure 4 | Exploratory analyses and Bayesian posterior densities for the effect of condition on task accuracy. Positive Bayesian posterior estimates indicate higher accuracy in the effort than the other condition. Model specification: accuracy $\sim$ condition + pre-training baseline accuracy + objective task difficulty.

Figure 5 | Effort preference correlations. Effort preferences for different trial types in different study sections correlated strongly with one another.

Individual differences

In all three conditions, participants' effort preferences for the inhibition task in the pre-training, training (reward and probe trials), and post-training sections correlated positively ($r_s > .35$, $p_s < .001$; Fig. 5). For example, those who chose the hard task more frequently in the pre-training block also chose the hard task more frequently in training and post-training sections. We fitted eight models (corresponding to the eight registered hypotheses; Supplementary Fig. 5) to explore whether condition and pre-training baseline effort preference interacted to predict effort preference, and found some evidence for interaction for Hypotheses 1 and 5 ($BFs > 1$). The effect of condition (effort versus performance) was stronger for participants with higher effort preferences at baseline on reward trials in the training section, $b = 0.21$ [0.02, 0.38], $d = 0.75$, $BF = 1.02$ ($p = .025$), and on the post-training inhibition trials, $b = 0.29$ [0.13, 0.45], $d = 1.23$, $BF = 43.55$ ($p < .001$), but not on the post-training updating trials ($BFs \leq 0.06$). That is, participants in the effort (vs. performance) condition who chose the hard task more frequently during pre-training also chose the hard task more frequently on the rewarded trials and post-training inhibition block.

To further identify heterogeneous treatment effects, we fitted 8 causal forests^{73,74} (one for each hypothesis; Supplementary Fig. 6). This non-parametric algorithm allows for data-driven covariate selection and flexible modelling of interactions in high dimensions, which guard against spurious heterogeneity or interactions common in classical approaches⁷⁵.

Following previous work⁷⁶, we first trained (separately for each hypothesis) a pilot forest on seven covariates (age, pre-training effort preference, Need for Cognition⁷⁷, Distress Tolerance⁷⁸, Conscientiousness⁷⁹, Grit⁸⁰, and Implicit Theories of Intelligence⁸¹). We then trained a second forest on the top three covariates that saw a reasonable number of splits in the first forest, which allows the second forest to make more splits on the most important covariates. Supplementary Figure 6 shows, for each hypothesis, how the conditional average treatment effects vary (mostly non-linearly) as a function of each covariate while holding constant the other two covariates at their mean values.

For the reward trials (Hypotheses 1 and 2), higher pre-training baseline and Need For Cognition were associated with stronger treatment effects, but the intervention was less effective for older participants. For the probe trials (Hypotheses 3 and 4), higher Need for Cognition, Grit, and age was associated with *weaker* treatment effects (i.e., the intervention was less effective for participants with higher values for these three covariates). For the post-training inhibition trials (Hypotheses 5 and 6), higher pre-training baseline was associated with stronger treatment effects (effects for other covariates were less clear). Finally, for the post-training updating (far-transfer) trials (Hypotheses 7 and 8), the effects of the covariates on the treatment effect were less consistent (Supplementary Fig. 6). Together, these results suggest the manipulation may be particularly effective for certain people—especially those who are already predisposed to engaging in harder tasks.

Discussion

Together, the eight registered analyses provide evidence for the idea that people can learn to value effort^{19,82,83}. The strong results for Hypotheses 1 and 2 ($BFs > 100$) indicate that participants successfully learned the task and reward contingencies in our paradigm. When rewarded for choosing hard over easy tasks, participants actively sought out the hard task, and importantly, did not sacrifice task performance.

For the remaining six models, the Bayesian analyses found the weakest evidence (but the mass of the posterior distributions were still mostly positive) for Hypotheses 3 and 4 ($BFs = 0.52, 0.83$), which predicted that effort preferences on the probe (i.e., unrewarded) trials during the training section would be higher in the effort condition (relative to the other two conditions). We found limited evidence for within-block near-transfer effects: Increases in effort preferences on rewarded trials may have somewhat carried over or "transferred" to the probe trials, though we would have expected the carryover effects to be stronger considering that the probe trials were randomly interleaved with rewarded trials. These relatively weak effects suggest that participants in the effort condition—who had chosen to perform the hard task more frequently—might be strategically "easing off" on the probe trials (e.g., labour-leisure trade-off^{84,85}). The exploratory finding that the treatment effects on probe trials were *weaker* for participants with higher Need for Cognition and Grit scores is consistent with the idea of strategic "easing off".

The evidence for Hypotheses 5 and 6 were stronger ($BFs > 100, 0.92$), suggesting that the positive effects of rewarding effort on effort preference were stronger in the post-training section where rewards were entirely absent. These results suggest two possibilities: First, consolidation may have occurred between the training and post-training sections to facilitate these across-block transfer effects; second, there was less dynamic, trial-by-trial strategic "easing off" in this section (relative to the training section), thereby allowing treatment effects to reveal themselves more clearly. Future work is needed to understand these possibilities, which will provide insights into how our manipulation increases people's willingness to engage in cognitive work over time (e.g., changes in cost-benefit valuation^{86–91}).

We also found evidence consistent with Hypotheses 7 and 8 ($BFs = 6.19, 1.24$), which focused on effort preferences on the updating task. Because participants did not train on this task, which also required a different set of cognitive processes than the trained inhibition task, this task served as a "far-transfer" measure⁹². In other words, rewarding participants for choosing the hard inhibition task in the training section made them more willing to engage with a hard task that requires a different set of cognitive processes (see also⁹³).

Hypotheses 6 and 8 test our claims most stringently (effort versus neutral condition contrast in the post-training section for the inhibition and updating tasks, respectively), and the evidence for these two hypotheses are weaker but still positive. Nevertheless, as stated in our registered analytic plan, the generally positive results across hypotheses—if we consider not just the Bayes factors but also the posterior distributions, effect sizes, and frequentist p-values—provide evidence for both training (Hypotheses 1 to 4) and transfer (Hypotheses 3 to 8) effects. The latter is noteworthy given that transfer-effects have been quite elusive in prior work^{34,35,33,32, but see⁹⁴}. We also

found that in the post-training sections (both inhibition and updating tasks), participants in the performance condition chose the hard task less frequently relative to their own pre-training choices and participants in the effort and neutral conditions.

The condition differences in effort preference across study sections indicate that participants in the performance and neutral conditions became less willing (relative to their own pre-training baseline) to choose the hard task (inhibition and updating tasks) over time ¹⁴. Performance-contingent rewards may have reduced the intrinsic value of effort ^{57–59}, especially when rewards were removed in the post-training sections, resulting in choice patterns resembling “depletion” ⁷¹. That is, even though performance-contingent rewards reliably affect task performance in many other studies ^{95,96}, such rewards might also encourage effort-avoidant behaviour. Most importantly, participants in the effort condition mostly maintained their pre-training effort preferences in the post-training sections, suggesting that rewarding effort may be preferable if the goal is to reduce effort-avoidance.

Despite the overall positive evidence for our hypotheses, there is heterogeneity in treatment effects. Results from exploratory linear interaction models and non-parametric causal forests ^{73,74} suggest the rewarding effort may be the most effective for those who may already be predisposed to choosing harder tasks prior to the experiment; note however, that the treatment effects were still positive for those who were less predisposed. Future work is needed to more precisely identify factors that moderate treatment effects, which can help design tailored short- and long-term interventions in practical settings such as schools.

The main limitations of our study were the resource limitations (total registered N = 750) and the relatively brief manipulation (i.e., the 40 rewarded trials in the training section lasted only about 7 minutes). Extending the present study to longitudinal and developmental studies in both quasi-experimental or field settings with stronger and repeated manipulations is a promising direction for future work.

Our results suggest that directly reinforcing behaviours that emphasise the value of engaging in effortful challenges can change people's preferences for challenge and exerting effort. Effort does not always have to feel aversive ⁷²—if rewarded appropriately over time, effort could even feel valuable and meaningful ^{19,97}. People might find more value and meaning in exerting effort if they have been exposed to environments that reinforce the value of effort ⁸². Our reward-effort manipulation easily translates into interventions that can be implemented in educational and clinical settings. For example, schools and universities could operate mostly on binary pass-fail grading systems (effort condition in our study) instead of the letter-grade system (performance condition), and students and children should be encouraged to and rewarded for attempting challenging tasks ^{43,47,98}. Such an approach could help individuals learn the value of effort, develop resilience and self-regulatory skills, and improve outcomes for individuals and societies.

Methods

Participants

846 participants completed the study ($n_{\text{effort}} = 282$, $n_{\text{performance}} = 281$, $n_{\text{neutral}} = 283$). Following our registered sampling plan (Bayes factor > 10 for Hypotheses 6 and 8, or a total of 750 participants), we excluded 85 participants (see Sampling plan section for exclusion details). The final sample size was 761 ($n_{\text{effort}} = 254$, $n_{\text{performance}} = 255$, $n_{\text{neutral}} = 252$).

Ethics information

The research complied with the University of Toronto Research Ethics in Human Research Unit's regulations. All participants provided informed consent in accordance with the regulations of the ethics unit. Participants were recruited via the online recruitment platform Prolific and were paid approximately £9.66/hr and received a cash bonus (between £0.50 and £5), which was determined by the number of points earned.

Pilot data

Results from a pilot study using the dot-motion inhibition task demonstrated the feasibility of the proposed paradigm because the value functions (Fig. 1c) led to different effort preferences that are consistent with Hypothesis 1. Participants in the effort condition ($n = 59$) had higher effort preferences than those in the performance condition ($n = 60$), $b = 13.89$, 95% HPD = [5.62, 22.06], $d = 0.63$, $BF = 29.79$ ($p < .001$). That is, participants who experienced the reward-effort value function (Fig. 1c, left panel) chose the hard version of the dot-motion inhibition task more frequently than those who experienced the reward-performance value function (Fig. 1c, right panel). See Supplementary Figure 1.

Design

This experimental paradigm used a mixed (between-within subjects) design. All participants completed three sections in this order (*within*-subject): pre-training, training, and post-training. The pre-training section presented participants with two unrelated cognitive tasks in two separate blocks to evaluate participants' baseline effort preferences on these two tasks (Fig. 1a). Next, the training section consisted of a single block and was the only block in the experiment that delivered rewards (Fig. 1b). Participants were randomly assigned (*between*-subject assignment) to the effort, performance, or neutral condition. Finally, the post-training section had two blocks that assessed whether the reward manipulation (Fig. 1c) during the previous training section affected effort preferences on the two tasks presented during pre-training (Fig. 1a). Data collection was performed blind to the conditions of the experiment; data analysis, however, was not performed blind.

Pre-training section and post-training section

Both the pre-training the post-training sections each consisted of two blocks of demand selection tasks ¹⁴ (Fig. 1a). In each section, two cognitive tasks requiring primarily inhibition and updating abilities were presented in two separate blocks (order counterbalanced across participants). Each pre-training block had 40 trials. Behavioural indices (e.g., choices, accuracy, reaction time) obtained from the two pre-training blocks provided baseline measures of effort preferences and performance on the two cognitive tasks. The post-training section (presented after the training section) was similar to the pre-training section, but only had 20 trials in each block. This block allowed us to measure how behaviour changes as a function of our experimental manipulation in the training section. On each trial in each block, participants pressed either the left or right key to select the cue shown on the left or right of the display (3000 ms response deadline), which represented either the easy or hard version of the cognitive task. Participants then performed the task they selected, and no performance feedback was provided at the end of each trial (Fig. 1a). The mappings between cues and task difficulty were counterbalanced across participants and the locations (left or right) of the cues were randomly determined on every trial.

The inhibition task in one of the demand selection blocks was a Simon-like dot-motion conflict task ^{99,100}. After choosing a cue (reflecting either the easy or hard version of the inhibition task), participants performed three repetitions of the selected task (Fig. 1a). On each repetition, participants saw an array of colored moving dots and had up to 1500 ms to respond. Depending on their choice (easy or hard task), they had to press either the left or right key to indicate the dot motion *direction* (leftward or rightward; 300 dots with 75-100% motion coherence, sampled from a uniform distribution) or the *colour* of the dots (presented in one of four colours, with two colours mapped to each key). Another 20 to 50 (sampled from a uniform distribution) distractor dots in a different colour moved in a direction that was consistent with or opposite to the majority of the dots. The easy version required little controlled attention because participants simply had to indicate the motion direction of the majority of the dots while ignoring the colour of the dots. However, the hard version required controlled attention because participants had to indicate the colour of the dots while overriding their automatic tendency to indicate motion direction; critically, the key for the correct colour response could be congruent with the dot motion direction or it could be incongruent with the dot motion direction (65% chance it was incongruent on each repetition). Colour-response mappings were counterbalanced across participants.

The updating task in the other demand selection block was a working memory and attention control task ¹⁰¹. After choosing one of the cues (reflecting either the easy or hard version of the updating task), participants performed the selected task (Fig. 1a). On each trial, participants added a digit to three serially presented digits. The easy version required participants to add 0 to each digit (e.g., 7, 8, 6) and to use the left, right, up, or down key to choose the correct response (i.e., 786) out of four similar responses (3000 ms response deadline). The hard version required participants to add 3 or 4 to each digit (e.g., add 4 to each digit: 7 becomes 1, 8 becomes 2, 6 becomes 0, etc.); two digits were used for the hard task to minimise practice effects and the digit was selected randomly at the start of each trial.

Training section

In the training section, participants completed one block (60 trials) of the demand selection task with the inhibition task (but not the updating task; see explanation below). At the beginning of each trial, participants saw one of two cues—reward trial cue or probe trial cue (Fig. 1b)—explicitly signalling to participants the presence or absence of rewards on that trial, respectively. Next, participants chose to perform either the easy or hard version of the inhibition task (represented by the same cues seen in the pre-training section), and performed their chosen task. Finally, they saw either the reward or probe trial feedback, depending on whether it was a reward or probe trial, respectively. Rewarded ($n = 40$) and probe ($n = 20$) trials were pseudo-randomly interleaved in this block, such that the first trial was always a reward trial and probe trials occurred after every one to three reward trials.

To test our hypotheses, we randomly assigned participants to the effort, performance, or neutral condition. The three conditions had identical task instructions, structure, sequence, and cues (reward, probe, feedback); the only difference was how the points were delivered on rewarded trials (Fig. 1b, Reward trial cue). On these trials, participants in the effort condition experienced the reward-effort value function (Fig. 1c, first panel). Those in the performance condition experienced the reward-performance function (Fig. 1c, second panel). Those in the neutral condition received the same amount of rewards for choosing the easy or hard task and regardless of their performance (Fig. 1c, third panel).

Specifically, on rewarded trials (Fig. 1b, reward trial cue), participants in the effort condition received more points for choosing the hard relative to the easy task (e.g., 370 vs. 230 points, respectively, plus jitter drawn from a normal distribution, $N(\mu = 0, \sigma = 5)$), regardless of their reaction times (RTs). However, participants in the performance condition received points that scaled with their own RTs (plus jitter drawn from the same normal distribution as above), regardless of their choice (easy or hard version of the task). Participants in the neutral condition received the same points (plus jitter) regardless of their choice or performance. Critically, the expected reward across conditions was identical. Because participants performed three repetitions of the inhibition task after each choice (Fig. 1b), the feedback (Fig. 1b, reward trial feedback panel) received at the end of each rewarded trial was the mean of the points received on three repetitions of the inhibition task. On rewarded trials, 0 points were given for incorrect or missed responses (RT deadlines were participant-specific; see next paragraph).

Participant-specific RT criteria were used to allocate points on rewarded trials, ensuring participants' RTs were evaluated against their own RT benchmarks on the inhibition task. These benchmarks were based on each participant's RTs across the easy and hard versions of the task. In addition, if participants responded incorrectly or too slowly, they received 0 points (i.e., no “empty praises”). This approach ensured incentive compatibility and that participants had to perform relatively well, even when RT performance was uncorrelated with points earned (Fig. 1c, first and third panels). That is, choosing the hard task and then slacking off was a bad strategy that likely resulted in no rewards.

The distribution of RTs used for benchmarking included RTs on only correct responses (i.e., correct RTs) of the inhibition task in the pre-training section demand

selection task (Fig. 1a); very fast and slow correct RTs (± 1 times the median absolute deviation) were also excluded¹⁰². RT deadlines were determined separately for each participant, and the easy and hard versions of the inhibition task had the same deadline (i.e., maximum RT [after applying exclusion criteria above] plus 150 ms).

This RT-benchmarking procedure helped to ensure each participant had neither too much nor too little time to respond on each trial. For example, two participants—one fast, one slow (with median RTs of 500 ms and 700 ms, respectively)—would have different RT deadlines (e.g., 750 ms and 950 ms respectively), but both would receive the same number of points on a trial if they responded correctly at their own median RTs (500 ms and 700 ms, respectively). Thus, task difficulty and rewards were tailored to each participant's ability, such that participants received, on average, the same number of points, regardless of individual differences in average RTs.

We chose to train participants on only the inhibition task but not the updating task for two reasons. First, our paradigm required participants to feel efficacious on reward-performance trials (Fig 1c, middle panel)—exerting more effort should lead to more accurate and faster responses, and the inhibition task allowed for a much tighter coupling between performance and rewards obtained than the updating task. Second, we are interested in whether the effects of training on a highly controlled inhibition task generalise to the updating task, which is more ecologically valid because solving mathematics problems resemble real-life problems more than indicating the motion or colour of dots.

Procedure

Participants were recruited to take part in a study titled “What are your cognitive preferences and abilities?” and saw the following study description: “We are examining how you make decisions and perform different cognitive tasks. You’ll be doing cognitive tasks on the computer and answering a few questions about yourself after completing the tasks.”

At the beginning of the experiment, participants were told that they were adventurous space explorers who had to complete various missions in space. Before the pre-training section (Fig. 1a), participants practised and learned the four different cues associated with the easy and hard versions of the inhibition and updating tasks. Each practice block (maximum 80 trials per block) terminated when participants either completed all trials or performed well above chance levels in the preceding 20 trials ($\geq 80\%$ and 70% correct in the last 20 trials on the inhibition and updating tasks, respectively; see exclusion criteria in Sampling plan section below). During practice, participants were encouraged to be adventurous by trying different tasks instead of always choosing one task. Participants then completed the two actual blocks in the pre-training section (Fig. 1a).

After which, they began the training section, where they first learned the cues associated with the reward and probe trials (Fig. 1b): “If you see a landed spaceship, an alien will be delivering rewards. If the spaceship hasn’t landed, you WON’T be receiving rewards.” They were told that in the upcoming mission, their “goal is to earn as many points as possible” and that the earned points would be converted to a cash bonus at the end of the study. They also read the following instructions: “How many points you could earn from the aliens depends on some combination of WHICH ROCKET YOU

CHOOSE or/and HOW WELL YOU PERFORM (accuracy and reaction time). If you respond inaccurately or too slowly, you will receive 0 points for that response. So if the alien gave you very few points, it's likely because you made too many mistakes or/and were slow. Therefore, to maximize your earnings, try to use the feedback/points you receive from the aliens to improve HOW WELL YOU PERFORM and inform WHICH ROCKET TO CHOOSE in the future". After completing 10 practice trials (5 rewarded trials, 5 probe trials) where they were encouraged to explore different strategies, they completed the actual training section (Fig. 1b).

Before beginning the post-training section, participants were briefly reminded of the four cues (from the pre-training section) associated with the easy and hard versions of the inhibition and updating tasks. To inform participants that no rewards would be delivered in this section, they were explicitly told that "the aliens have retreated so they won't be around to deliver points or rewards" before they completed the two blocks in this section.

Finally, they completed demographic, personality, and debriefing questionnaires: Need for Cognition ⁷⁷, Distress Tolerance ⁷⁸, Conscientiousness ⁷⁹, Grit ⁸⁰, and Implicit Theories of Intelligence ⁸¹. They were then compensated for their participation and points earned during the training section.

Sampling plan

We used the Bayes Factor Design Analysis approach ¹⁰³ to determine the number of participants required to provide compelling evidence for Hypotheses 6 and 8 (Table 1, Design Table), which predicted that relative to participants in the neutral condition, those in the effort condition would choose the hard over the easy version of the inhibition task (Hypothesis 6) and updating task (Hypothesis 8) more frequently in the post-training section. This criterion reflected our belief that these two hypotheses tested our ideas most directly.

We used the sequential design with maximum participant approach to recruit additional participants until either (a) the Bayes factor (BF) provided strong evidence for the null hypothesis ($BF < 0.1$) or alternative hypothesis ($BF > 10$), or (b) a total of 750 participants had been reached. That is, the BF for Hypotheses 6 and 8 must each exceed one of the thresholds; if not, we stopped only after we have recruited 750 participants in total. We computed BFs using the BayesFactor package ¹⁰⁴ for the R Environment for Statistical Computing ¹⁰⁵, and used the default JZS prior ¹⁰⁶. We compared the full linear regression model (outcome ~ condition + baseline) against the null model (outcome ~ baseline) to determine whether the BF had exceeded either threshold.

Participants who failed to meet the criteria below were excluded from the sequential sampling procedure and all analyses. They were excluded if they performed poorly during pre-training practice blocks. Specifically, those who did not respond with at least 80% and 70% correct in the last 20 practice trials on the inhibition and updating tasks, respectively, were excluded from all analyses ($n = 26$). We also excluded participants who performed better (i.e., faster median correct RT) on the hard relative to the easy version of either of the two cognitive tasks in the last 20 practice trials ($n = 61$). This criterion was necessary because our paradigm hinged on the fact that the hard

tasks should require more effort (i.e., as reflected in slower median correct RT) than the easy tasks, and thus, more rewards would be necessary to offset the costs associated with performing the hard tasks.

To further ensure data quality was high, we asked participants to indicate honestly the extent to which they had tried to follow all task instructions on a continuous slider scale with three equally-spaced anchors (never, mostly, always); participants who responded below the midpoint of the scale (mostly) were excluded from all analyses ($n = 1$). As a final check, participants were asked to recall the four tasks (easy and hard versions of the inhibition and updating tasks) and described what the corresponding cues looked like; participants who failed to describe the cues correctly were also excluded from all analyses ($n = 4$).

Analysis plan

We excluded trials with no responses (i.e., because participants failed to respond in time) before performing any aggregation. To compute effort preferences on the different trial types in the different blocks, we calculated for each participant the proportion of choices whereby they chose the hard task. To test each hypothesis in Table 1, we fitted a Bayesian linear regression to test the effect of condition (effort versus performance or neutral condition) on effort preferences. To obtain the direct effect¹⁰⁷ of condition, we included effort preferences in the pre-training demand selection blocks as a covariate (i.e., linear regression with one covariate and the condition regressor; i.e., analysis of covariance). Specifically, effort preferences on the pre-training *inhibition* task was included as the covariate in Hypotheses 1 to 6, whereas effort preferences on the pre-training *updating* task was the covariate in Hypotheses 7 and 8.

We fitted the models with the R package brms¹⁰⁸ and reported the following statistics that were calculated from the posterior samples: beta estimate, its Bayesian 95% highest-posterior-density (HPD) interval, and Cohen's d effect size. We used default brms priors (i.e., improper flat priors over the reals) because with sufficient data, priors are unlikely to influence parameter estimates (but they always influence the BF). For each effect, we also reported the BF, which was computed using the BayesFactor package and the default JZS prior^{104,106}. $BF = 1$ indicates the data do not favour either the experimental or null hypothesis. BFs between 3 and 10 provide moderate evidence for the experimental hypothesis, whereas BFs between 0.3 and 0.1 provide moderate evidence for the null hypothesis. BFs greater than 10 or smaller than 0.1 provide at least strong evidence for the experimental and null hypothesis, respectively¹⁰³. We also reported frequentist probability values.

Protocol Registration

Stage 1 protocol: <https://doi.org/10.6084/m9.figshare.14230283.v1>

Data availability statement

Data and materials repository:

https://osf.io/9unj5/?view_only=b03a5e40151346f7bd0f5e66c588a047

Code availability statement

Code repository:

https://osf.io/9unj5/?view_only=b03a5e40151346f7bd0f5e66c588a047

References

1. Disabato, D. J., Goodman, F. R. & Kashdan, T. B. Is grit relevant to well-being and strengths? Evidence across the globe for separating perseverance of effort and consistency of interests. *Journal of Personality* **87**, 194-211 (2019).
2. Shoda, Y., Mischel, W. & Peake, P. K. Predicting adolescent cognitive and self-regulatory competencies from preschool delay of gratification: Identifying diagnostic conditions. *Developmental Psychology* **26**, 978-986 (1990).
3. Watts, T. W., Duncan, G. J. & Quan, H. Revisiting the marshmallow test: A conceptual replication investigating links between early delay of gratification and later outcomes. *Psychological Science* **29**, 1-19 (2018).
4. Credé, M., Tynan, M. C. & Harms, P. D. Much ado about grit: A meta-analytic synthesis of the grit literature. *Journal of Personality and Social Psychology* **113**, 492-511 (2017).
5. de Ridder, D. T., Lensvelt-Mulders, G., Finkenauer, C., Stok, F. M. & Baumeister, R. F. Taking stock of self-control: a meta-analysis of how trait self-control relates to a wide range of behaviors. *Personality and Social Psychological Review* **16**, 76-99 (2012).
6. Moffitt, T. E. et al. A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences* **108**, 2693-2698 (2011).
7. Inzlicht, M., Werner, K. M., Briskin, J. L. & Roberts, B. W. Integrating models of self-regulation. *Annual review of psychology* **72**, 319-345 (2021).
8. Jackson, C. K. What do test scores miss? The importance of teacher effects on non-test score outcomes. *Journal of Political Economy* **126**, 2072-2107 (2018).
9. Kwon, H. W. The sociology of grit: Exploring grit as a sociological variable and its potential role in social stratification. *Sociology Compass* **11**, e12544 (2017).
10. Roberts, B. W., Kuncel, N. R., Shiner, R., Caspi, A. & Goldberg, L. R. The power of personality: The comparative validity of personality traits, socioeconomic status, and cognitive ability for predicting important life outcomes. *Perspectives on Psychological Science* **2**, 313-345 (2007).
11. Tough, P. How kids learn resilience. *The Atlantic* **317**, 56-66 (2016).
12. Hull, C. L. *Principles of behavior: An introduction to behavior theory* (Appleton-Century-Crofts, 1943).
13. Kool, W. & Botvinick, M. Mental labour. *Nature Human Behaviour* **2**, 899-908 (2018).
14. Kool, W., McGuire, J. T., Rosen, Z. B. & Botvinick, M. M. Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General* **139**, 665-682 (2010).
15. Kurzban, R. The sense of effort. *Current Opinion in Psychology* **7**, 67-70 (2016).
16. Shenhav, A. et al. Toward a rational and mechanistic account of mental effort. *Annual Review of Neuroscience* **40**, 99-124 (2017).
17. Westbrook, A., Kester, D. & Braver, T. S. What is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. *PLoS One* **8**, e68210 (2013).

18. Friedrich, A. M. & Zentall, T. R. Pigeons shift their preference toward locations of food that take more effort to obtain. *Behavioural Processes* **67**, 405-415 (2004).
19. Inzlicht, M., Shenhav, A. & Olivola, C. Y. The effort paradox: Effort is both costly and valued. *Trends in Cognitive Sciences* **22**, 337-349 (2018).
20. Lydall, E. S., Gilmour, G. & Dwyer, D. M. Rats place greater value on rewards produced by high effort: An animal analogue of the "effort justification" effect. *Journal of Experimental Social Psychology* **46**, 1134-1137 (2010).
21. Norton, M. I., Mochon, D. & Ariely, D. The IKEA effect: When labor leads to love. *Journal of Consumer Psychology* **22**, 453-460 (2012).
22. Dunn, T. L., Gaspar, C. & Risko, E. F. Cue awareness in avoiding effortful control. *Neuropsychologia* **123**, 77-91 (2019).
23. Dunn, T. L., Inzlicht, M. & Risko, E. F. Anticipating cognitive effort: roles of perceived error-likelihood and time demands. *Psychological Research* **83**, 1033-1056 (2017).
24. Constantinidis, C. & Klingberg, T. The neuroscience of working memory capacity and training. *Nature Reviews Neuroscience* **17**, 1-12 (2016).
25. Enriquez-Geppert, S., Huster, R. J. & Herrmann, C. S. Boosting brain functions: Improving executive functions with behavioral training, neurostimulation, and neurofeedback. *International Journal of Psychophysiology* **88**, 1-16 (2013).
26. Jaeggi, S. M., Buschkuhl, M., Jonides, J. & Shah, P. Short- and long-term benefits of cognitive training. *Proceedings of the National Academy of Sciences* **108**, 10081-10086 (2011).
27. Melby-Lervåg, M. & Hulme, C. Is working memory training effective? A meta-analytic review. *Developmental Psychology* **49**, 270-291 (2013).
28. Melby-Lervag, M., Redick, T. S. & Hulme, C. Working memory training does not improve performance on measures of intelligence or other measures of "far transfer": Evidence from a meta-analytic review. *Perspectives on Psychological Science* **11**, 512-534 (2016).
29. Owen, A. M. et al. Putting brain training to the test. *Nature* **465**, 775-778 (2010).
30. Shipstead, Z., Redick, T. S. & Engle, R. W. Is working memory training effective? *Psychological Bulletin* **138**, 628-654 (2012).
31. Simons, D. J. et al. Do "brain-training" programs work? *Psychological Science in the Public Interest* **17**, 103-186 (2016).
32. von Bastian, C. C., Guye, S. & De Simoni, C. How strong is the evidence for the effectiveness of working memory training. *Cognitive and Working Memory Training: Perspectives from Psychology, Neuroscience, and Human Development* **58**, 1-23 (2019).
33. Miles, E. et al. Does self-control improve with practice? Evidence from a six-week training program. *Journal of Experimental Psychology: General* **145**, 1075-1091 (2016).
34. Friese, M., Frankenbach, J., Job, V. & Loschelder, D. D. Does self-control training improve self-control? A meta-analysis. *Perspectives on Psychological Science* **12**, 1077-1099 (2017).
35. Lee, B. M. & Kemmelmeier, M. How reliable are the effects of self-control training?: A re-examination using self-report and physical measures. *PLoS One* **12**, e0178814 (2017).

36. Alan, S., Boneva, T. & Ertac, S. Ever failed, try again, succeed better: Results from a randomized educational intervention on grit. *The Quarterly Journal of Economics* 1-130 (2019).
37. Lazowski, R. A. & Hulleman, C. S. Motivation interventions in education: A meta-analytic review. *Review of Educational Research* **86**, 602-640 (2016).
38. Yeager, D. S. et al. Using Design Thinking to Improve Psychological Interventions: The Case of the Growth Mindset During the Transition to High School. *Journal of Educational Psychology* **108**, 374-391 (2016).
39. Yeager, D. S. et al. Teaching a lay theory before college narrows achievement gaps at scale. *Proceedings of the National Academy of Sciences* **113**, E3341-8 (2016).
40. Foliano, F., Rolfe, H., Buzzeo, J., Runge, J. & Wilkinson, D. Changing mindsets: effectiveness trial. *National Institute of Economic and Social Research* (2019).
41. Li, Y. & Bates, T. Does growth mindset improve children's IQ, educational attainment or response to setbacks? Active-control interventions and data on children's own mindsets. (2017).
42. Sisk, V. F., Burgoyne, A. P., Sun, J., Butler, J. L. & Macnamara, B. N. To what extent and under which circumstances are growth mind-sets important to academic achievement? Two meta-analyses. *Psychological Science* **29**, 549-571 (2018).
43. Yeager, D. S. et al. A national experiment reveals where a growth mindset improves achievement. *Nature* **573**, 364-369 (2019).
44. Crespi, L. P. Quantitative variation of incentive and performance in the white rat. *The American Journal of Psychology* **55**, 467-517 (1942).
45. Eisenberger, R. Learned industriousness. *Psychological Review* **99**, 248-267 (1992).
46. O'Doherty, J. P., Cockburn, J. & Pauli, W. M. Learning, Reward, and Decision Making. *Annual Review of Psychology* **68**, 73-100 (2017).
47. Leonard, J. A., Lee, Y. & Schulz, L. E. Infants make more attempts to achieve a goal when they see adults persist. *Science* **357**, 1290-1294 (2017).
48. Gunderson, E. A. et al. Parent praise to 1- to 3-year-olds predicts children's motivational frameworks 5 years later. *Child Development* **84**, 1526-1541 (2013).
49. Lucca, K., Horton, R. & Sommerville, J. A. Infants rationally decide when and how to deploy effort. *Nature Human Behaviour* **4**, 372-379 (2020).
50. Boksem, M. A. & Tops, M. Mental fatigue: Costs and benefits. *Brain Research Reviews* **59**, 125-139 (2008).
51. Botvinick, M. & Braver, T. Motivation and cognitive control: From behavior to neural mechanism. *Annual Review of Psychology* **66**, 83-113 (2015).
52. Kurniawan, I. T., Guitart-Masip, M., Dayan, P. & Dolan, R. J. Effort and valuation in the brain: the effects of anticipation and execution. *Journal of Neuroscience* **33**, 6160-6169 (2013).
53. Massar, S. A., Lim, J., Sasmita, K. S. & Chee, M. W. Rewards boost sustained attention through higher effort: A value-based decision making approach. *Biological Psychology* **120**, 21-27 (2016).
54. Padmala, S. & Pessoa, L. Reward reduces conflict by enhancing attentional control and biasing visual cortical processing. *Journal of Cognitive Neuroscience* **23**, 3419-3432 (2011).

55. Umemoto, A. & Holroyd, C. B. Task-specific effects of reward on task switching. *Psychological Research* **79**, 698-707 (2015).
56. Varazzani, C., San-Galli, A., Gilardeau, S. & Bouret, S. Noradrenaline and dopamine neurons in the reward/effort trade-off: a direct electrophysiological comparison in behaving monkeys. *Journal of Neuroscience* **35**, 7866-7877 (2015).
57. Deci, E. L., Koestner, R. & Ryan, R. M. A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin* **125**, 627-668 (1999).
58. Kohn, A. *Punished by rewards: The trouble with gold stars, incentive plans, A's, praise, and other bribes*. (Houghton Mifflin Harcourt, Boston, MA, 1999).
59. Tang, S. & Hall, V. C. The overjustification effect: A meta-analysis. *Applied Cognitive Psychology* **9**, 365-404 (1995).
60. Henderlong Corpus, J. & Lepper, M. R. The effects of person versus performance praise on children's motivation: Gender and age as moderating factors. *Educational Psychology* **27**, 487-508 (2007).
61. Mueller, C. M. & Dweck, C. S. Praise for intelligence can undermine children's motivation and performance. *Journal of Personality and Social Psychology* **75**, 33-52 (1998).
62. Amsel, A. Frustrative nonreward in partial reinforcement and discrimination learning: Some recent history and a theoretical extension. *Psychological Review* **69**, 306-328 (1962).
63. McCuller, T., Wong, P. T. P. & Amsel, A. Transfer of persistence from fixed-ratio barpress training to runway extinction. *Animal Learning & Behavior* **4**, 53-57 (1976).
64. Eisenberger, R., Carlson, J. & Frank, M. Transfer of persistence to the acquisition of a new behaviour. *Quarterly Journal of Experimental Psychology* **31**, 691-700 (1979).
65. Eisenberger, R., Carlson, J., Guile, M. & Shapiro, N. Transfer of effort across behaviors. *Learning and Motivation* **10**, 178-197 (1979).
66. Eisenberger, R., Mitchell, M. & Masterson, F. A. Effort training increases generalized self-control. *Journal of Personality and Social Psychology* **49**, 1294-1301 (1985).
67. Göllner, R. et al. Is doing your homework associated with becoming more conscientious. *Journal of Research in Personality* **71**, 1-12 (2017).
68. Claro, S., Paunesku, D. & Dweck, C. S. Growth mindset tempers the effects of poverty on academic achievement. *Proceedings of the National Academy of Sciences* **113**, 8664-8668 (2016).
69. Tough, P. *How children succeed: Grit, curiosity, and the hidden power of character* (Houghton Mifflin Harcourt, 2012).
70. Lin, H., Werner, K. M. & Inzlicht, M. Promises and perils of experimentation: The mutual-internal-validity problem. *Perspectives on Psychological Science* **16**, 854-863 (2021).
71. Lin, H., Saunders, B., Friese, M., Evans, N. J. & Inzlicht, M. Strong effort manipulations reduce response caution: A preregistered reinvention of the ego-depletion paradigm. *Psychological Science* 1-17 (2020).

72. Wu, R., Ferguson, A. & Inzlicht, M. Do humans prefer cognitive effort over doing nothing? (2021).
73. Athey, S., Tibshirani, J. & Wager, S. Generalized random forests. *The Annals of Statistics* **47**, (2019).
74. Athey, S. & Wager, S. Estimating Treatment Effects with Causal Forests: An Application. *Observational Studies* **5**, 37-51 (2019).
75. Wager, S. & Athey, S. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* **113**, 1228-1242 (2018).
76. Basu, S., Kumbier, K., Brown, J. B. & Yu, B. Iterative random forests to discover predictive and stable high-order interactions. *Proc Natl Acad Sci U S A* **115**, 1943-1948 (2018).
77. Cacioppo, J. T. & Petty, R. E. The need for cognition. *Journal of Personality and Social Psychology* **42**, 116 (1982).
78. Simons, J. S. & Gaher, R. M. The Distress Tolerance Scale: Development and validation of a self-report measure. *Motivation and Emotion* **29**, 83-102 (2005).
79. John, O. P. & Srivastava, S. The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research* **2**, 102-138 (1999).
80. Duckworth, A. L. & Quinn, P. D. Development and validation of the Short Grit Scale (GRIT-S). *Journal of personality assessment* **91**, 166-174 (2009).
81. De Castella, K. & Byrne, D. My intelligence may be more malleable than yours: the revised implicit theories of intelligence (self-theory) scale is a better predictor of achievement, motivation, and student disengagement. *European Journal of Psychology of Education* **30**, 245-267 (2015).
82. Campbell, A. V., Chung, J. M. H. & Inzlicht, M. Meaningfulness of effort: Deriving purpose from really trying. (2022).
83. Clay, G., Mlynski, C., Korb, F. M., Goschke, T. & Job, V. Rewarding cognitive effort increases the intrinsic value of mental labor. *Proceedings of the National Academy of Sciences* **119**, (2022).
84. Dora, J., van Hooff, M. L. M., Geurts, S. A. E., Kompier, M. A. J. & Bijleveld, E. The effect of opportunity costs on mental fatigue in labor/leisure trade-offs. *Journal of Experimental Psychology: General* **151**, 695 (2022).
85. Kurzban, R., Duckworth, A., Kable, J. W. & Myers, J. An opportunity cost model of subjective effort and task performance. *Behav Brain Sci* **36**, 661-679 (2013).
86. Ritz, H. & Shenhav, A. Humans reconfigure target and distractor processing to address distinct task demands. *bioRxiv* (2021).
87. Ritz, H., Leng, X. & Shenhav, A. Cognitive Control as a Multivariate Optimization Problem. *J Cogn Neurosci* **34**, 569-591 (2022).
88. Shenhav, A., Prater Fahey, M. & Grahek, I. Decomposing the motivation to exert mental effort. *Current Directions in Psychological Science* **30**, 307-314 (2021).
89. Westbrook, A. et al. Dopamine promotes cognitive effort by biasing the benefits versus costs of cognitive work. *Science* **367**, 1362-1366 (2020).
90. Westbrook, A., Frank, M. J. & Cools, R. A mosaic of cost-benefit control over cortico-striatal circuitry. *Trends Cogn Sci* S1364-6613(21)00115 (2021).

91. Westbrook, A. et al. Economic choice and heart rate fractal scaling indicate that cognitive effort is reduced by depression and boosted by sad mood. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* (2022).
92. Miyake, A. & Friedman, N. P. The Nature and Organization of Individual Differences in Executive Functions: Four General Conclusions. *Curr Dir Psychol Sci* **21**, 8-14 (2012).
93. Judd, N. & Klingberg, T. Training spatial cognition enhances mathematical learning in a randomized study of 17,000 children. *Nat Hum Behav* **5**, 1548-1554 (2021).
94. Scherer, R., Siddiq, F. & Sánchez Viveros, B. The cognitive benefits of learning computer programming: A meta-analysis of transfer effects. *Journal of Educational Psychology* **111**, 764-792 (2019).
95. Frömer*, R., Lin*, H., Wolf, C. K. D., Inzlicht, M. & Shenhav, A. Expectations of reward and efficacy guide cognitive control allocation. *Nature Communications* **12**, 1-11 (2021).
96. Lin, H., Ristic, J., Inzlicht, M. & Otto, A. R. The Average Reward Rate Modulates Behavioral and Neural Indices of Effortful Control Allocation. *Journal of Cognitive Neuroscience* 1-14 (2022).
97. Inzlicht, M. & Campbell, A. V. Effort feels meaningful. *J Exp Psychol Gen* (2022).
98. Dweck, C. S. & Yeager, D. S. Mindsets: A view from two eras. *Perspectives on Psychological science* **14**, 481-496 (2019).
99. Danielmeier, C., Eichele, T., Forstmann, B. U., Tittgemeyer, M. & Ullsperger, M. Posterior medial frontal cortex activity predicts post-error adaptations in task-related visual and motor areas. *Journal of Neuroscience* **31**, 1780-1789 (2011).
100. Ritz, H. & Shenhav, A. Parametric control of distractor-oriented attention. *Proceedings of the 41st Annual Meeting of the Cognitive Science Society* 967-973 (2019).
101. Kahneman, D., Tursky, B., Shapiro, D. & Crider, A. Pupillary, heart rate, and skin resistance changes during a mental task. *Journal of Experimental Psychology* **79**, 164-167 (1969).
102. Leys, C., Ley, C., Klein, O., Bernard, P. & Licata, L. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology* **49**, 764-766 (2013).
103. Schönbrodt, F. D. & Wagenmakers, E. J. Bayes factor design analysis: Planning for compelling evidence. *Psychonomic Bulletin & Review* **25**, 128-142 (2018).
104. Morey, R. D. & Rouder, J. N. BayesFactor: Computation of Bayes Factors for Common Designs. R package version 0.9.12-4.2. (2018).
105. R Core Team. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. (2019).
106. Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D. & Iverson, G. Bayesian t tests for accepting and rejecting the null hypothesis. *Psychon Bull Rev* **16**, 225-237 (2009).
107. Pearl, J. Lord's Paradox revisited – (Oh Lord! Kumbaya!). *Journal of Causal Inference* **4**, (2016).
108. Bürkner, P. C. brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software* **80**, 1-28 (2017).

Acknowledgements

This work was supported by a grant from the Natural Sciences and Engineering Research Council of Canada (NSERC RGPIN-2019-05280) awarded to M.I. The funders had no role in study design, and did play any role in data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

H.L., A.W., and M.I. contributed to the conception and design of the work. A.W. and M.I. provided critical oversight and feedback of the work. H.L. and F.F. programmed the experiment and collected the data. H.L. wrote the manuscript. A.W. and M.I. provided critical feedback.

Competing interests

The authors declare no competing interests.

Table 1. Design Table

Question	Hypothesis	Sampling plan	Analysis	Interpretation given to different outcomes
1. Do effort preferences on <i>rewarded trials</i> in the <i>training section</i> differ between the <i>effort</i> and <i>performance</i> conditions?	Effort preferences will be higher in the <i>effort</i> condition than <i>performance</i> condition.	We will stop sampling or recruiting participants when the Bayes factors for Hypotheses 6 and 8 (i.e., rows 6 and 8 of this table) exceed a threshold (10 or 0.1) or when a total of 750 participants has been recruited.	The outcome measure is effort preferences (% hard choices) on <i>rewarded trials</i> in the training section. We will fit a Bayesian linear regression model with condition as the main regressor and effort preferences on the pre-training inhibition task as the covariate.	The results are consistent with the hypothesis if Bayes factor ≥ 3 . The results are consistent with the null hypothesis if the Bayes factor ≤ 0.3 . Specifically, Bayes factors larger than 1, 3, or 10 (or less than 1, 0.3, 0.1) provide anecdotal, moderate, and strong evidence for the experimental hypothesis (or null hypothesis).
2. Do effort preferences on <i>rewarded trials</i> in the <i>training section</i> differ between the <i>effort</i> and <i>neutral</i> conditions?	Effort preferences will be higher in the <i>effort</i> condition than <i>neutral</i> condition.	Same as above.	Same as the row above.	Same as above.
3. Do effort preferences on <i>probe (unrewarded) trials</i> in the <i>training section</i> differ between the <i>effort</i> and <i>performance</i> conditions?	Effort preferences will be higher in the <i>effort</i> condition than <i>performance</i> condition.	Same as above.	The outcome measure is effort preferences (% hard choices) on <i>probe trials</i> in the training section. We will fit a Bayesian linear regression model with condition as the main regressor and effort preferences on the pre-training inhibition task as the covariate.	Same as above.
4. Do effort preferences on <i>probe (unrewarded) trials</i> in the <i>training section</i> differ between the <i>effort</i> and <i>neutral</i> conditions?	Effort preferences will be higher in the <i>effort</i> condition than <i>neutral</i> condition.	Same as above.	Same as the row above.	Same as above.
5. Do effort	Effort preferences	Same as above.	The outcome	Same as above.

preferences on the <i>inhibition task</i> in the <i>post-training section</i> differ between the <i>effort</i> and <i>performance</i> conditions?	will be higher in the <i>effort</i> condition than <i>performance</i> condition.		measure is effort preferences (% hard choices) on the inhibition task in the post-training section. We will fit a Bayesian linear regression model with condition as the main regressor and effort preferences on the pre-training inhibition task as the covariate.	
6. Do effort preferences on the <i>inhibition task</i> in the <i>post-training section</i> differ between the <i>effort</i> and <i>neutral</i> conditions?	Effort preferences will be higher in the <i>effort</i> condition than <i>neutral</i> condition.	Same as above.	Same as the row above.	Same as above.
7. Do effort preferences on the <i>updating task</i> in the <i>post-training section</i> differ between the <i>effort</i> and <i>performance</i> conditions?	Effort preferences will be higher in the <i>effort</i> condition than <i>performance</i> condition.	Same as above.	The outcome is effort preferences (% hard choices) on the updating task in the post-training section. We will fit a Bayesian linear regression model with condition as the main regressor and effort preferences on the pre-training updating task as the covariate.	Same as above.
8. Do effort preferences on the <i>updating task</i> in the <i>post-training section</i> differ between the <i>effort</i> and <i>neutral</i> conditions?	Effort preferences will be higher in the <i>effort</i> condition than <i>neutral</i> condition.	Same as above.	Same as the row above.	Same as above.

Supplementary information

In the pilot study, participants were randomly assigned to either the effort or performance condition (Supplementary Fig. 1). Each participant first completed a baseline block (40 trials) where no rewards were offered. On each trial, they chose to perform the easy or hard version of the dot-motion inhibition task (equivalent to Block 1 in Fig. 1a). They then completed a reward block (40 rewarded trials) where they saw the reward trial cue and feedback (Fig. 1b) on every trial. Those in the effort condition experienced only the reward-effort value function (Fig. 1c, left panel), whereas those in the performance condition experienced only the reward-performance value function (Fig. 1c, right panel).

In the main text (Pilot data section), we reported that effort preferences in the reward block were higher in the effort than performance condition, demonstrating the feasibility of the proposed paradigm. Below, we report the results from additional models. When we include effort preferences in the baseline as a covariate in our model, we found strong evidence for the effect of condition, $b = 14.64$, 95% HPD = [7.28, 22.33], $d = 0.73$, $BF = 144.02$ ($p < .001$). When we modeled effort preferences as a change score (reward block minus baseline block), the results were similar, $b = 15.45$, 95% HPD = [7.54, 23.72], $d = 0.70$, $BF = 107.86$ ($p < .001$), indicating the robustness of the condition manipulation. We also modeled effort preferences as a function of condition, block, and condition-block interaction. There was an interaction effect, $b = 15.39$, 95% HPD = [5.29, 25.19], $d = 0.38$, $BF = 10.14$ ($p = .005$), which was driven by an increased in effort preferences across blocks in the effort condition, $b = 15.64$, 95% HPD = [7.64, 23.21], $d = 0.75$, $BF = 225.90$ ($p < .001$), but no change in effort preferences in the performance condition, $b = 0.13$, 95% HPD = [-6.56, 7.57], $d = 0.01$, $BF = 0.19$ ($p = .971$). This null effect suggests the performance condition serves as a good control or baseline condition because, overall, the manipulation does not increase or decrease effort preferences.

Supplementary Figure 1 | Pilot results. Participants were randomly assigned to either the effort (left panel) or performance (right panel) condition. They first completed a baseline block (dot-motion inhibition task) where no rewards were offered before completing a reward block where rewards were delivered at the end of each trial. In the reward block, those in the effort condition experienced only the reward-effort value function, whereas those in the performance condition experienced only the reward-performance value function. Effort preference (proportion of choices whereby participants chose the hard task) in the reward block was higher for the effort than performance condition. Each dot is one participant. The three horizontal lines in each violin plot are the 0.25, 0.50, and 0.75 quantiles. Each gray line connects the baseline and reward blocks' effort preferences for one participant: Positive and negative slopes indicate that the condition manipulation increased and decreased effort preferences, respectively.

Supplementary Figure 2 | Bayesian posterior estimates for models predicting effort preference and controlling for pre-training effort preference and pre-training task accuracy. Positive condition (“cond”) estimates are hypothesis-consistent results (i.e., higher effort preference in the effort than the other condition). Model specification: effort preference ~ condition * (pre-training baseline + pre-training task accuracy).

Supplementary Figure 3 | Exploratory analyses and Bayesian posterior densities for the effect of condition on task reaction time. Positive Bayesian posterior estimates indicate longer reaction times in the effort than the other condition. Model specification: reaction time ~ condition + pre-training baseline reaction time + objective task difficulty.

Supplementary Figure 4 | Task accuracy for reaction time quartiles. Task accuracy for easy and hard tasks. Mean and 95% CI are shown for both the y-axis (task accuracy) and x-axis (reaction time quartile). The y-axis is truncated and begins at 0.85 because mean accuracies were very high across all quartiles.

Supplementary Figure 5 | Exploratory analyses and Bayesian posterior densities for the interaction effect between condition and pre-training effort preference. Positive Bayesian posterior estimates indicate that the effect of condition on effort preference is stronger for participants with higher pre-training effort preference.

Supplementary Figure 6 | Casual forests examining heterogeneous treatment effects.
 Predicted conditional average treatment effects for each hypothesis plotted against values of covariates while holding constant the other two covariates at their mean values.