1 Apparent preferences for cognitive effort fade when multiple forms
2 of effort and delay are interleaved in a foraging environment

3 Claudio Toro-Serey^{1,2}, Gary A. Kane^{1,2}, & Joseph T. McGuire^{1,2}

4 1. Department of Psychological and Brain Sciences, Boston University, Boston, USA

5 2. Center for Systems Neuroscience, Boston University, Boston, USA

6 Corresponding authors: Claudio Toro-Serey (ctoroserey@mclean.harvard.edu) & Joseph T. McGuire
7 (jtmcg@bu.edu)

Abstract

Cognitive and physical effort are typically regarded as costly, but demands for effort also seemingly boost the value of prospects under certain conditions. One contextual factor that might influence the perceived value of effort is the mix of different demand types a decision maker encounters in a given environment. Here, we embedded both cognitive and physical effort in a “prey selection” foraging task, which required decision makers not only to evaluate the magnitude and delay of a focal prospective reward but also to estimate the general opportunity cost of time. In two experiments, participants encountered prospective rewards that required equivalent intervals of cognitive effort, physical effort, or unfilled delay. Monetary offers varied per trial, and the two experiments differed in whether the type of effort or delay cost was the same on every trial (between-participant manipulation, $n=21$ per condition), or varied across trials (within-participant manipulation, $n=48$). When each participant faced only one type of cost, cognitive effort persistently produced the highest acceptance rate compared to trials with an equivalent period of either physical effort or unfilled delay. We theorized that if cognitive effort were intrinsically rewarding, we would observe the same pattern of preferences when participants foraged for varying cost types in addition to rewards. Contrary to this prediction, in the within-participant experiment, an initially higher acceptance rate for cognitive effort trials disappeared over time amid an overall decline in acceptance rates as participants gained experience with all three conditions. Our results indicate that cognitive demands may reduce the discounting effect of delays, but not because decision makers assign intrinsic value to cognitive effort. Rather, the results suggest that a cognitive effort requirement might influence contextual factors such as subjective delay duration estimates, which can be recalibrated if multiple forms of demand are interleaved.

Introduction

Evaluating whether to pursue a desirable outcome often entails assessing the time and effort it will require. In psychology and economics, demands for both time and effort have traditionally been identified as costs that detract from the net subjective value of rewards. Decision makers tend to prefer low-effort and immediate rewards over larger rewards that demand longer delays and higher effort (Ainslie, 1975; Frederick et al., 2002; Hull, 1943; Kool & Botvinick, 2018; Kool et al., 2010; Shenhav et al., 2017; Treadway et al., 2009; Walton et al., 2006; Westbrook & Braver, 2015; Westbrook et al., 2013). However, people also sometimes appear to favor effortful courses of action in everyday decisions. Experimental and theoretical work has chiefly explained this phenomenon by suggesting a reciprocal relationship between effort and outcomes, whereby exerting effort boosts the perceived value of ensuing outcomes (Hernandez Lallement et al., 2014; Kacelnik & Marsh, 2002; Mochon et al., 2012), and positive outcomes imbue value into the preceding effortful behavior (i.e. learned industriousness, Eisenberger, 1992). These proposals raise questions about what features can make an effortful task attractive in its own right, and how contextual factors can alter its perceived costs (Inzlicht et al., 2018). Understanding what governs the intrinsic costs and rewards of effort could improve our ability to motivate completion of effortful daily activities such as schoolwork or physical exercise.

Recent work has proposed that the subjective costs of delay and cognitive effort are modulated by the value of the alternative ways one’s time and cognitive resources could otherwise be used, implying that decision makers track *opportunity costs* (Fawcett et al., 2012; Kurzban et al., 2013; Otto & Daw, 2019). Preferences for or against effortful courses of action could therefore depend on the presence of other types of demands and opportunities in the same choice environment (Kool & Botvinick, 2018). For example, exerting high effort for a reward might be aversive, but the same level of work might become appealing if the only alternative is the boredom of an equivalent period of passive waiting (for instance, a person might choose to take a cumbersome way home rather than wait for a delayed train, even if the resulting time of arrival is the same). Effects of context on effort preferences have been illustrated by experiments on charitable giving, in which high-effort donation method (such as a run) attracted larger donations than a low-effort method (such as a picnic) when either was available alone (the “martyrdom effect”), but the high-effort event was disfavored when both options were offered as alternatives (Olivola & Shafir, 2013).

Standard models in neuroeconomics and behavioral economics formalize the effects of effort or delay on valuation in terms of discounting functions (Ainslie, 1975; Frederick et al., 2002; Green et al., 1994; Kable & Glimcher, 2007; Klein-Flügge et al., 2015; Kool et al., 2010; Westbrook et al., 2013). Discounting functions appear to have a similar form for delay and effort (Massar et al., 2015; Prevost et al., 2010; Seaman et al., 2018), although there is ongoing debate about how the shape of the discounting function might vary

across types of demands (Arulpragasam et al., 2018; Białaszek et al., 2017; Chong et al., 2017; Hartmann et al., 2013; Klein-Flugge et al., 2016; Kool & Botvinick, 2018; Prevost et al., 2010) and the degree to which discount rates correlate across cost domains (Seaman et al., 2018; Westbrook et al., 2013). A limitation of the discounting approach is that it does not directly probe how aspects of the decision context, such as opportunity costs, might influence cost evaluation, and constrains the understanding of behavioral demands to one in which such demands necessarily reduce the value of prospects.

Single choice foraging paradigms, in which decision makers choose to accept or reject individually presented prospects, provide a natural way to investigate contextual influences on cost evaluation (Constantino & Daw, 2015; Garrett & Daw, 2020; Krebs et al., 1977; Mobbs et al., 2018; Stephens & Krebs, 1986). Such paradigms require weighing a prospective reward against the perceived opportunity cost of time. Specifically, foragers must decide whether to spend their limited time by harvesting a currently available resource or by seeking out more profitable (or less costly) alternatives instead. This could involve, for example, deciding when to leave a patch of depleting richness (e.g. when to quit a job that provides diminishing opportunities for growth), or choosing whether to pursue or forgo offers presented serially (e.g. whether to accept your first job offer or wait for new potential ones). Here we focus on this latter type of foraging, traditionally called *prey selection* (Krebs et al., 1977), as it provides a means to examine how choices to accept or reject individual prospects are influenced by contextual knowledge about other potentially obtainable alternatives in the same environment.

Foraging has recently attracted high scientific interest due to its ecological validity, which is rooted in evolutionarily conserved choice behaviors (Hayden, 2018; Mobbs et al., 2018). The prey selection foraging paradigm allows for the estimation of optimal, reward maximizing choice patterns on the basis of the opportunity cost of time (Charnov, 1976; Krebs et al., 1977). The opportunity cost depends on the richness of the environment, which can be experimentally manipulated by balancing the amount of time it takes to harvest an offer (*handling time*) and to search for a new one (*travel time*). We can then interpret observed deviations from such optimal behavior, and examine how foraging behavior changes when the handling time is filled with cognitive or physical effort in comparison with an unfilled delay (thus disentangling effort from delay duration). In contrast to multi-alternative economic choice paradigms, which often treat choices as independent events, the sequential nature of foraging encourages individuals to consider both focal and global elements of the decision context. Extensive previous work on modeling foraging decisions provides computational tools to formalize relevant decision variables (Stephens & Krebs, 1986).

Across two behavioral experiments, we tested the overarching hypothesis that preferences would vary depending on whether individuals faced a single type or multiple types of behavioral demands (i.e. cognitive effort, physical effort, and delay). Experiment 1 used a between-subject design in which each participant foraged for rewards of varying magnitudes that imposed varying durations of a single form of demand.

Experiment 2 used a within-subject design in which individuals foraged for varying forms of demand in addition to varying reward magnitudes. First, we examined whether the discounting effects consistently observed in multi-alternative economic choice paradigms persisted when these options were presented serially, and whether the resulting foraging patterns depended on the demands imposed to obtain rewards. The existence of a well-defined optimal, reward-maximizing choice strategy helped us identify when demands boosted or detracted from the value of rewards (signified by tendencies to overharvest poor offers or underharvest advantageous offers, respectively). Second, we probed whether the perceived cost (or value) of demands was intrinsic or context-dependent. Varying the level of exposure to multiple types of demands across experiments allowed us to assess whether cost evaluation was dependent on the availability of alternative avenues of action. Further, including cognitive demands and physical demands let us examine whether such effects might apply differently to distinct types of effort.

Our results provided mixed support for our pre-registered hypotheses, which are detailed in <https://osf.io/2rsgm/registrations>. For Experiment 1, we hypothesized that 1) decision makers would integrate reward and timing information in order to approximate optimal foraging strategies, as reflected by acceptance rates; 2) acceptance rates would be higher for trials that demanded physical effort and cognitive effort than for unfilled delay or effortless engagement (based on results from pilot data and previous findings from single-alternative choice paradigms; e.g., Olivola & Shafir, 2013. See also Eisenberger, 1992; Inzlicht et al., 2018); and 3) acceptance rates for each type of demand would remain stable over time, and participants would rarely exhibit within-trial reversals (accepting a trial but then quitting without receiving the reward). Consistent with our predictions, we found that individuals could approximate optimality by favoring higher rewards and shorter delays. We also found that people in the Cognitive group accepted more trials than optimal, whereas Physical and Wait groups under-accepted, partially matching our predictions but raising the possibility that cognitive effort might add value to outcomes (or attenuate subjective temporal costs). Finally, acceptance rates (and differences in acceptance rates across demands) were stable within each cost condition, in line with expectations. These findings motivated the design of Experiment 2, in which effort and delay trials were intermixed so that rejecting a trial of one type could potentially lead to a trial of the other type. For Experiment 2, we hypothesized that 1) acceptance rates would once again approximate optimality, and would show effects of reward magnitude similar to those seen in Experiment 1; 2) in contrast to Experiment 1, participants would accept unfilled-delay trials at a higher rate than trials with either form of effort; and 3) that choices would be consistent and stable over time. As before, decision makers approached the optimal strategy for the foraging environment, resembling the pattern observed in Experiment 1. Contrary to our second and third hypotheses, we found that acceptance rates were initially higher for cognitive effort than for delay, but this apparent preference disappeared over time as decision makers foraged for types of demands in

addition to rewards.

Computational modeling results suggested that the normative predictions of traditional foraging models were insufficient to explain our data. Instead, the data supported a model in which demands imposed a bias on the estimated opportunity cost of time, which converged as individuals gained experience with interleaved types of demands. We hypothesize that this bias arises from modulations of subjective time. For instance, cognitive effort could have become appealing by virtue of compressing the perceived trial time, a perception that could be recalibrated with exposure to alternative demands. A subjective shortening of experienced durations during cognitive effort could be due to the ongoing recruitment of working memory (e.g. preventing individuals from estimating the elapsed time), or because the discrete events within the cognitive task subdivided the time interval. These results expand our understanding of the attractiveness of effort, suggesting potential reward-independent factors that could be leveraged to motivate effort engagement in diverse, everyday scenarios.

Experiment 1: Between-subject Manipulation of Subjective Costs During Foraging

Methods

Participants

The study was preregistered with the Open Science Framework (<https://osf.io/2rsgm/registrations>). The data and code for all analyses can be found in https://github.com/ctoroserey/Cost_studies. All experimental procedures were approved by the Boston University Institutional Review Board, and written consent was acquired for all participants. For Experiment 1, we recruited individuals until the planned number of 84 eligible participants was achieved (58 Female, median age = 21, range = 18 - 31; number excluded before reaching goal = 8). The sample size was determined by means of power analysis (ANOVA), using a significance level of 0.05, power of 0.8, an effect size of $f = 0.45$ (estimated from a pilot study), and three groups (one for each cost type). The resulting per-group sample was 20, which we increased to 21 in order to match three possible block orders. We added an extra group of 21 participants who experienced a minimally effortful condition in order to determine whether effort or mere task engagement was driving our results.

Participant data sets could be excluded on the basis of any of the following preregistered criteria: 1) Withdrawal: if the participant did not complete the study (1 participant); 2) Inattentiveness: a catch trial was placed at the end of each experimental block, asking participants to press a key within 3 seconds (time requirement based on pilot study response times). A participant who failed two or more of these checks was

excluded and replaced (5 participants). 3) Improbable choice behavior: The task was structured so that the largest reward amount should always be accepted. A participant who quit every trial in at least one block was assumed not to have followed or understood task instructions, or to have disengaged from the task altogether (0 participants). 4) Performance: Participants were forced to travel if they made 2 mistakes in a cognitive effort trial (see task procedures below), or if they gripped below threshold during physical trials. Any participant with more than 30% forced travels was excluded (1 participant).

Foraging Task

The experimental task was run using PsychoPy 2 (v1.85.1, Peirce et al., 2019) on a Macbook Pro laptop. Participants foraged for monetary rewards in an environment in which each trial required physical effort, cognitive effort, or passive waiting for a set period of time (the “handling time”) in order to obtain the reward (Fig. 1A). Their goal was to maximize their gains within a fixed total amount of playing time (six 7-minute blocks). On each trial, a monetary offer was displayed for 2 s, and participants had the opportunity to expend time and/or effort during the handling time in order to earn it. Upon completion of a trial, participants saw a 2 s feedback message displaying the reward obtained, which was followed by a travel time to the next offer. Alternatively, the participant could quit at any point during the handling time by pressing the spacebar, and immediately start traveling. Participants received their earnings in the form of a real monetary bonus that was added to a fixed \$12 participation amount (possible range of task-based earnings = \$0 - \$15.50, rounded up to the nearest quarter).

Participants were divided into four groups. Each participant in the first three groups was assigned to one of three types of costs (cognitive effort, physical grip-force effort, or delay), and a fourth group experienced a low-effort physical task that required gripping above a negligible force level. This last condition was added to test the possibility that any differences in acceptances between effort and delay were driven by task engagement alone. Each group was unaware that other cost conditions existed. Participants exerted physical effort by gripping a handheld dynamometer (Biopac Systems, United States) using their dominant hand. Gripping requirements in the physical effort condition were calibrated at 20% of maximum voluntary contraction (acquired at the beginning of the session). The low-effort condition was equivalent except that the grip-force threshold was set to 5%. Cognitive effort entailed switching among Stroop, dot motion judgment, and flanker tasks. Task switching has previously been shown to impose subjective costs (Kool et al., 2010). In the Stroop task, one of three color names was displayed on the screen (red, blue or green), with a font color that was either congruent (e.g. the word red presented in red) or incongruent (e.g. the word red presented in blue). Participants had to select the color of the font (i.e., they had to suppress their tendency to read

the word). For the motion judgment task, 100 solid white dots moved on the screen. A fraction of these dots moved coherently to the left or right, while the rest moved in random directions (coherence could be either 30% or 40%, uniformly sampled). Participants had to respond with the direction of the coherent movement. In the flanker task, rows of arrowheads pointed either to the left or the right (maximum of 3 rows, 3 to 13 arrowheads per row). Participants responded with the direction of the center arrowhead, which could point in the same or opposite direction of its neighbors. The tasks were configured so that responses always involved a left or right key press (e.g. for Stroop, two colored circles were presented at each side of the screen). During the handling time, cognitive tasks and their configurations were randomly sampled, and were presented for 1 s followed by a 1 s inter-stimulus interval. Participants were asked to respond within each task’s presentation time (i.e., there was a 1 s response deadline). Before the experiment, participants trained in each cognitive task until they correctly performed six consecutive trials of each type. While in the handling time, if participants failed to maintain an above-threshold grip during the physical effort task (following the 1 s grace period) or made two mistakes during the cognitive task, they were forced to travel and missed the reward.

There were three block types, in which handling times of 2, 10, or 14 seconds were paired with travel times of 14, 6, and 2 seconds, respectively. All combinations added up to 16 seconds, meaning that accepted trials took the same amount of time in all conditions, but rejected trials were shorter in environments with shorter travel times. Timing parameters were held constant within each 7-min block. Each block type was presented twice in pseudorandom order (total session length = 42 minutes). Reward amounts varied per trial (4, 8, or 20 cents with equal probability), with the constraint that every reward was presented twice every six trials. This prevented sequences from being dominated by a single amount during any window of time. Timing information was disclosed at the beginning of each block, and each trial’s reward amount was displayed during a 2 s offer window before the trial began. Participants received training prior to the experimental session and were told about all possible timing conditions and reward amounts in order to avoid any biases that could arise through experience-dependent learning (Dundon et al., 2020; Garrett & Daw, 2020). Finally, we encouraged participants to carefully evaluate their options by informing them that accepting all offers would not maximize their rewards.

Figure 1: Task design and optimal behavior. A: General foraging trial structure. On each trial, participants were offered an opportunity to earn money by sustaining effort or waiting during the handling time (2, 10, or 14 s). The end of a trial was followed by a travel time (handling and travel times always added up to 16 s). Participants could skip unfavorable trials and immediately start traveling to a potentially better offer. In Experiment 1, the type of demand was fixed per participant and handling time varied per block. In Experiment 2, handling time was fixed at 10 s, but a combination of effort and delay trials changed per block. B: Possible earnings per second for each acceptance threshold (i.e. the smallest amount accepted) for each handling time in Experiment 1. Circles denote the reward-maximizing threshold for each block, which is described in the table. Experiment 2 presented only the 10 s handling time, for which it was optimal to skip all 4-cent offers.

Foraging theory posits that the decision to accept a delayed reward should depend on the opportunity cost of time to be incurred in obtaining it, which in turn depends on the richness of the environment (Charnov, 1976; Stephens & Krebs, 1986). In this study, the richness of the environment was manipulated by the length of the handling and travel times. Shorter travel times produced richer environments by making it possible to save more time by skipping low-reward trials. Since time combinations were fixed per block, we calculated each block’s optimal accept/reject strategy by computing the expected rate of gain (g^*) from all decision strategies according to the following equation:

$$g^* = \frac{\sum p_i R_i}{3T_{\text{offer}} + \sum p_i (T_{\text{handling}} + T_{\text{reward}}) + 3T_{\text{travel}}}, p \in \{0, 1\}$$

where R and p are the reward magnitude and acceptance probability of offer i , respectively, T_{handling} is the handling time for the block, T_{travel} is the travel time, T_{offer} is the duration of the offer window (2 s), and T_{reward} is the duration of the window displaying the reward earned (2 s, only presented for completed trials). This gives the average reward per second attainable in each block as a function of the acceptance threshold (i.e. the lowest amount the participant is willing to accept in a given timing environment). Since an ideal forager should always either accept or reject any given category of prey (Krebs et al., 1977), probabilities other than 0 and 1 were not considered. The resulting possible strategies were to accept only 20 cents, accept 8 and 20 cents, or accept everything. Fig. 1B shows possible earnings per second for each choice strategy, as well as the lowest amount participants should accept in order to maximize their rewards (circled dots). For example, a participant in a 10 s handling-time block should accept 8-cent and 20-cent rewards (and reject 4-cent rewards) to maximize their reward rate; a participant in a 14 s handling-time block should accept only 20-cent rewards, and a participant in a 2 s handling-time block should accept everything. Note that accepting every offer was often detrimental to earnings.

Analyses

All analyses were performed in R 3.5.2 (R Core Computing Team, 2017). First we tested whether decision makers integrated delay and reward information. To address the prediction that participants would be more likely to accept higher reward and shorter handling time trials, we fit a mixed-effects logistic regression to predict trial-wise acceptances (using the lme4 package, Bates et al., 2015), giving a random intercept to each subject to account for biases in participant behavior. We included continuous regressors for handling time and reward, as well as two covariates that probed the influence of recent history on choices. The first was a term containing the number of consecutive unobtained offers prior to a given trial (unobtained offers could be

due to quits or forced travels). The second regressor tracked the sum of the reward amounts on the previous n offers. We identified the n that minimized Akaike’s Information Criterion (AIC) among 6 versions of the model (ranging from 2 previous trials up to 7, beyond which the model failed to converge), and report all coefficients from the winning model, which used the sum of the previous 6 offers.

Our foraging task was configured such that over- and under-accepting were detrimental to total earnings. To confirm this, we fit a general linear model with constant, linear, and quadratic terms to estimate the correspondence between proportion of trials completed (independent variable) and total earnings (dependent variable). A significant quadratic coefficient thus would signal that the task statistics operated generally as expected. Next, to determine the optimality of the group’s decisions, we examined whether each cost type produced a bias to over or under-accept offers of 4 and 8 cents (assuming that 20cent-offers would always be accepted, given our design). The reward-maximizing strategy was to always reject these offers for handling times of 14 s, reject 4 cents and accept 8 cents for 10 s, and accept both for 2 s, yielding a combined optimal overall acceptance rate of 50%. We performed a two-sided one-sample chi-squared test of proportions against the null probability of 0.5 for each type of cost (with trials pooled across participants in each group). Therefore, a significant difference indicated that participants either over- or under-accepted rewards.

The analyses above differed slightly from the pre-registered plan. Rather than pooling coefficients from individual logistic regressions, we favored obtaining a single coefficient from a mixed model that accounted for participant-level random effects. The chi-squared test was adopted from the Experiment 2 preregistration, as we found it better suited to the question than the original formulation (i.e. individual tests for each combination of experimental parameters).

Next, we compared acceptance rates across cost conditions. We first performed a one-way ANOVA on the proportion of trials accepted using group as a factor. We then compared the proportion completed across all 4 groups using non-parametric permutation contrasts (6 tests). For each test, on each of 5000 permutation iterations, participants’ group labels were randomly shuffled without replacement, and the difference in mean acceptance rates across iterations created an empirical null distribution. The unpermuted group mean difference was then evaluated against this permuted distribution. The same approach was used to test differences in total earnings.

In order to further look at the effect of handling time, offer amount, and cost condition, we computed the probability of accepting a trial with a mixed-effects logistic regression. Our a priori model of interest included all three variables as fixed main effects, and a random intercept per subject. Cost condition was modeled with three categorical terms, with the fourth condition as the reference condition. We ran three versions of the model with different reference conditions, in order to test all pairwise differences among the four cost conditions. We then examined whether this a priori model outperformed both simpler and more

complex models. We used both AIC and Bayesian Information Criterion (BIC) to determine the model that minimized the negative log-likelihood while penalizing the addition of parameters. The regressors included in the 8 candidate models were: 1) intercept only; 2) cost condition only; 3) handling time only; 4) offer amount only; 5) handling time and offer amount; 6) cost condition, handling time, and offer amount main effects (a priori model from above); 7) adding a handling-time-by-offer-amount interaction; and 8) all possible two-way interactions. We predicted that model 6 would have the lowest AIC and BIC. Nested models with similar AIC were statistically compared using an analysis-of-deviance. The significance test was computed as the probability of the reduction in deviance, based on a chi-square distribution with degrees of freedom equal to the difference in the number of parameters between models.

Reward-maximizing prey foragers should either accept or reject offers based on their profitability relative to the available alternatives (Krebs et al., 1977; Stephens & Krebs, 1986). However, decision makers often curtail persistence before reaching a delayed prospect based on a continuous reevaluation of the prospect's value or a lapse in self control (Ainslie, 1975; McGuire & Kable, 2012, 2015; Mischel & Ebbesen, 1970). Therefore, we examined whether participants were consistent (not quitting trials that were originally accepted) and stable (maintaining similar acceptance rates over time) in their choices. To evaluate consistency, we visually examined survival curves indicating at what point during the handling time participants quit a trial; large numbers of late quit times would indicate a lack of consistency. We interrogated stability by computing each participant's total proportion of acceptances in the first and second half of the experiment for every block type, comparing the mean proportion of acceptances using paired permutations (5000 iterations). We then tested whether the observed differences in costs were still present in each half of the experiment by applying the winning model from the mixed-effects logistic regression analysis to each half separately.

Results

Decision Makers Integrate Delay and Reward Information

Four groups of participants faced with different behavioral demands (physical effort, cognitive effort, low-effort task engagement, or passive delay) chose their preferred strategy to maximize rewards in foraging environments of varying richness. We will refer to the four groups below as Physical, Cognitive, Easy, and Wait. Environmental richness was dictated by the time it took to obtain a reward (the handling time) and the time between trials (the travel time). We hypothesized 1) that participants would approximate reward-maximizing behavior by preferentially accepting higher rewards and shorter delays; 2) that participants confronted with effortful demands would be more likely to accept trials than those faced with passive delay, regardless of handling time and reward amount; and 3) that choices would be stable over time and consistent

within trials.

We performed a mixed-effects logistic regression to address the first hypothesis, pooling across the four conditions. Since participants in the effort groups were forced to travel when they failed to perform above threshold (see Methods), we distinguish between acceptances (not explicitly quitting a trial during the handling period) and completions (successfully obtaining a reward). Even so, the proportion of forced travels was low for the Cognitive group (5% on average across participants) and 0 for the Physical group (i.e. no forced travels in this group). Overall, participants accepted an average of 64% of trials (SD = 17%, range = 35% to 100%). In line with our predictions, larger reward amounts significantly increased acceptance probabilities ($\beta = 0.49$, SE = 0.01, $p < 0.001$), whereas longer delays decreased them ($\beta = -0.34$, SE = 0.01, $p < 0.001$). Two additional regressors showed that having missed out on consecutive rewards decreased the probability of acceptance ($\beta = -0.06$, SE = 0.03, $p < 0.001$), and that participants became more selective when recent offer history was richer ($\beta = -0.01$, SE = 0, $p < 0.001$). The counterintuitive reduction in acceptances following consecutive unobtained rewards was driven by the Wait group (removing this group from the model yielded a $p > 0.1$).

Figure 2: Overall choice behavior in Experiment 1. A: Proportion accepted per cost. B: Total number of dollars earned by each group by the end of the experiment (not including the participation reward). C: The relationship between proportion accepted and total earned. Consistent with the foraging design, participants who over and under accepted earned the least. These results suggest that people were willing to pursue cognitively effortful actions at the expense of earnings.

The next set of analyses examined the hypothesis that both Cognitive and Physical groups would uniformly accept more trials than the Wait and Easy groups. Fig. 2A shows that the Cognitive group consistently

accepted more offers, which resulted in lower earnings (middle). One-way ANOVAs showed that differences among groups were significant both for overall proportion accepted ($F(3, 80) = 6.94, p < 0.001$), and total earnings ($F(3, 80) = 8.07, p < 0.001$). Post-hoc permutations (5000 iterations) comparing mean acceptance rates between every pair of costs confirmed that acceptance rates of the Cognitive group (mean = 0.74, SD = 0.13) were higher than those of the Physical (mean = 0.6, SD = 0.15; $p = 0.004$, Cohen's D = 1.02) and Wait groups (mean = 0.55, SD = 0.11; $p < 0.001$, Cohen's D = 1.55). It also showed that those faced with the easy task accepted more than those in the Wait group, although responses from the Easy group were more variable (mean Easy = 0.68, SD = 0.19; $p = 0.012$, Cohen's D = 0.85) (all other $p > 0.05$). In line with its higher acceptance rates, the Cognitive group had significantly lower earnings (mean = 13.55, SD = 0.84) than either the Physical group (mean = 14.03, SD = 0.53; $p = 0.036$, Cohen's D = 0.69) or the Wait group (mean = 14.5, SD = 0.58; $p < 0.001$, Cohen's D = 1.32). In addition, the Wait group earned more than the Physical ($p = 0.008$, Cohen's D = 0.85) and Easy groups (mean Easy = 13.94, SD = 0.52; $p = 0.002$, Cohen's D = 1.02).

We next examined the optimality of these decisions. The foraging task was configured so that there was a single reward-maximizing strategy per block type, and participants were informed that accepting everything would not maximize their gains. Accordingly, participants who over- or under-accepted earned the least (Fig. 2C; general linear model with quadratic term, $F = 25.65, \beta = -8.99, SE = 2.42, R^2 = 0.39, p < 0.001$). We performed group-level chi-squared tests of proportions (with trials pooled across participants) to examine whether each cost biased decision makers to under- or over-accept 4 and 8 cent offers, which collectively should have been accepted 50% of the time across the three timing environments. We found a significant over-acceptance of these trials by the Cognitive group (proportion accepted = 0.59; $\chi^2 = 73.89, p < 0.001$), but under-acceptances by the Physical group (proportion accepted = 0.39; $\chi^2 = 123.81, p < 0.001$) and Wait group (proportion accepted = 0.31; $\chi^2 = 359.13, p < 0.001$), with no significant difference for the Easy group. Together, the findings discussed so far indicate that although participants reacted appropriately to reward and timing contingencies, departures from optimality were systematically influenced by what form of cost was at stake. The findings raise the possibility that cognitive effort can boost the value of an offer, as reflected by higher acceptances at the expense of earnings.

Figure 3: Comparison of acceptance rates across cost types for each combination of reward amount and handling time. A: Mean acceptance rate in each group for every combination of handling time and reward (with SEM). Grey points indicate the reward-maximizing acceptance rate for each combination, which was always either 0 or 1. B: Matrix of coefficients that resulted from switching the reference cost condition in a mixed-effects logistic regression model. Each element represents how much more likely a group was to accept offers compared to the reference condition (row). White asterisks denote the p-value for significant comparisons (one, two, and three asterisks correspond to p-values under 0.05, 0.01, and 0.001, respectively). C: AIC and BIC values for a mixed-effects model comparison, ranging from a baseline intercept-only model to a full-interaction model. Both metrics yielded similar values across models.

We next probed whether these results held uniformly across reward and timing parameter combinations. Fig. 3A shows the mean proportion of acceptances ($\pm$ SEM) per combination of handling time, reward, and cost type, with reward-maximizing acceptance rates depicted by the gray circles. Qualitatively, the figure confirms two important predictions. First, participants adapted to the richness of each timing block, gravitating towards optimality regardless of the cost they faced. Second, effects of cost type were consistent across the three timing conditions, and acceptances were uniformly highest for the Cognitive group. We tested these observations with a mixed-effects logistic regression model, which estimated the probability of accepting an offer as a function of handling time, reward amount, and cost type (see Methods for details). The model showed significant main effects of handling time ($\beta = -0.34$, $SE = 0.01$, $p < 0.001$) and offer amount ($\beta = 0.49$, $SE = 0.01$, $p < 0.001$). Comparisons among all conditions are shown on Fig. 3B. The Cognitive group was significantly more likely to accept offers than the Physical group ($\beta = 1.5$, $p = 0.005$) or the Wait group ($\beta = 2.11$, $p < 0.001$), but not the easy group.

These observations were not better explained by simpler or more complex regression models. Fig. 3C shows AIC and BIC values for models of increasing complexity, starting with an intercept-only configuration (see Methods for details). Here we focus on AIC, as both metrics yielded comparable results. As predicted, the a priori model with all main effects performed better than those relying on a single parameter (a priori model pseudo $R^2 = 0.82$), as well as the model with only main effects for handling time and reward amount ($AIC_{\text{a-priori}} = 3416$, $AIC_{\text{handling/reward}} = 3427$, $\chi^2(3) = 16.87$, $p < 0.001$).

Contrary to our expectations, the a priori model was outperformed by models that added an interaction between handling and reward ($AIC_{\text{HR_interaction}} = 3394$, $\chi^2(1) = 24.16$, $p < 0.001$) as well as one that considered all possible interactions ($AIC_{\text{all_interactions}} = 3300$, $\chi^2(7) = 129.58$, $p < 0.001$). However, upon examination, neither of these models offered further interpretable insights on the observed differences across costs. The overall pattern of results supported the hypotheses that handling and reward amounts were integrated in the expected fashion, and that the rate of acceptances was affected by the demand faced. The results offered only partial support for the hypothesis that effort would induce higher acceptance rates; higher acceptance rates were observed for cognitive effort but not for physical effort.

Consistency and Stability of Choices

Next, we tested the hypothesis that foragers would exhibit within-trial consistency and across-trial stability in their choices. Participants were deemed consistent if they tended not to quit after engaging in a trial (that is, if they pressed the spacebar to reject a trial either quickly or not at all). To examine this, we plotted survival curves of trial rejection times across all participants (Fig. 4A), with black crosses signaling censored

387 observations either for completed trials (at 2 s, 10 s, and 14 s) or forced travels. These curves show that
 388 most decisions to quit happened within the first second into the handling time, and participants rarely quit
 389 afterwards (only 4% of quits occurred later than 1 s). The pattern of results in the Physical and Easy groups
 390 reflected that participants sometimes quit by letting the 1 s initial grace period expire without gripping
 391 (rather than by pressing the spacebar); such trials were coded as quits rather than as forced travels. Plotting
 392 the quitting time distributions across blocks for the first 1 s of the handling time (ECDFs on Fig. 4A) showed
 393 that participants in the Wait and Cognitive groups responded faster as the experiment progressed. These
 394 response time observations suggested clarity in participant choices.

Figure 4: Choice consistency and stability in Experiment 1. A: Survival curve plots of trial rejection times.
 B: Distributions of trial rejection times within the first 1 s of the handling time. C: Acceptance rate in the
 second half of the session as a function of acceptance rate in the first half of the session. Dashed lines indicate
 the mean for each group. D: proportion accepted over time per group.

Stability was defined as the maintenance of similar acceptance rates throughout the experimental session. The scatterplot in Fig. 4C shows the overall proportion accepted by each participant before and after the midpoint of the session. Paired permutations showed no change in acceptance rates for the Wait and Easy groups ($p > 0.05$), but we found significantly lower acceptance rates in the second half of the session for the Physical group ($p = 0.016$, Cohen's $D = 0.27$) and the Cognitive group ($p < 0.001$, Cohen's $D = 0.56$). Fig. 5A shows that the reduction in acceptance rates in effort groups over time was steeper for longer handling times, a hallmark of effortful demands (Treadway et al., 2009). We then replicated the mixed-effects logistic regression from the previous section separately for each half of the session, which confirmed that differences among demands were mostly stable (Fig. 5B; a significant difference between Physical and Easy groups during the second half of the session was due to the decrease in acceptance rates by the Physical group). Together, the consistency and stability of behavior suggested that decision makers had a clear representation of the environment.

Figure 5: Acceptance behavior over time in Experiment 1. A: Offer acceptance rates per handling time, reward amount, cost type, and session half, showing that effortful demands (Cognitive and Physical) were less frequently accepted over time, especially for the longest handling time (14 s). B: Mixed-effects logistic regression model coefficients denoting the comparison of acceptance rates across cost conditions for each half of the experimental session separately (mirrored along the diagonal). Even though acceptance rates decreased in the effort groups, the differences in acceptance rates across cost conditions were mostly preserved.

Experiment 2: Within-subject Comparison of Costs

Experiment 1 demonstrated that human decision makers could appropriately integrate reward and timing information to make single-alternative accept/reject decisions in a foraging environment. In addition, acceptance rates were higher in a group of participants for whom rewards were associated with demands for cognitive effort than in groups for whom rewards were associated with physical effort or mere delay. One possible interpretation of the findings of Experiment 1 might be that cognitive effort was simply intrinsically rewarding. For instance, participants might have enjoyed performing the cognitive tasks and preferred them to a boring, unfilled delay. We reasoned that if this were the case, participants ought to continue preferentially accepting cognitive effort trials in an environment in which effort and non-effort trials were randomly intermixed within-subject, where rejecting a delay trial would potentially lead to an opportunity to accept a subsequent effort trial. On the other hand, if cognitive effort were not preferred, interleaving different types of demand in the same environment might lead to a different pattern of choices. Experiment 2 tested which types of trials decision makers would preferentially accept when they were allowed to forage for different types of demand in addition to different sizes of rewards. The objective was to address how the presence of multiple alternative forms of demand in the same choice environment would impact the subjectively assessed costs of effort and delay.

Methods

Participants

We collected data from a new sample of 48 participants (39 Female, median age = 21, range = 18 - 36; number excluded before reaching goal = 6). Sample size was once again determined by means of a power analysis (repeated measures ANOVA), using a significance level of 0.05, power of 0.8, an anticipated effect size of $f = 0.5$, and four factors (one for each demand-type condition; see below for details). The resulting sample size was 45, which we increased to 48 in order to balance the potential order of blocks. We excluded participants using the same criteria as in Experiment 1 (4 participants were excluded because they failed the inter-block attention checks).

Within-participant Manipulation of Cost Type

The original between-subjects task from Experiment 1 was modified so that participants foraged for cost types in addition to rewards. Every participant experienced all forms of cost (delay, physical effort, and cognitive effort). There were two block types, in which delay trials were interspersed with either physical-

effort or cognitive-effort trials. This setup prevented participants from having to rapidly switch between response modalities across trials (i.e. keyboard and handgrip). Each block combination (Physical/Wait and Cognitive/Wait) was experienced three times by each participant in 7-minute-long, interleaved blocks, for a total session duration of 42 minutes. Half of the participants experienced a block sequence that started with the Physical/Wait block and the other half experienced a sequence that started with the Cognitive/Wait block. Timing parameters were held fixed throughout the experiment and were matched to the middle condition from Experiment 1 (10 s handling time and 6 s travel time). Participants were informed that this timing combination would be constant prior to beginning the experiment. The upcoming effort type was disclosed at the beginning of each block, and reward/cost offers were displayed at the start of each trial (e.g. “8 cents for physical effort”). Unlike the previous experiment, participants could express their decision to quit by pressing the spacebar during the offer window (not only during the handling time), in which case the travel time would begin immediately at the end of the 2 s offer window. Participants were trained in each type of demand until they reached the same criteria as in Experiment 1.

Analyses

The analytical pipeline was mostly preserved from Experiment 1. We first tested whether decision makers integrated reward information. We performed a mixed-effects logistic regression with regressors for offer amount, number of consecutive offer misses, and the sum of the n previously observed offers (we identified n by comparing the AIC of different history lengths). We then confirmed that under- and over-acceptance were detrimental to total earnings by fitting the same quadratic linear model discussed above. The optimal strategy was to reject 4-cent offers and accept 8-cent offers, yielding an ideal overall acceptance rate of 50% (not counting 20-cent offers, which should have always been accepted). For these trials, we performed a two-sided one-sample chi-squared test of proportions against the null probability of 0.5 for each cost condition.

To compare acceptance rates among cost conditions, we first estimated acceptance probabilities using a mixed-effects logistic regression. The model included cost condition and reward amount as fixed main effects, and participant as a random intercept (handling time was not modeled, as participants experienced a single time combination). Cost condition was modeled with three categorical terms, with the fourth condition as the reference condition (two effort conditions, and a delay condition paired with each effort type). We ran three versions of the model with different reference conditions, in order to test all relevant pairwise differences among the four cost conditions. We compared the a priori model to alternative parameterizations using AIC and BIC, varying model complexity as follows: 1) intercept only; 2) cost condition only; 3) reward only; 4) cost condition and reward main effects (a priori model); and 5) adding a two-way interaction of cost

condition and reward. We predicted that model 4 would have the last appreciable decrease in the negative log-likelihood. Models with similar AIC were formally compared using the analysis-of-deviance approach from Experiment 1.

Next, we probed within-trial consistency and across-trial stability of choice behavior. First, we visualized the number of quit responses during the handling time with a survival curve, and plotted the quit response time distribution as a function of block to examine choice consistency over time. Next, we computed the proportion of quits that occurred during the choice window versus during the handling time. Earlier quits were seen as signifying greater consistency. We then computed each participant’s total acceptance rate in the first two and last two blocks, and compared them using a paired permutation analysis (5000 iterations). This analysis did not include the third and fourth block because the alternating, counterbalanced sequence of block types would have led to unequal numbers in the first three and last three blocks for individual participants.

Results

Decision Makers Integrate Reward Information

In Experiment 2, we tested decision makers’ preferences for cognitive effort or physical effort relative to unfilled delay delay. We modified the prey-selection foraging task, so each participant experienced every type of demand in blocks that interleaved wait trials with either physical or cognitive effort (with a fixed 10 s handling time and 6 s travel time, to allow for comparisons between experiments). We separately analyzed Wait trials from the two block types, labeling the two sets of Wait trials according to the type of effort with which they were paired (“Wait-C” for cognitive, and “Wait-P” for physical). We hypothesized that 1) reward offer amounts would influence decisions in the same way as in Experiment 1; 2) interleaving effort and delay would produce higher acceptance rates for wait trials than for cognitive or physical effort trials, in contrast to the pattern observed in Experiment 1. and 3) that choice patterns would remain stable over time and consistent within trials.

We first tested the hypothesis that participants would successfully integrate reward and timing information. A mixed-effects logistic regression showed that larger reward amounts significantly increased acceptance probabilities ($\beta = 0.762$, $SE = 0.024$, $p < 0.001$), and that participants became more selective if they had recently observed large offers ($\beta = -0.015$, $SE = 0.002$, $p < 0.001$). Model comparisons suggested that participants were influenced by the 7 most recent rewards. Unlike Experiment 1, participants were unaffected by the preceding number of consecutive unobtained rewards ($p = 0.434$). As before, a linear model predicting earnings showed that participants who completed too few or too many trials earned the least, in line with the experimental design (Fig. 6A; general linear model with quadratic term, $F = 11.21$, $\beta = -5.53$, $SE = 2.71$, R^2

502 = 0.33, $p = 0.047$). When assessing the optimality of acceptances for 4 and 8 cents (99% of 20 cent offers were
 503 correctly accepted), chi-squared tests against the optimal rate of 50% showed that participants under-accepted
 504 physical effort trials (proportion accepted = 0.46, $\chi^2 = 6.53$, $p = 0.01$; all other $p > 0.05$). Despite this
 505 deviation, participants earned a similar total amount per cost condition on average (mean Cognitive = 3.13,
 506 SD = 0.52; mean Wait-C = 3.56, SD = 0.27; mean Physical = 3.44, SD = 0.25; mean Wait-P = 3.59, SD =
 507 0.24). These results convey that, in contrast to Experiment 1, participants were similarly successful across
 508 the different cost conditions.

Figure 6: Within-subject acceptance behavior. A: Quadratic relationship between trial completions and total earnings. B: Proportion of acceptances of each reward per cost condition. C: Matrix of mixed-effects coefficients showing comparisons among cost conditions (while controlling for reward amount). Cognitive and Physical trials were compared to one another and were each compared to Wait trials from the same block (Wait-C and Wait-P, respectively). We omitted comparisons to Wait trials in the opposite block type (gray cells), which have little interpretability and were not addressed in our predictions. Asterisks mark statistically significant comparisons. D: AIC and BIC values for models of increasing complexity (AllMain refers to the a priori model which included main effects of cost condition and reward offer amount.)

Next, we addressed the hypothesis that passive waiting trials would be preferentially accepted relative to effortful trials. Even though the overall reward acceptance pattern was similar to the 10 s handling time condition in Experiment 1 (Fig. 6B), differences among cost conditions were less evident. Moreover, these differences contradicted our predictions, as cognitive effort continued to be accepted at the highest rate. A mixed effects logistic regression confirmed once again that increases in offer amount made acceptances significantly more likely ($\beta = 0.76$, $SE = 0.02$, $p < 0.001$), but that the only significant difference across cost conditions was a higher acceptance rate for cognitive-effort offers than physical-effort offers (Fig. 6C; $\beta = 0.38$, $p < 0.001$). Neither Cognitive trials nor Physical trials had significantly different acceptance rates from the Wait trials that were interleaved in the same block (Cognitive versus Wait-C: $\beta = 0.06$, $p = 0.565$; Physical versus Wait-P: $\beta = -0.18$, $p = 0.11$). Model comparisons showed that a model with all main effects explained a large portion of the variance ($R^2 = 0.82$), and there were similar AIC and BIC values for any model including reward amount (Fig. 6D). An analysis of deviance among the three models showed that adding cost condition to reward amount significantly improved the fit ($AIC_{\text{reward}} = 1859$, $AIC_{\text{a-priori}} = 1852$, $\chi^2(3) = 13.24$, $p = 0.004$) and that including a term for the interaction between reward amount and cost condition improved the fit still further ($AIC_{\text{Interaction}} = 1830$, $\chi^2(3) = 28.34$, $p < 0.001$). These results contradicted our hypothesis that Wait trials would be favored, but also did not indicate participants had a straightforward preference for Cognitive trials relative to the corresponding Wait-C trials. Instead, cost condition had only relatively small effects on acceptance rates when cost conditions varied within-participant. To shed further light on the dynamics underlying the small overall effect, we next examined how it varied over the course of the experimental session.

Consistency and Stability of Choices

Consistent with Experiment 1, a survival analysis showed that participants made most of their choices quickly within the offer window (Fig. 7A). The median percentage of quits during the handling time was 2.61% ($SD = 7.67$), with only 2 participants quitting over 20% of trials during the handling time. Cumulative quitting time distributions showed that responses during the offer window became faster over time (Fig. 7B). Once in the handling time, the overall percentage of forced travels was low for both Cognitive trials (median = 5.99%, $SD = 9.59$) and Physical trials (median = 0%, $SD = 3.62$). The results suggest that decision makers were consistent in their adopted strategies, and were able to perform the tasks well.

Figure 7: Choice consistency and stability in Experiment 2. A: Survival curves showing trial rejection times for each cost condition. B: Distribution of trial rejection times within the offer window. C: Acceptance rate in the last two blocks as a function of acceptance rate in the first two blocks. Dashed lines indicate the mean for each cost. D: Acceptance rates in the three successive occurrences of each block type.

However, against our predictions, participants accepted significantly fewer trials in the second half of the experiment, regardless of cost condition (paired permutations; Cognitive: $\text{mean}_{\text{pre}} = 0.76$, $\text{SD}_{\text{pre}} = 0.17$, $\text{mean}_{\text{post}} = 0.65$, $\text{SD}_{\text{post}} = 0.21$, $p < 0.001$, Cohen's $D = 0.6$; Physical: $\text{mean}_{\text{pre}} = 0.71$, $\text{SD}_{\text{pre}} = 0.2$, $\text{mean}_{\text{post}} = 0.65$, $\text{SD}_{\text{post}} = 0.21$, $p = 0.016$, Cohen's $D = 0.27$; Wait-C: $\text{mean}_{\text{pre}} = 0.73$, $\text{SD}_{\text{pre}} = 0.2$, $\text{mean}_{\text{post}} = 0.66$, $\text{SD}_{\text{post}} = 0.22$, $p = 0.009$, Cohen's $D = 0.31$; Wait-P: $\text{mean}_{\text{pre}} = 0.72$, $\text{SD}_{\text{pre}} = 0.2$, $\text{mean}_{\text{post}} = 0.66$, $\text{SD}_{\text{post}} = 0.22$, $p = 0.008$, Cohen's $D = 0.3$; Fig. 8). The reduction in acceptances was not driven by a performance decline, as the proportion of forced travels was not significantly different between the first and second half of the experiment (paired permutation with 5000 iterations, $p = 0.161$). Notably, the pattern of choices

from the first part of the experiment resembled that from the corresponding 10 s handling-time condition in Experiment 1 (compare Fig. 5A, middle to Fig. 8A). Mixed effects logistic regressions performed on each half separately confirmed that the relative preference for cognitive effort seen previously was preserved at the beginning of the experiment (Fig. 8B, lower triangle), but that it became undetectable over time amid an overall decline in acceptance rates as participants gained experience with cognitive effort, physical effort, and passive waiting trials (upper triangle). This fading preference for cognitive effort was observed regardless of which block participants experienced first, although its magnitude appeared to be reduced for those who experienced the physical block first (Fig. 8, bottom)

Together, our findings show that the stable preferential acceptance of cognitive effort trials that we observed during single-demand foraging slowly faded if alternative demands were introduced. This finding raises the question of how features of the cognitive task might be re-evaluated as a function of other experiences in the same decision context. We probed this question by fitting a series of computational models to data from both Experiments 1 and 2.

Figure 8: Acceptance rates over time in Experiment 2. A: Proportion accepted in the first two and last two blocks of the experiment. An initial preference for cognitive effort trials faded amid a global decrease in acceptance rates. B: Mixed effects coefficient matrix for early and late trials separately. The pattern for early trials – but not for late trials – resembled what was observed in Experiment 1 (Figure 3B). C: Results presented separately for participants who experienced the Cognitive/Wait block type first or who experienced the Physical/Wait block type first.

Computational Modeling of Foraging Behavior

We developed a computational model to identify cognitive operations that could explain the complex patterns observed across Experiments 1 and 2. We identified four key aspects of the behavioral results that an adequate computational model should be able to capture: 1) modest under- and over-harvesting of offers relative to the optimal strategy as a function of reward offer amount (e.g. Fig. 3A); 2) dependence of choices on recent offer history; 3) higher acceptance rates if all reward offers were associated with cognitive effort (Experiment 1); and 4) gradual abolishment of the tendency to over-accept cognitive effort trials if other forms of demand were interleaved in the same environment (Experiment 2). These conditions excluded a number of potential explanatory candidates that were not modeled, such as learned industriousness, boredom, fatigue, or asymmetric learning of environmental richness (see Discussion).

Methods

Model Fitting and Model Comparison Procedures

For each model, both population-level and individual subject-level parameters were fit simultaneously using an Expectation-Maximization algorithm (Huys et al., 2011, 2012). Briefly, it was assumed that, at the population level, subject-level values for each parameter were distributed according to a Gaussian distribution, defined by a mean and standard deviation. Accordingly, point estimates for the individual subject parameters were found via maximum a posteriori (MAP) estimation using the population distribution for each parameter as the priors, and the variance in each individual subject parameter estimate was calculated by taking the inverse of the Hessian matrix at the MAP estimate. Next, the mean of the prior was updated by taking the weighted mean of individual subject estimates, and the variance of the prior was updated by taking the weighted variance across the parameter estimates of individual subjects. This process was repeated until the mean and variance of the prior for all parameters converged (a change of less than .1% across successive iterations). Full mathematical details of the algorithm can be found in Huys et al. (2011, 2012), and the algorithm was implemented using a custom written R package, GaussExMax, available on GitHub (github.com/gkane26/GaussExMax).

Model comparison was performed by calculating Bayesian Information Criterion at the group level (integrated BIC or iBIC, Huys et al., 2011, 2012; Kane et al., 2019). iBIC penalizes the likelihood of the data given model parameters, $p(d|\text{model})$, for model flexibility (the number of parameters, k), and the size of the penalty depends on the number of observations o :

$$iBIC = \log p(d|\text{model}) + \frac{k}{2} \log(o)$$

A Laplace approximation was used to find the log marginal likelihood, the likelihood of the data for each subject given the population level distribution over each parameter (Huys et al., 2011; Huys et al., 2012; Kane et al., 2019).

Model Definitions

We began by fitting a *base model* derived from the normative predictions of the Marginal Value Theorem (MVT, Charnov, 1976), which states that a delayed reward should be accepted if its magnitude surpasses the opportunity cost of time incurred in obtaining it (see Reward-maximizing Strategy for details). This evaluation rule can be expressed as

$$V_{\text{base}_i} = R_i - \gamma T_{\text{handling}_i} \quad (1)$$

such that the reward R of trial i is favorable if it is larger than the average amount one could earn per second in the environment γ (the opportunity cost) during the handling time T_{handling_i} , with high values of γ promoting selectivity. We fit this and all subsequent models using a softmax decision rule to determine the probability of accepting the offer on a given trial

$$P(\text{accept})_i = \frac{1}{1 + e^{-(1/\beta)(V_{\text{base}_i})}} \quad (2)$$

The value of γ in this base model was computed as an individual’s empirical rate of earnings by the time of each offer, with the free temperature parameter β indexing decision noise. Differences among the between-subject cost conditions in Experiment 1 could then be due to the impact of each cost on decision noise, which could have also driven the gradual convergence of acceptance rates in Experiment 2.

A second possibility is that individuals updated the estimate of gamma at different rates when faced with different demands. Previous accounts have demonstrated biases that can impact the learning of the opportunity cost in foraging paradigms (Dundon et al., 2020; Garrett & Daw, 2020). We therefore tested a second candidate model, an *adaptive model* that allowed γ to evolve over time, adapting a learning rule from Constantino & Daw (2015)

$$\gamma_{i+1} = (1 - (1 - \alpha)^{\tau_i}) \frac{R_i A_i}{\tau_i} + (1 - \alpha)^{\tau_i} \gamma_i \quad (3)$$

where A_i is an acceptance indicator (0 or 1), and τ_i is the duration of the trial. The duration was given by

$$\tau_i = T_{\text{handling}_i} A_i + T_{\text{travel}_i} \quad (4)$$

such that the handling time was counted only for accepted trials. Higher values of the learning rate α (ranging from 0 to 1) resulted in a larger update toward the current trial's rate of earnings weighted by its duration (first term in the equation). Free parameters for this adaptive model were the learning rate α and the initial value of γ in addition to the temperature β . These modifications allowed for the tracking of recent reward history, while maintaining the same evaluation rule

$$V_{\text{adaptive}_i} = R_i - \gamma_i T_{\text{handling}_i} \quad (5)$$

where values of γ were now updated on each trial i according to Eq. 3, rather than by the empirical rate of earnings.

A third possibility was that rather than impacting the learning rate, features specific to each cost condition additively biased the estimate of γ . Such bias could remain stable if only one type of demand was experienced (Experiment 1), but adjust over time as biases from other tasks were experienced (Experiment 2). For example, such biases could reflect an individual's global estimation of the richness of the environment, could reflect any direct utility or disutility of the experience during the handling time, or could reflect a focal estimation of the passage of time experienced during performance of each task (e.g. how long 10 seconds subjectively feels). Altered perceived time durations have occasionally been associated with levels of cognitive engagement (Csikszentmihalyi, 2014; Wearden, 2016), and would be expected to impact the subjective estimate of a trial's reward rate. Intuitively, the subjective duration of a fixed handling time could be recalibrated toward a consensus estimate as decision makers experienced different types of pre-reward demands. Our third candidate model, *biased model*, added a demand-specific term to γ that was gradually integrated across tasks for those participants who experienced multiple cost conditions. The overall model took the following form

$$V_{\text{biased}_i} = R_i - (\gamma + C_i)T_{\text{handling}_i} \quad (6)$$

where γ was the mean total rate of earnings per subject (mean = 0.55 cents per second, SD = 0.02). Negative values of C_i would therefore suggest a boost in offer value produced by a particular form of demand. On each trial i , C_i was the cost bias produced by each demand c_{demand} weighted against a global bias estimate c_{global}

$$C_i = c_{\text{global}}w_i + c_{\text{demand}}(1 - w_i) \quad (7)$$

with w_i controlling how early participants started integrating the consensus cost c_{global} into their decisions. The value of this weight was given by

$$w_i = \left(\frac{i}{\max(i)} \right)^{\alpha_w} \quad (8)$$

where the term in parentheses was the proportion of the experiment completed so far (note that the organization of the experimental session was fully disclosed), with high values of the free parameter α_w denoting a late integration of the global converging bias. Since participants in Experiment 1 experienced a single type of demand, the values of c_{demand} and c_{global} were the same and α_w did not play a role. This resulted in fitting only two free parameters, c_{demand} and β , for that data set. For Experiment 2, we fixed each value of c_{demand} (i.e. $c_{\text{cognitive}}$, c_{wait} , and c_{physical}) to the mean of the fitted parameters from each demand group from Experiment 1, and fit three free parameters: the global bias c_{global} , the convergence rate α_w , and temperature β . This dynamic allowed for differences in acceptance rates across cost conditions to remain stable in Experiment 1 but gradually collapse in Experiment 2.

Finally, we asked if the biases imposed by each form of demand could be isolated from a flexible computation of the opportunity cost of time. Our fourth candidate model, an *adaptive + biased model*, combined the main features of the adaptive model and the biased model, such that γ was estimated on each trial per Eq. 3, with demands imposing distinct and potentially converging biases as described in Eq. 6-8. The form of this fourth model was

$$V_{\text{adaptive} + \text{biased}_i} = R_i - (\gamma_i + C_i)T_{\text{handling}_i} \quad (9)$$

We fit four free parameters to Experiment 1 data (the temperature β , learning rate α , initial value of γ , and demand-specific bias c_{demand}) and five free parameters to Experiment 2 data (the temperature β , learning rate α , initial value of γ , consensus bias c_{global} , and the rate of bias integration α_w). The resulting parameters were compared across cost groups (in Experiment 1) and across experiments using ANOVAs, and the c_{demand} biases computed in Experiment 1 were compared against 0 using one-sample t-tests.

Results

The fourth model produced the lowest iBIC in both experiments (Fig. 9A; Experiment 1 iBIC = 3552.2; Experiment 2 iBIC = 1818.57), indicating that the tenets of the Marginal Value Theorem were necessary but not sufficient to explain the observed adaptive behavior. Fig. 9B-C show parameter fits from the winning model for both experiments. The learning rate α and the bias c_{demand} were of particular interest, as these parameters distinguished whether different forms of demand globally affected the integration of the environmental richness or directly modulated an offer's perceived cost, respectively. Analyses of variance showed significant differences among Experiment 1 groups only for c_{demand} ($F(2, 60) = 7.01, p = 0.002$), with α being similar for all groups. Pairwise post-hoc permutations confirmed that the Cognitive group (mean = -0.08, SD = 0.14) had lower values for the bias parameter than either the Wait group (mean = 0.11, SD = 0.18; $p = 0.001$; Cohen's D = 1.16) or the Physical group (mean = 0.09, SD = 0.21; $p = 0.004$; Cohen's D = 0.95). Parameter estimates were significantly greater than zero for the Wait group ($t(20) = 2.79, p = 0.011$), and marginally significantly lower than zero for the Cognitive group ($t(20) = -2.51, p = 0.021$).

Figure 9: Computational modeling results. A: iBIC values for four candidate models of choice behavior for both experiments. Lower values indicate better fit. B: Estimates of the bias parameter c -demand for participants in the Physical, Cognitive, and Wait cost conditions of Experiment 1. Negative values suggest the Cognitive group acted as if cognitive demands attenuated the perceived opportunity cost of the delay. C: Estimates of other model parameters indicated consistent foraging behavior across conditions and experiments. D: Synthetic behavior for Experiment 2 generated by the winning model. Error bars correspond to the empirical mean acceptance rate $\pm$ SEM per cost type, whereas lines and shaded areas show the model-generated mean acceptance rate and its SEM, respectively. The model's behavior recapitulated the initial apparent preference for cognitive effort and the eventual convergence of acceptance rates by the end of the session.

The mean values of c_{demand} from each group were carried over to the model of Experiment 2, and were aggregated over time according to a weighted decay into a final c_{global} value. As can be seen in Fig. 9C, other parameter values were comparable across experiments, pointing to consistency in the way decision makers approached our foraging scenarios. In addition, values of α_w suggested that c_{demand} estimates converged towards the second half of the experiment (mean = 1.75, SD = 2.07), gradually reducing bias by the end of the experimental session (c_{global} mean = -0.02, SD = 0.06). Model fits provided good qualitative resemblance to the observed choice dynamics in each experiment. This is illustrated in Fig. 9, in which 95% confidence intervals from within-subject fits are overlaid on empirical choices from Experiment 2.

Discussion

Understanding the sporadic attractiveness of effort remains elusive. In this study, we sought to identify factors underlying this phenomenon, and tested the hypothesis that the perceived cost or value of effort could depend on whether the decision environment contains opportunities to achieve the same outcome through different means. We conducted two foraging experiments that exposed individuals to a single or intermixed types of demands (cognitive effort, physical effort, or passive delay). We found a tendency toward engaging with cognitively effortful prospects when individuals faced a single form of demand, which faded over time when other costs were also experienced.

Our findings provided partial support for our pre-registered hypotheses (<https://osf.io/2rsgm/registrations>). In accord with our predictions, decision makers in both experiments successfully approximated reward-maximizing strategies, as evidenced by acceptance rates near optimal levels. This is in line with previous foraging experiments, where the main deviation from optimality has been a tendency to overharvest (Constantino & Daw, 2015; Wikenheiser et al., 2013). In contrast, our prediction that effort requirements would lead to greater acceptance rates only held for the Cognitive group, and this tendency did not reverse as predicted in Experiment 2. These hypotheses originally stemmed from prevalent proposals that incentivized work should become more appealing when presented in isolation (Eisenberger, 1992; Inzlicht et al., 2018), but that people would opt for the easiest way to achieve desired outcomes when multiple alternatives are available (Hull, 1943; Kool et al., 2010; Olivola & Shafir, 2013). Instead, our novel use of foraging paradigms, in addition to our inclusion of types of effort that are often studied separately, highlighted that reward-independent task characteristics can promote or inhibit the pursuit of effortful prospects.

The observed choice dynamics ruled out a number of candidate explanations. For instance, if the subjective costs and benefits of effort or delay were intrinsic, we would not have expected to observe the gradual convergence in acceptances when all three forms of demand were intermixed in the same environment.

This line of reasoning also speaks against a direct role of boredom in our results. The aversiveness of boredom can lead people to attempt any available activity in order to avoid it (Westgate, 2020; Wilson et al., 2014). In our paradigm, such a tendency would have been expected to lead to a persistent preference for either effortful condition. Lastly, the gradual convergence of acceptance rates across cost conditions rules out the possibility that our sample was mostly composed of demand seekers. Some individuals express a strong need for cognition, which can bias studies of mental labor (Cacioppo et al., 1996; Sayalı & Badre, 2019). Even though the prevalence of this trait in our college-level population could have driven the preferences for cognitive effort observed in both experiments, we would have expected such individuals to maintain this preference in the face of alternatives. Another unlikely candidate is fatigue. Even though participants decreased their acceptance rates for effort in both experiments, we would have expected participants in Experiment 2 to increase engagement with Wait trials, substituting Wait trials for effort trials in order to preserve their earnings. Our observations are unlikely to be due to the risk of failing to successfully complete effortful trials, as participants remained accurate throughout both experiments, and efficacy was consistent across cost conditions and experiments (Frömer et al., 2021). Finally, we cannot conclude that cognitive effort gained value from its rewarding outcomes, as has been proposed in the past (Eisenberger & Cameron, 1996). This phenomenon would not have been expected to benefit a single type of demand differentially because cognitive effort, physical effort, and unfilled delay were rewarded at the same rate.

Traditional foraging models, in which a reward is weighed against the earnings foregone by pursuing it (Charnov, 1976; Constantino & Daw, 2015), were also not sufficient to explain our results. Instead, our results were consistent with a computational model in which each type of demand distinctly biased the estimation of the opportunity cost, and such biases could be gradually recalibrated over time as decision makers gained experience with additional demands. In this formulation, passive delay and physical effort added to the opportunity cost, whereas cognitive effort reduced it. This raises the question of how these biases arise. We hypothesize that demands altered the perceived duration of the delay interval. Immersive activities can subjectively speed up the passage of time (Csikszentmihalyi, 2014), and this effect appears to scale with the amount of cognitive processing required by a task (Wearden, 2016). Moreover, adjusting for time sensitivity has been shown to improve fits for delay discounting models (McKerchar et al., 2009). Therefore, even though participants knew the nominal handling times, cognitive effort might have promoted higher acceptance rates by subjectively compressing the perceived durations. We speculate that this construal of elapsed time could then be adjusted with the experience of alternative demands, as participants learned how long the 10-second handling time felt in the context of various activities. This could help explain why preferences for cognitive effort were smaller for those participants in Experiment 2 who started with the physical effort block, as they had already begun to form an estimate of the handling time before they experienced their first cognitive

734 effort trial. The cognitive effort task might have shortened subjective durations for several reasons, either
735 because it divided the handling time into discrete, 2 s trials, or because taxing cognitive resources distracted
736 individuals from estimating the passage of time (Zakay & Tsal, 1989). These hypotheses will have to be
737 experimentally probed in the future.

738 Relying on foraging paradigms enabled us to address several gaps in the existing literature of effort-based
739 decision making. For example, even though engaging with a demanding offer encompasses the evaluation,
740 pursuit, and receipt of the outcome, investigations of the pursuit and receipt stages have been limited. Many
741 studies have sought to isolate the evaluation stage by presenting individuals with offers that are not realized
742 until the end of the experimental session, deliberately setting aside the question of how the ongoing experience
743 of cost and outcome affects subsequent choices (Apps et al., 2015; Chong et al., 2016; Kable & Glimcher,
744 2007; Massar et al., 2015). However, discrepancies between expected and experienced value are well known
745 (Kahneman & Thaler, 2006), and choosing immediate or low-effort outcomes could be due to factors other
746 than time preference, such as a perceived risk of being unable to redeem the future rewards, or hoping to find
747 better offers in the near future (Fawcett et al., 2012; Frederick et al., 2002; Sozou, 1998). We circumvented
748 these pitfalls by providing participants with an opportunity to actively experience the demands they accepted,
749 and giving them full knowledge of the available options. Another challenge in comparing effort and delay is
750 that increments in the required effort are often paired with a longer execution time. We avoided this issue by
751 imposing consistent handling times for all demands, allowing us manipulate temporal durations and demands
752 for effort independently. Finally, the existence of well-defined optimal strategies provided a useful reference
753 point for evaluating choice behavior. The apparent attractiveness of cognitive effort was reflected both by a
754 higher proportion of acceptances at the expense of earnings, and by parameter estimates capturing bias in
755 the estimation of opportunity costs.

756 A potential limitation of this work is that we examined only a single combination of cognitively effortful
757 tasks. Future work should investigate whether other ways of operationalizing cognitive and physical effort
758 would be more aversive or more engaging and would resultingly have different effects on decision making.
759 A potentially relevant factor is the way in which we calibrated cognitive and physical tasks. For physical
760 demands, we adjusted the gripping threshold to people’s maximum strength, whereas we trained participants
761 in the cognitive tasks until they reached criterion. Training might have improved response automaticity and
762 reduced error-related risk. In addition, attributes such as the efficacy to complete a task, or the diversity of
763 its components, have been reported to promote engagement (Bandura, 2010; Eisenberger & Cameron, 1996).
764 However, hand grip tasks have frequently been used to probe physical effort in previous work (Chong et al.,
765 2017; Hogan et al., 2018; Prevost et al., 2010), and task switching has been shown to impose subjective costs
766 (Kool et al., 2010; Shenhav et al., 2017). Furthermore, the observed decrease in acceptance rates over time

for both effort domains is consistent with patterns seen for other forms of effortful exertion (Treadway et al., 2009).

The present findings have important implications for everyday and clinical scenarios. Excessive aversion to delay and effort can manifest as impulsivity and apathy, respectively (Ainslie, 1975; Bonnelle et al., 2016; Chong et al., 2016; Cummings, 1993), fluctuations in which are hallmark symptoms of various mental illnesses and substance use disorders (Paulus, 2007). Lack of motivation to exert effort or wait for delayed prospects is a distinguishing factor of conditions such as depression and schizophrenia (Barch et al., 2019). In depression, medicated patients in remission are more sensitive to effort than control participants, and this sensitivity is predictive of relapse (Berwian et al., 2020). Encouraging patients to engage in activities has been found to be a helpful treatment for depression (Dimidjian et al., 2006). Tailored reward schedules can successfully achieve this (Eisenberger & Cameron, 1996). However, since depressed patients also show insensitivity to rewards (Huys et al., 2013), identifying what makes effort appealing in its own right can help address treatment-resistant cases. Our work suggests that voluntary effortful engagement can be encouraged by removing cognitively less-effortful paths to reward from the immediate decision context, and suggests it would be useful for future work to test directly for effects of effortful tasks on concurrent time estimation. These same principles can potentially be applied to increasing productivity at work or adherence to physical exercise routines.

Conclusion

In this study, we have examined the factors that motivate engagement with different behavioral demands, emphasizing how engagement can be influenced by features of the decision environment. Complementing work that has focused on the role of incentives and retrospective justifications (Eisenberger & Cameron, 1996; Inzlicht et al., 2018), our work highlights outcome-independent and task-specific features that can motivate the pursuit of effortful courses of action.

Acknowledgments

We thank Andrew Westbrook, the Shenhav lab, and the Daw lab for helpful discussions. This work was supported by grants from the National Science Foundation (BCS-1755757) and the Office of Naval Research (MURI award N00014-16-1-2832 and DURIP award N00014-17-1-2304).

References

- Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82(4), 463–496. <https://doi.org/10.1037/h0076860>
- Apps, M., Grima, L. L., Manohar, S., & Husain, M. (2015). The role of cognitive effort in subjective reward devaluation and risky decision-making. *Scientific Reports*, 5, 16880. <https://doi.org/10.1038/srep16880>
- Arulpragasam, A. R., Cooper, J. A., Nuutinen, M. R., & Treadway, M. T. (2018). Corticoinsular circuits encode subjective value expectation and violation for effortful goal-directed behavior. *Proceedings of the National Academy of Sciences of the United States of America*, 115(22), E5233–E5242. <https://doi.org/10.1073/pnas.1800444115>
- Bandura, A. (2010). Self-Efficacy. In *The corsini encyclopedia of psychology*. <https://doi.org/10.1002/9780470479216.corpsy0836>
- Barch, D. M., Pagliaccio, D., Luking, K., Moran, E. K., & Culbreth, A. J. (2019). Pathways to Motivational Impairments in Psychopathology: Common Versus Unique Elements Across Domains. In M. Neta & I. J. Haas (Eds.), *Emotion in the mind and body* (pp. 121–160). https://doi.org/10.1007/978-3-030-27473-3_5
- Bates, D., Mächler, M., Zurich, E., Bolker, B. M., & Walker, S. C. (2015). Fitting Linear Mixed-Effects Models Using lme4. *JSS Journal of Statistical Software*, 67. <https://doi.org/10.18637/jss.v067.i01>
- Berwian, I. M., Wenzel, J. G., Collins, A. G., Seifritz, E., Stephan, K. E., Walter, H., & Huys, Q. J. (2020). Computational Mechanisms of Effort and Reward Decisions in Patients with Depression and Their Association with Relapse after Antidepressant Discontinuation. *JAMA Psychiatry*, 77(5), 513–522. <https://doi.org/10.1001/jamapsychiatry.2019.4971>
- Białaszek, W., Marcowski, P., & Ostaszewski, P. (2017). Physical and cognitive effort discounting across different reward magnitudes: Tests of discounting models. *PLoS ONE*, 12(7), 1–25. <https://doi.org/10.1371/journal.pone.0182353>
- Bonnelle, V., Manohar, S., Behrens, T., & Husain, M. (2016). Individual Differences in Premotor Brain Systems Underlie Behavioral Apathy. *Cerebral Cortex*, 26(2), 807–819. <https://doi.org/10.1093/cercor/bhv247>
- Cacioppo, J. T., Petty, R. E., Feinstein, J. A., & Jarvis, W. B. G. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, 119(2), 197–253. <https://doi.org/10.1037/0033-2909.119.2.197>
- Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2), 129–136. [https://doi.org/10.1016/0040-5809\(76\)90040-X](https://doi.org/10.1016/0040-5809(76)90040-X)
- Chong, T.-J., Bonnelle, V., & Husain, M. (2016). Quantifying motivation with effort-based decision-making paradigms in health and disease. In *Progress in brain research* (1st ed., Vol. 229, pp. 71–100).

<https://doi.org/10.1016/bs.pbr.2016.05.002>

Chong, T. T., Apps, M., Giehl, K., Sillence, A., Grima, L. L., & Husain, M. (2017). Neurocomputational mechanisms underlying subjective valuation of effort costs. *PLoS Biology*, 15(2), 1–28. <https://doi.org/10.1371/journal.pbio.1002598>

Constantino, S. M., & Daw, N. D. (2015). Learning the opportunity cost of time in a patch-foraging task. *Cognitive, Affective & Behavioral Neuroscience*, 15(4), 837–853. <https://doi.org/10.3758/s13415-015-0350-y>

Csikszentmihalyi, M. (2014). *Flow and the foundations of positive psychology*. Springer.

Cummings, J. L. (1993). Frontal-Subcortical Circuits and Human Behavior. *Archives of Neurology*, 50, 873–880.

Dimidjian, S., Hollon, S. D., Dobson, K. S., Schmaling, K. B., Kohlenberg, R. J., Addis, M. E., ... Jacobson, N. S. (2006). Randomized trial of behavioral activation, cognitive therapy, and antidepressant medication in the acute treatment of adults with major depression. *Journal of Consulting and Clinical Psychology*, 74(4), 658–670. <https://doi.org/10.1037/0022-006X.74.4.658>

Dundon, N. M., Garrett, N., Babenko, V., Cieslak, M., Daw, N. D., & Grafton, S. T. (2020). Sympathetic involvement in time-constrained sequential foraging. *Cognitive, Affective and Behavioral Neuroscience*. <https://doi.org/https://doi.org/10.3758/s13415-020-00799-0>

Eisenberger, R. (1992). Learned industriousness. *Psychological Review*, 99(2), 248–267. <https://doi.org/10.1037/0033-295X.99.2.248>

Eisenberger, R., & Cameron, J. (1996). Detrimental Effects of Reward - Reality or myth? *The American Psychologist*, 51(11), 1153–1166. <https://doi.org/10.1037/0003-066X.51.11.1153>

Fawcett, T. W., McNamara, J. M., & Houston, A. I. (2012). When is it adaptive to be patient? A general framework for evaluating delayed rewards. *Behavioural Processes*, 89(2), 128–136. <https://doi.org/10.1016/j.beproc.2011.08.015>

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2), 351–401.

Frömer, R., Lin, H., Wolf, C. D., Inzlicht, M., & Shenhav, A. (2021). Expectations of reward and efficacy guide cognitive control allocation. *Nature Communications*, 12(1), 1–11.

Garrett, N., & Daw, N. D. (2020). Biased belief updating and suboptimal choice in foraging decisions. *Nature Communications*, 11(1), 1–12. <https://doi.org/10.1038/s41467-020-16964-5>

Green, L., Fristoe, N., & Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review*, 1(3), 382–389. <https://doi.org/10.3758/BF03213979>

Hartmann, M. N., Hager, O. M., Tobler, P. N., & Kaiser, S. (2013). Parabolic discounting of monetary

rewards by physical effort. *Behavioural Processes*, 100(October 2013), 192–196. <https://doi.org/10.1016/j.beproc.2013.09.014>

Hayden, B. Y. (2018). Economic choice: the foraging perspective. *Current Opinion in Behavioral Sciences*, 24, 1–6. <https://doi.org/10.1016/j.cobeha.2017.12.002>

Hernandez Lallement, J., Kuss, K., Trautner, P., Weber, B., Falk, A., & Fliessbach, K. (2014). Effort increases sensitivity to reward and loss magnitude in the human brain. *Social Cognitive and Affective Neuroscience*, 9(3), 342–349. <https://doi.org/10.1093/scan/nss147>

Hogan, P. S., Galaro, J. K., & Chib, V. S. (2018). Roles of Ventromedial Prefrontal Cortex and Anterior Cingulate in Subjective Valuation of Prospective Effort. *Cerebral Cortex*, 1–14. <https://doi.org/10.1093/cercor/bhy310>

Hull, C. L. (1943). *Principles of Behavior* (R. M. Elliott, Ed.). New York, NY: Appleton-Century.

Huys, Q. J., Cools, R., Gölzer, M., Friedel, E., Heinz, A., Dolan, R. J., & Dayan, P. (2011). Disentangling the roles of approach, activation and valence in instrumental and pavlovian responding. *PLoS Computational Biology*, 7(4). <https://doi.org/10.1371/journal.pcbi.1002028>

Huys, Q. J., Eshel, N., O’Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: How the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, 8(3). <https://doi.org/10.1371/journal.pcbi.1002410>

Huys, Q. J., Pizzagalli, D. A., Bogdan, R., & Dayan, P. (2013). Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. *Biology of Mood & Anxiety Disorders*, 3(1), 1–16. <https://doi.org/10.1186/2045-5380-3-12>

Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The Effort Paradox: Effort Is Both Costly and Valued. *Trends in Cognitive Sciences*, 22(4), 337–349. <https://doi.org/10.1016/j.tics.2018.01.007>

Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, 10(12), 1625–1633.

Kacelnik, A., & Marsh, B. (2002). Cost can increase preference in starlings. *ANIMAL BEHAVIOUR*, 63, 245–250. <https://doi.org/10.1006/anbe.2001.1900>

Kahneman, D., & Thaler, R. H. (2006). Anomalies: Utility maximization and experienced utility. *Journal of Economic Perspectives*, 20(1), 221–234. <https://doi.org/10.1257/089533006776526076>

Kane, G. A., Bornstein, A. M., Shenhav, A., Wilson, R. C., Daw, N. D., & Cohen, J. D. (2019). Rats exhibit similar biases in foraging and intertemporal choice tasks. *eLife*, 8, 1–22. <https://doi.org/10.7554/elife.48429>

Klein-Flugge, M. C., Kennerley, S. W., Friston, K., & Bestmann, S. (2016). Neural Signatures of Value Comparison in Human Cingulate Cortex during Decisions Requiring an Effort-Reward Trade-off. *Journal of Neuroscience*, 36(39), 10002–10015. <https://doi.org/10.1523/JNEUROSCI.0292-16.2016>

- Klein-Flügge, M. C., Kennerley, S. W., Saraiva, A. C., Penny, W. D., & Bestmann, S. (2015). Behavioral Modeling of Human Choices Reveals Dissociable Effects of Physical Effort and Temporal Delay on Reward Devaluation. *PLOS Computational Biology*, 11(3), e1004116. <https://doi.org/10.1371/journal.pcbi.1004116>
- Kool, W., & Botvinick, M. (2018). Mental labour. *Nature Human Behaviour*, 2(December), 899–908. <https://doi.org/10.1038/s41562-018-0401-9>
- Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision Making and the Avoidance of Cognitive Demand. *Journal of Experimental Psychology: General*, 139(4), 665–682. <https://doi.org/10.1037/a0020198>
- Krebs, B. Y. J. R., Erichsen, J. T., & Webber, M. I. (1977). Optimal prey selection in the great tit (*Parus major*). *Animal Behavior*, 25(2), 30–38.
- 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(06), 661–679. <https://doi.org/10.1017/S0140525X12003196>
- Massar, S. A. A., Libedinsky, C., Weiyan, C., Huettel, S. A., & Chee, M. W. L. (2015). Separate and overlapping brain areas encode subjective value during delay and effort discounting. *NeuroImage*, 120, 104–113. <https://doi.org/10.1016/j.neuroimage.2015.06.080>
- McGuire, J. T., & Kable, J. W. (2012). Decision makers calibrate behavioral persistence on the basis of time-interval experience. *Cognition*, 124(2), 216–226. <https://doi.org/10.1016/j.cognition.2012.03.008>
- McGuire, J. T., & Kable, J. W. (2015). Medial prefrontal cortical activity reflects dynamic re-evaluation during voluntary persistence. *Nature Neuroscience*, 18(5), 760–769.
- Mckelchar, T. L., Green, L., Myerson, J., Pickford, T. S., Hill, J. C., & Stout, S. C. (2009). A comparison of four models of delay discounting in humans. *Behavioural Processes*, 81, 256–259. <https://doi.org/10.1016/j.beproc.2008.12.017>
- Mischel, W., & Ebbesen, E. B. (1970). Attention in delay of gratification. *Journal of Personality and Social Psychology*, 16(2), 329–337. <https://doi.org/10.1037/h0029815>
- Mobbs, D., Trimmer, P. C., Blumstein, D. T., & Dayan, P. (2018). Foraging for foundations in decision neuroscience: insights from ethology. *Nature Reviews Neuroscience*, 1. <https://doi.org/10.1038/s41583-018-0010-7>
- Mochon, D., Norton, M. I., & Ariely, D. (2012). Bolstering and restoring feelings of competence via the IKEA effect. *International Journal of Research in Marketing*, 29(4), 363–369. <https://doi.org/10.1016/j.ijresmar.2012.05.001>
- Olivola, C. Y., & Shafir, E. (2013). The Martyrdom Effect: When Pain and Effort Increase Prosocial Contributions. *Journal of Behavioral Decision Making*, 26(1), 91–105. <https://doi.org/10.1002/bdm.767>

- Otto, A. R., & Daw, N. D. (2019). The opportunity cost of time modulates cognitive effort. *Neuropsychologia*, 123(May 2018), 92–105. <https://doi.org/10.1016/j.neuropsychologia.2018.05.006>
- Paulus, M. P. (2007). Decision-making dysfunctions in psychiatry - Altered homeostatic processing? *Science*, 318(5850), 602–606. <https://doi.org/10.1126/science.1142997>
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., ... Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195–203. <https://doi.org/10.3758/s13428-018-01193-y>
- Prevost, C., Pessiglione, M., Metereau, E., Clery-Melin, M.-L., & Dreher, J.-C. (2010). Separate Valuation Subsystems for Delay and Effort Decision Costs. *Journal of Neuroscience*, 30(42), 14080–14090. <https://doi.org/10.1523/JNEUROSCI.2752-10.2010>
- R Core Computing Team. (2017). *R: A Language and Environment for Statistical Computing*. Retrieved from <https://www.r-project.org/>
- Sayali, C., & Badre, D. (2019). Neural systems of cognitive demand avoidance. *Neuropsychologia*, 123(June 2018), 41–54. <https://doi.org/10.1016/j.neuropsychologia.2018.06.016>
- Seaman, K. L., Brooks, N., Karrer, T. M., Castrellon, J. J., Perkins, S. F., Dang, L. C., ... Samanez-Larkin, G. R. (2018). Subjective value representations during effort, probability and time discounting across adulthood. *Social Cognitive and Affective Neuroscience*, 13(5), 449–459. <https://doi.org/10.1093/scan/nsy021>
- 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), annurev-neuro-072116-031526. <https://doi.org/10.1146/annurev-neuro-072116-031526>
- Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society B: Biological Sciences*, 265(1409), 2015–2020. <https://doi.org/10.1098/rspb.1998.0534>
- Stephens, D. W., & Krebs, J. W. (1986). *Foraging Theory*. Princeton University Press.
- Treadway, M. T., Buckholtz, J. W., Schwartzman, A. N., Lambert, W. E., & Zald, D. H. (2009). Worth the 'EEfRT'? The effort expenditure for rewards task as an objective measure of motivation and anhedonia. *PLoS ONE*, 4(8), 1–9. <https://doi.org/10.1371/journal.pone.0006598>
- Walton, M., Kennerley, S., Bannerman, D., Phillips, P., & Rushworth, M. (2006). Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, 19(8), 1302–1314. <https://doi.org/10.1016/j.neunet.2006.03.005>
- Wearden, J. (2016). *The psychology of time perception* (pp. 1–261). <https://doi.org/10.1057/978-1-137-40883-9>
- Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive, Affective, & Behavioral Neuroscience*, 15(2), 395–415. <https://doi.org/10.3758/s13415-015-0334-y>

- Westbrook, A., Kester, D., & Braver, T. S. (2013). What Is the Subjective Cost of Cognitive Effort? Load, Trait, and Aging Effects Revealed by Economic Preference. *PLoS ONE*, 8(7). <https://doi.org/10.1371/journal.pone.0068210>
- Westgate, E. C. (2020). Why Boredom Is Interesting. *Current Directions in Psychological Science*, 29(1), 33–40. <https://doi.org/10.1177/0963721419884309>
- Wikenheiser, A. M., Stephens, D. W., & Redish, A. D. (2013). Subjective costs drive overly patient foraging strategies in rats on an intertemporal foraging task. *Proceedings of the National Academy of Sciences*, 110(20), 8308–8313. <https://doi.org/10.1073/pnas.1220738110>
- Wilson, T. D., Reinhard, D. A., Westgate, E. C., Gilbert, D. T., Ellerbeck, N., Hahn, C., ... Shaked, A. (2014). Just think: The challenges of the disengaged mind. *Science*, 345(6192), 75–77. <https://doi.org/10.1126/science.1250830>
- Zakay, D. A. N., & Tsal, Y. (1989). Awareness of attention allocation and time estimation accuracy. *Bulletin of the Psychonomic Society*, 27(3), 209–210.