Effort expenditure links control of vigor with decision-making

Shruthi Sukumar[1], Reza Shadmehr[2], Alaa A. Ahmed[1]
University of Colorado Boulder[1] & Johns Hopkins University[2]

Introduction: Imagine you reach for a pencil to write down a brand-new research idea after just finding out one of your papers was accepted to a journal. Now re-imagine this scenario, but instead, you had been notified your paper was rejected. How would the reach for that pencil be different? Would one be faster than the other? To answer this question, we sought to determine how changes in prior rate of reward due to past actions affected the vigor of reaching movements. We began with a normative model for movement vigor proposed by Yoon et al. (2018) that builds on a classical ecological framework called the Marginal Value Theorem (MVT) proposed by E. Charnov (1976). In the classical MVT, Charnov prescribed a simple rule for distributing harvest time in patches—the forager should stop harvesting and consequently leave a patch when the marginal capture rate of the patch falls below the average capture rate of the environment (Fig 1a). Yoon et al. augmented this model to similarly prescribe the optimal movement duration between patches as the duration at which the rate of movement expenditure equals the capture rate of the environment in magnitude (Fig 1b). Therefore, the essence of the generalized MVT is that when foraging, actions—harvest or travel—and their durations are selected to maximize the average rate of capture or the overall quality of an environment. In a higher quality environment, like one in which you found out your paper was accepted for publication, shorter durations of actions are to be selected to maintain high capture rate; harvests will have shorter durations and, more importantly for us, movements will exhibit higher vigor. A slower than optimal movement duration carries an inherent opportunity cost, as it represents time better spent on more rewarding actions. We sought to modulate the capture rate of the environment by changing the travel effort of reaching movements. Our primary goal was to examine whether changes in history of effort affect the choice of reach vigor.

Methods: We developed a reach-based foraging protocol in which human subjects performed arm reaches as the means of travelling between patches. Once they arrived at a patch, they attained reward (points that were exchanged for money) by producing force on the handle grip (Fig 2a). To harvest reward, subjects moved the cursor into a red patch (Fig 2b) wherein they had to increase their grip force to a threshold ($F_g = 30N$; Fig 2c). Once this minimum grip force was attained, they were provided with reward at a rate that declined the longer they stayed. They could leave the current patch and move on to the next one at any time (Fig 2d–f). To vary effort expenditure of travel, the robot simulated various masses in two environments. In one environment, on each trial the subjects reached against a low effort mass (0kg). In another environment, they reached against a high effort mass (3.5kg). To determine the effect of effort history on movement vigor, we inserted probe trials in both environments. In probe trials, mass was always 2kg (Fig 2g). We asked whether modulation of effort history affected reach vigor, as well as harvest duration. Theory predicts that increased effort expenditure would reduce the global capture rate, leading to reduced reach vigor, and increased duration of harvest.

Results: Indeed, subjects performed faster reaches in probe trials corresponding to the low effort environment as opposed to those in the high effort environment, as MVT predicted (Fig 3). Peak velocity was higher in probe trials belonging to the low effort environments (Fig 3a, b, d; $p<0.001$), and travel duration was correspondingly lower (Fig 3c; $p<0.001$). Additionally, we did not observe subjects significantly increase or decrease their peak velocity (Fig 3e; $p = 0.647$) or travel duration ($p = 0.289$) over the course of any given block of ten probe trials. Therefore, subjects moved faster between patches in probe trials following a history of low effort trials, compared to probe trials following a history of high effort trials.

Our data also showed history of effort altering harvest behavior. Subjects were faster at ramping up their grip force as they began to harvest in probe trials belonging to the low effort environment as compared to the high effort environment. This is seen in how quickly they ramped-up their force (peak force rate, Fig 4a, $p<0.001$) and the duration of this ramp-up period (harvest reaction time, Fig 4b, $p=0.017$). We did not see any significant modulation in the number of berries harvested between the two environments ($p =0.240$). Finally, we computed the “giving-up” duration as the time between the last berry collected and when they decided to stop waiting for an additional berry and drop their force below the threshold. We saw that this giving-up duration, normalized to the duration between the last two berries, was slightly longer in the high effort environment as compared to the low effort environment (Fig 4c; $p =0.0105$).

In summary, both movement vigor, and harvest decisions, were influenced by effort history. Following a history of high effort, subjects reduced their reach vigor, took longer to begin harvesting reward, and waited longer to give up harvesting. This suggests that movement vigor and decision-making are linked via the history of effort expenditure.

Figure 1: Generalized MVT Predictions

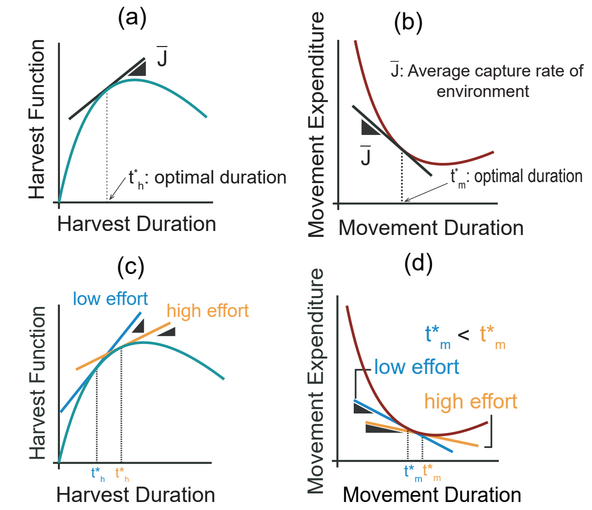

Figure 1: Prediction for (a) harvest duration and (b) movement duration of the generalized marginal value theorem. $\bar{J}$: average capture rate of the environment. Optimal harvest duration t_h* is the point at which the slope of the harvest intake curve equals $\bar{J}$. Optimal movement duration t_m* is the point at which the magnitude of the slope of the movement expenditure curve equals $\bar{J}$. (c) After experiencing a history of high effort in an environment (lower $\bar{J}$), the harvest duration in the current patch is smaller. (d) After a history of high effort, the movement duration selected for the current movement is higher.

Figure 2: Experimental protocol

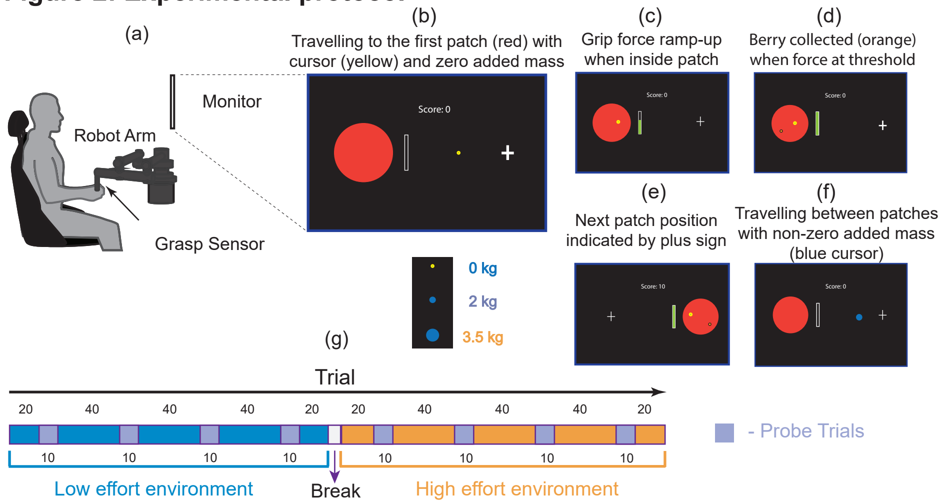

Figure 2: (a) Subjects were seated in front of a computer monitor while grasping the end of a robotic manipulandum with a grasp sensor.(b) Subjects are cued a red 'patch' into which they need to move the cursor so they can collect reward. (c) Once in the patch they can start increasing grip force to begin harvesting berries to the required force. (d) Once required force is applied, berries are harvested with an audio-visual stimulus indicating decreasing berry rate; for each berry subjects see, an orange circle and hear a high-pitched beep. Force had to be maintained above threshold during berry collection. (e) Subjects could move to the next cued location of the patch at any time to harvest reward; cumulative score is displayed at all times. (f) Depending on the environment, subjects experienced added mass while moving between patches indicated by a blue solid circle obscuring the cursor whose size indicated how heavy the arm was going to be. (g) Experimental protocol for the two environments as experimental blocks with number of trials indicated.

Figure 3: History of effort modulated movement vigor

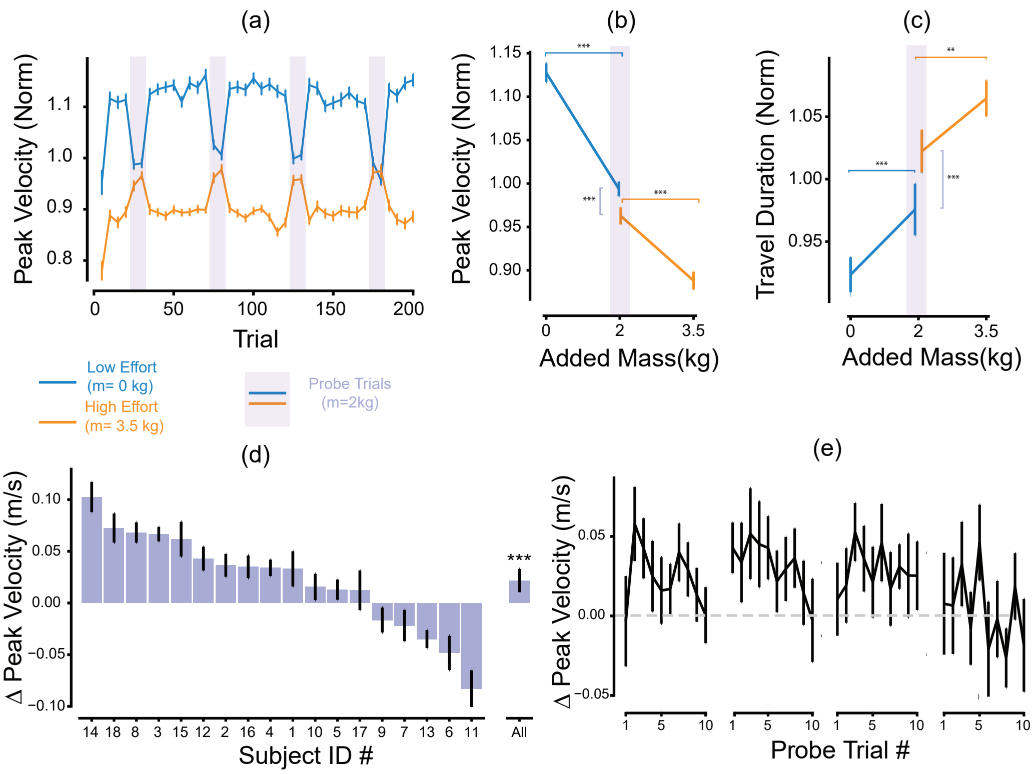

Figure 3: (a) Peak velocity, normalized to subject average, plotted as a function of trial for the low (blue) and high (yellow) effort environments. Data for five consecutive trials is binned. Purple regions are probe trials (m=2kg). (b,c) Normalized Peak velocity(b) and Travel Duration(c) plotted with respect to added mass on the current trial. (d) Average difference in peak velocity ΔPV between corresponding probe trials in low effort environment and high effort environment (Low - High) for each subject. Most subjects move on average faster in the probe trials belonging to the low effort environment. (e) ΔPV for probe trials plotted with respect to probe trial number. Peak velocity of probes in the low effort environments was higher than that in the high effort for most probe trials on average. The gray dashed line represents zero difference; each sub-block of ten probe trials is represented separately but in order.

Figure 4: History of effort slows harvest actions

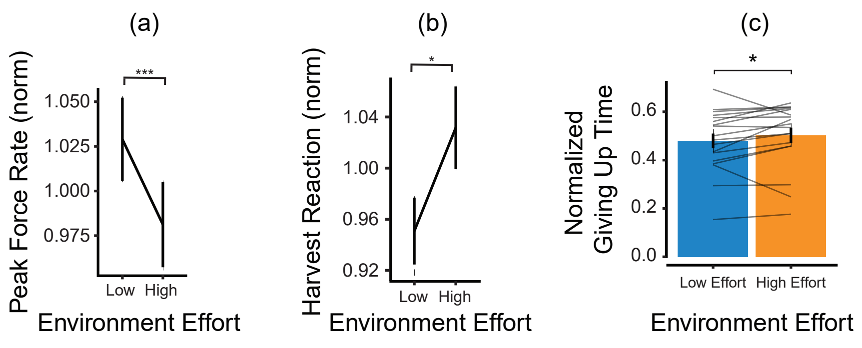

Figure 4: (a) Peak force rate – the maximum rate of force generation during ramp-up –normalized to subject averages is lower in probe trials belonging to the high effort environment. (b) Harvest reaction time – the duration from patch entry upto reaching the grip threshold – correspondingly increases in the probe trials belonging to the high effort environment as opposed to the low effort environment. (c) Giving-up time (normalized to the wait time between the last two berries) increased in the high effort environment.